# $\mathcal{ECHOS}$ -Enhanced Capacity 802.11 Hotspots

Arunchandar Vasan
Department of Computer Science
University of Maryland, College Park
Email: arun@cs.umd.edu

Ramachandran Ramjee, Thomas Woo
Bell Labs, Lucent Technologies
Email: ramjee.woo@bell-labs.com

***Abstract*— The total number of hotspot users around the world is expected to grow from 9.3 million at the end of 2003 to 30 million by the end of 2004 according to researcher Gartner. Given the explosive growth in hotspot wireless usage, enhancing capacity of 802.11-based hotspot wireless networks is an important problem. In this paper, we make two important contributions. We first present the AP-CST algorithm that dynamically adjusts the Carrier Sense Threshold (CST) in order to allow more flows to co-exist in current 802.11 architectures. We then extend the current hotspot engineering paradigm by allowing every cell and AP access to all available channels. These cells are then managed by the RNC-SC algorithm running in a centralized Radio Network Controller. This algorithm assigns mobile stations to appropriate cells/channels and adjusts transmit power values dynamically, thereby exploiting spatial heterogeneity in distribution of users at the hotspots. Through detailed and extensive simulations, we show that the performance of 802.11-based hotspots can be improved by up to 195% per-cell and 70% overall.**

## I. INTRODUCTION

The total number of hotspot users around the world is expected to grow from 9.3 million at the end of 2003 to 30 million by the end of 2004 according to researcher Gartner [19]. Given the explosive growth in hotspot wireless usage, enhancing capacity of 802.11-based hotspot wireless networks is an important problem. Improving performance adaptively is especially critical since such hotspots are typically characterized by unpredictable load with a large number of users accessing network connectivity in a relatively small physical area.

Hotspot deployments typically operate in infrastructure mode where an Access Point (AP) provides connectivity to multiple mobile clients. Given the widespread adoption of 802.11-standard as is, we would like a solution that does not modify the standard. While other researchers have examined throughput enhancements in 802.11 infrastructure mode [4], [3], they assume that exactly one channel is available for optimization in each *cell*, i.e., an area covered by an AP(s), and operate within the conservative 802.11 floor reservation and carrier sensing mechanism [15].

In this paper, we introduce the $\mathcal{ECHOS}$ architecture to exploit the spatial heterogeneity of users and flows in order to improve 802.11 capacity in hotspots. In $\mathcal{ECHOS}$, we devise two algorithms, AP-CST and RNC-SC, that can improve 802.11 hotspot performance substantially.

AP-CST (Access Point (modifies) Carrier Sense Threshold) allows for multiple flows to co-exist in the same channel by dynamically modifying the carrier sensing threshold (CST). In 802.11, if the signal strength on the medium is higher than CST, the medium is termed busy, and stations delay their transmissions. Our approach is to allow this threshold to be programmable. Clearly, only the minimum strength at which a signal can be sensed is limited by the hardware. Increasing the threshold does not require changes to the hardware. An optimum carrier sense threshold allows flows which do not interfere to co-exist, i.e., transmit concurrently. As we shall see in Section III, typical settings for the static CST used today are very conservative (i.e. it precludes the parallel operation of flows even when they would not interfere with each other) and in many cases, supersedes the virtual carrier sensing using RTS/CTS. The AP-CST algorithm allows the AP to set its CST and those of its clients appropriately such that more flows can co-exist in the same channel without interference where possible. This solution addresses situations in a hotspot where neighboring cells are assigned the same channels due to the limited number of orthogonal channels (3) available in 802.11b/g. Furthermore, this solution can be implemented in practice at no extra hardware cost.

In RNC-SC (Radio Network Controller (uses) Secondary Channels) mode, we allow *each cell or AP access to all available channels*[1]. The RNC-SC algorithm executes in a centralized *Radio Network Controller* (RNC). In this algorithm, each cell uses one *primary* channel as done currently as well as uses the other two channels as *secondary* channels without causing co-channel interference. The RNC-SC algorithm computes the CSTs and transmit powers such that the performance of stations in the secondary channel improves without degrading the performance of stations in their primary channel. While RNC-SC results in extra hardware cost for the extra interfaces in each of the APs, the flexibility to dynamically use secondary channels allows it to improve hotspot performance beyond AP-CST, especially where there is heterogeneity in user distributions.

Note that implementing these algorithms at the mobile station does not require equipping each station with multiple 802.11 cards; it is possible to have just one 802.11 card and use a software-based solution such as MultiNet [8] to dynamically scan multiple channels and change associations. We only assume that the 802.11-based equipment is programmable through software such that certain parameters (CST and transmit power) can be dynamically adjusted. Interestingly, the trend for next generation wireless cards is to implement just

[1] In reality, the AP could be a single device with multiple interfaces; such devices are commercially available already, e.g. from [5]

0-7803-8968-9/05/$20.00 (C)2005 IEEE

the basic time-critical MAC functions in hardware while leaving their control and configuration to the operating system [8].

In summary, *the ability to dynamically allocate channels to stations and flexibly adapt parameters such as CST and/or transmit power provides us with a very powerful mechanism to adapt to heterogeneity and thereby improve performance.* $\mathcal{ECHOS}$ architecture is inspired by the design of wide-area CDMA-based wireless networks [21], where all cells operate using the same frequency band and a centralized radio network controller manages transmit power at the mobile station in order to minimize interference and maximize performance. Such a trend is also seen in 802.11 based design like in [9] . To the best of our knowledge, no one has so far proposed a multi-channel solution with dynamic CST in order to improve performance in 802.11 networks.

In this paper, we make two important contributions. We first present the AP-CST algorithm that dynamically adjusts the Carrier Sense Threshold (CST) in order to allow more flows to co-exist in current 802.11 architectures. We then extend the current hotspot engineering paradigm by allowing every cell and AP access to all available channels (three orthogonal channels in the case of 802.11b standard, the focus of this paper). These cells are then managed by the RNC-SC algorithm running in a centralized Radio Network Controller. This algorithm assigns mobile stations to appropriate cells/channels and adjusts transmit power values dynamically, thereby exploiting spatial heterogeneity in distribution of users at the hotspots. Through detailed and extensive simulations, we show that the performance of 802.11-based hotspots can be improved by both these algorithms by up to 195% per-cell and 70% overall.

The rest of the paper is organized as follows. We survey related work in Section II. In Section III, we present some definitions, motivate the case for a dynamic CST, and identify the optimum CST to be chosen by a transmitter. Section IV presents our proposed architecture and the algorithms involved. We evaluate the performance of the proposed architecture in Section V. Finally, we discuss enhancements and limitations of our approach in VI and conclude in Section VII.

## II. Related Work

Related work falls into two broad categories: extensions of the 802.11 MAC protocol to enhance throughput/fairness (primarily, in ad-hoc mode) and better planning and load balancing algorithms to enhance throughput in 802.11 infrastructure mode. We briefly review some of these contributions below.

In the area of 802.11 MAC enhancements, the basic observation that many researchers make is that the current 802.11 floor reservation and carrier sensing mechanism is conservative [15]. Thus, different techniques have been proposed to enhance the parallelism in transmissions between different sessions, primarily in mobile nodes using 802.11 to form ad-hoc networks [1], [6], [7], [10], [22] All these involve either modifications to the 802.11 MAC protocol or using out-of-band tones and thus, cannot be used to enhance the performance of the hundreds of millions of already deployed 802.11 cards and access points. Very recently, references [23] and [12] have also pointed out that varying CST can help boost performance. They approach the problem analytically in the case of ad-hoc networks.

While the use of multiple channels for throughput enhancement in 802.11 networks has been proposed in the context of ad hoc multi-hop wireless networks [16], [11], [20], given the dynamic and decentralized nature of wireless ad hoc networks, these solutions rely on each node making decisions based on its locally perceived medium characteristics and there is little scope for centralized coordination.

In the area of throughput enhancements of 802.11 infrastructure mode, the authors of [3] propose two enhancements to the Access Point (AP) association algorithms. In [4], the authors perform centralized coordination of APs in 802.11 PCF mode by allocating channels and time slots to APs (through graph coloring and centralized scheduling) to support fairness and QoS guarantees. However, these and most work in this area assume that each AP is capable of using only a single channel at a give time. In this paper, we take a completely different approach by allowing all APs to operate in all three channels as long as interference is not significant. The trend in the 802.11 equipment industry is also to introduce a centralized controller or switch [2], [17]. However, they do not advocate the use of multiple channels in each AP and the algorithms in their controller likely include proprietary modifications to the 802.11 standard at the AP.

## III. Observations on Carrier Sensing in 802.11

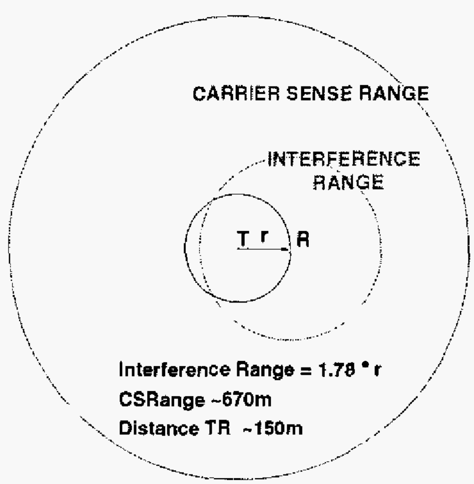


Fig. 1. Three ranges in wireless communication

In this section, we first make some observations about existing 802.11 carrier sensing. Then, we provide intuition about the optimum value of CST to be used by a transmitter. All simulation experiments in this paper were performed using the Qualnet simulator [18], with an accurate bit-error based model for the physical layer.

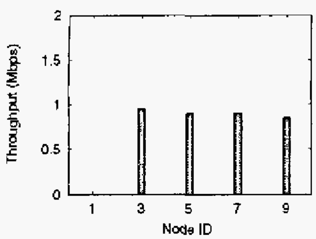


Fig. 2. RTS/CTS, -93dBm

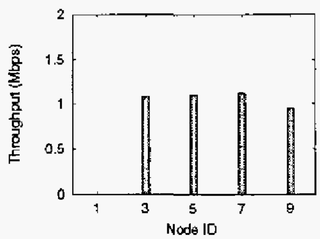


Fig. 3. No RTS/CTS, -93dBm

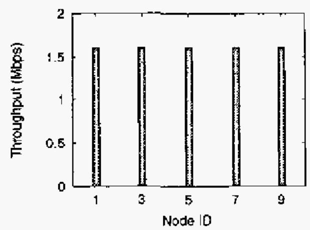


Fig. 4. RTS/CTS, -83dBm

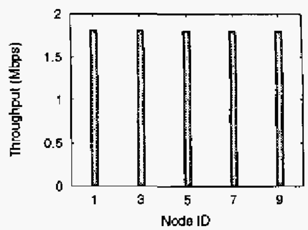


Fig. 5. No RTS/CTS, -83dBm

**Terminology** Let $T$ denote a transmitter and $R$ denote a receiver. We define the following ranges in 802.11:

- **Carrier Sense Range:** The region of space around $T$ in which the received strength of $T$'s transmission is greater than the Carrier Sense Threshold (CST) of the receiver.
- **Transmission Range:** The region of space around $T$ in which the received strength of $T$'s transmission is sufficient for successful reception to occur.
- **Interference Range:** The region of space around $R$, where any transmission by another source $T'$ would interfere with the frame transmitted by $T$ and cause its loss.

Suppose $T$ and $T'$ are two transmitters at distances $d_T$ and $d_I$, respectively from the receiver. $T'$ is an interferer to the transmission from $T$. Then the Signal to Noise Ratio (SNR) at the receiver is given by $d_I^4/d_T^4$ under the assumptions that both the transmitters transmit with the same power and that strength of the received signal falls off as $kP_0/d^4$, where $k$ is a suitable constant, $P_0$ is the transmission power, and $d$ is the distance from the signal source. For successful reception, the requirement is that the SNR be above a threshold ($\gamma$) which is typically 10 (10 dB). This yields the requirement:

$$\frac{d_I^4}{d_T^4} > \gamma$$

or $d_I > 1.78 d_T$. Therefore, for interference to occur, the required condition is that $d_I < 1.78 d_T$. Figure 1, drawn to scale, illustrates these terms. The relative distances shown in the figure have been obtained from the Qualnet simulator [18] for 802.11 transmission at 2Mbps with a CST of -93 dBm and transmit power of 15 dBm.

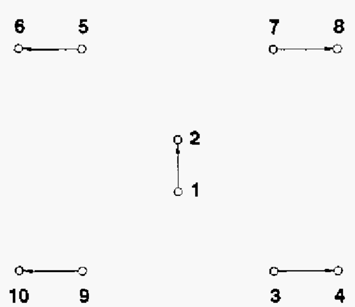


Fig. 6. RTS/CTS and/or default static CST resulting in low throughputs and severe unfairness.

**Observation I:** *The use of physical carrier sensing with the default fixed carrier sensing threshold can unnecessarily couple together several flows reducing per-flow and aggregate throughput.*

By coupling of flows, we mean that the flows sense each other's carrier and share the wireless medium. Figure 6 shows the topology considered. Nodes 5 and 7 sense only each other's carriers, 3 and 9 sense only each other, while 1 hears all transmissions. All transmitting stations are outside the interference ranges of all receiving stations. The workload considered is CBR traffic on all transmitters at rates sufficient to saturate the medium. We expect each transmitter in the pairs which sense each other (5 and 7, 3 and 9) to achieve half the maximum capacity, while node 1 is completely starved (as at any point of time, one of the other four transmissions go on). Figure 2 shows the throughput obtained with the use of RTS/CTS and Figure 3 shows the throughput obtained without the use of RTS/CTS.The throughput obtained with the use of RTS/CTS is lower due the overhead associated with the RTS/CTS frames. In both cases, flows which are coupled share the bandwidth approximately equally, while node 1 gets completely starved. As we mentioned before, all transmitters are outside the interference ranges of all the receivers. Therefore, we can safely schedule their transmissions simultaneously. To do this, we set the CST of all transmitting stations to be -83 dBm to decouple each flow from the others. The resulting throughput graphs are shown in Figure 5 when RTS/CTS is disabled and Figure 4 when RTS/CTS is enabled. It is clear that now each station obtains the maximum possible throughput and node 1 is no longer starved.

If the default fixed value of CST is too conservative, what should be the the value of the CST to be used by the transmitter? While the -83dBm value used in this experiment was chosen empirically, we would like to be able to determine the optimum threshold value dynamically. We answer this question with the following observation.

**Observation II:** *The optimum value of the CST is that at which the carrier sense range of the transmitter just covers the interference range of the receiver.*

Recall that if $d$ is the distance between $T$ and $R$,then the interference range of the receiver $R$ is a circle of radius $1.78d$ centered at $R$. The value of the CST needed at $T$ is the value needed to sense the carrier of any interfering source $I$ at or within a distance $(2.78)d$ from $T$. The strength of the signal from the farthest distance of any interfering source, $2.78d$, is

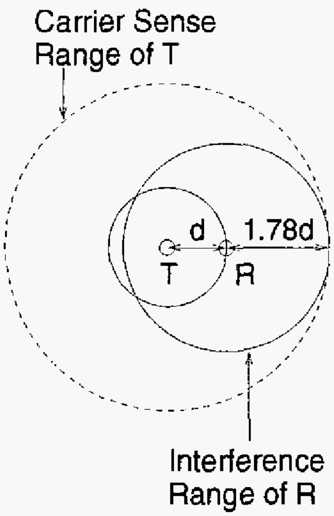


Fig. 7. The optimum value of CST is that at which the carrier sense range of $S$ just covers the interference range of $R$.

given by $k.P_0/(2.78d)^4 = k.P_0/d^4.\alpha = SS[T,R]/\alpha$, where $P_0$ is the transmit power and $SS[T,R]$ is the received strength of the signal from $T$ at $R$. Note that this can be generalized to any propagation model with a suitable value of $\alpha$. For propagation that decays as $d^4$, $\alpha$ is set to be $2.78^4$. The value of $\alpha$ can also be chosen empirically.

*Therefore, if $SS[T,R]$ denote the received signal strength from $T$ at $R$, then the optimum carrier sense threshold required at $T$ is given by $SS[T,R]/\alpha$, where $\alpha$ is a suitable constant determined by the propagation model.*

**Multiple Sources of Interference:** Due to the nature of carrier sensing, it can be shown the maximum number of independent interferers is bounded. If the receiver can provide feedback about the total amount of interference it sees, then the appropriate CST to be used by the transmitter can be determined. Thus, the model obtained for one source of interference can be generalized for multiple sources of interference as well.

We now use this observation to dynamically adjust CST values of clients and APs such that more flows can co-exist, resulting in higher per-flow and aggregate throughput in the hotspot. The details of these algorithms are described in the next section.

## IV. Architecture and Algorithms

In this section, we describe our $\mathcal{ECHOS}$ architecture and algorithms for a) dynamically identifying flows that can co-exist and b) allowing them to co-exist by setting optimum CST values based on exploiting our observations in the previous section.

### A. Overview

Consider a hotspot with four cells shown in Figure 8. The figure shows cells $C_0, C_1, C_2$ and $C_3$ with channels 0, 1, 0 and 2 respectively. The $\mathcal{ECHOS}$ architecture can operate in one of two modes, which we call AP-CST and RNC-SC, to improve performance.

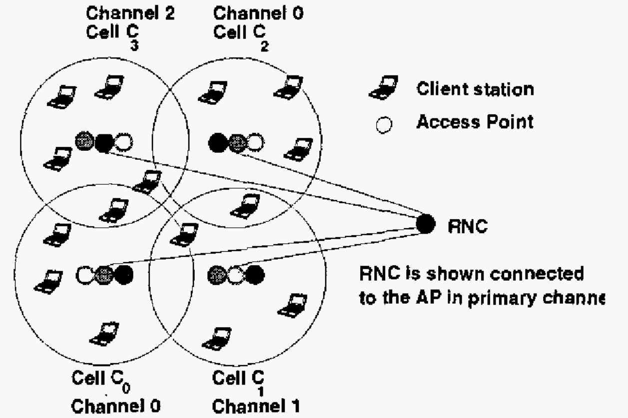


Fig. 8. The Echos Architecture. Each circle denotes the coverage range of the APs.

The mode AP-CST optimizes existing deployments by simply choosing the appropriate CST to be used by each of the APs and the stations of the hotspot. This algorithm is executed in a distributed manner in each of the APs and can be implemented in practice at no extra hardware cost. This algorithm helps cells in the hotspot that are assigned the same channels ($C_0$ and $C_2$ in this example) due to the limited number of orthogonal channels (3) available in 802.11b. This algorithm is explained in detail in subsection IV-B.

The mode RNC-SC is an enhancement to AP-CST and uses a centralized RNC which controls the entire hotspot (see Figure 8). Recall that this mode requires each AP to operate in all channels or there be co-located APs in all channels in each cell. For ease of presentation, we assume that there are APs in all three channels. While this approach entails additional hardware cost, it can improve performance significantly when the load distribution between the different cells in the hotspot is heterogeneous. In this mode, each cell $C$ is assigned a channel, called the *primary* channel, according to current best practice. Other available channels are called the secondary channels of $C$. The AP on the primary channel is called the primary AP and the AP on the secondary channels are called secondary APs. When the RNC finds $C$'s load to be high, it creates a *secondary cell* consisting of a secondary AP and some clients of $C$. A secondary cell is created only if there is no impact (in reality, minimal) on the primary cells both in terms of interference and reduction in throughput due to sharing the channel. In order to "insulate" the primary cell from a secondary cell, we set the Tx power of the secondary cell and the CST of both primary and secondary cells appropriately. This is explained in detail in subsection IV-C.

In both modes, all clients periodically report information about load and current signal conditions to the AP (explained in the next subsection in detail.) In RNC-SC, the APs then relay the information further to the RNC that manages the APs and the clients in a centralized manner.

| | Information needed for AP-CST |
|---|---|
| $SS[AP, s]$ | Signal Strength (SS) of AP at $s$ |
| $SSin_{min}[s]$ | Min SS of inside cell signals at s |
| $SSout_{max}[s]$ | Max SS of outside cell signals at s |
| $SS[s, AP]$ | SS of each client $s$ at AP |
| $SSout_{max}[AP]$ | Max SS of outside signals at AP |
| | Additional Info. for RNC-SC |
| Load | Average load in each cell |
| $SSout_{max}[AP, k]$ | Max SS of outside signals in Ch. k |
| $SSout_{max}[s, k]$ | Max SS of outside signals in Ch. k |

TABLE I

INFORMATION MAINTAINED AT AP AND RNC

```
AP-CST (run at each AP which senses another cell)
     ▷ Set CST of clients & compute CST_min ;
 1   CST_min ← ∞
 2   for each station s of cell C
 3       do
 4           CST[s] ← SS[s, AP]/α − ε
 5           CST[s] ← min(CST[s], SSin_min[s])
 6           CST_min ← min(CST_min, CST[s])
     ▷ Get lowest strength among all clients at AP.
 7   for each station s of cell C
 8       do
 9           SNR ← SS[AP, s]/SSout_max[s]
10           SNR_min ← min(SNR, SNR_min)
     ▷ Set the CST of AP.
11   if (SNR_min > γ)
12       then CST[AP] ← SSout_max[AP] + ε
13       else CST[AP] ← CST_min
```

## B. *Algorithm* AP-CST

Each AP keeps track of the parameters shown in Table I. Each of the above information can be obtained through the appropriate software interface to the 802.11 card. The only issue is of *distinguishing* between inside and outside cell transmissions. This is done as follows. Observe that any data transmission within a cell is to/from the AP. We assume that a station can hear most of the transmissions from its AP; otherwise the station should re-associate with a closer AP in the hotspot. Therefore, for any data transmission within the cell, a station would either hear the data frame (and/or RTS) or the ACK (and/or CTS) from the AP. Otherwise, the station assumes that the transmission was from outside the cell.

The basic idea behind AP-CST is that if a cell $C$ senses signals from outside the same cell, say, cell $D$, we can potentially improve the performance of cell $C$ by identifying those flows of $C$ which can co-exist with interference from $D$ and allowing them to proceed by adjusting the CST of stations/AP of cell $C$. Likewise, cell $D$ can also be optimized. This can be done because most flows, depending on the location of clients with respect to APs, have a certain degree of tolerance for additional interference. If, by reducing the CST, we can allow additional flows to operated without causing interference beyond the available tolerance of existing flows, we have improved the performance for those flows.

Algorithm AP-CST is run at each AP when it senses another cell in the same channel. It consists of two main steps:

- Sets CST of each client while avoiding creation of new hidden terminals
- Sets CST of AP while ensuring that all of its clients are still served

As can be seen from lines 2-5, the CST of each station $s$ is set to the maximum signal strength that can be ignored by $s$ while attempting to transmit, given the signal strength of $s$ at the $AP$; this value is given by $SS[s, AP]/\alpha$ as explained in Observation II. The CST is then adjusted so that the station $s$ can still hear other stations from inside its own cell, i.e, $SSin_{min}[s]$ (line 5). *This avoids creation of new hidden terminals.*

Assigning the CST of the AP requires a more conservative approach as the AP needs to serve the needs of all its clients. *Therefore, the minimum SNR of the AP at its clients decides the CST of the AP.* Lines 7-13 decide the optimum CST for the AP. If $SNR_{min}$ is sufficiently large, then the AP ignores all interference from outside the cell (line 12). Otherwise, it chooses the optimum CST value as in Observation II (line 13).

To summarize, each cell tries to "insulate" itself from any other co-channel cells by increasing the CST of the AP and the stations. The increase is determined by strength of the signals within the cell and constrained by the requirement that no new hidden terminals be created among the stations. The AP-CST algorithm can be implemented with CST values piggybacked on beacons. We next describe algorithm RNC-SC, which is an enhancement over AP-CST.

## C. *Algorithm* RNC-SC

Recall that, in this mode, the APs are managed by a centralized *Radio Network Controller* (RNC) and each AP has access to all available channels. In order to execute RNC-SC, the algorithm require additional parameters as shown in Table I. The additional parameters needed are load and signal strength values over all channels. In order to compute the latter, the RNC uses the values obtained by a periodic site-scan[13] by the clients of a cell on all the secondary channels. In a site-scan, all stations passively listen to the scanned channel for a specified duration and report the results back to the AP. This gives the RNC effective signal strength of transmissions from the primary cell in that channel at the stations/AP of secondary cell.

*a) Algorithm RNC-SC:* has two main steps:

- Determine if a cell is overloaded.
- Choose and switch a client to a secondary channel in overloaded cell, if possible.

*b) Measuring load and overload:* We use the MAC service time (i.e., the time between the instant a frame is submitted to the MAC for transmission and the time instant the ACK is received) seen by a node as the measure of load. This value is smoothed using an exponential filter and averaged over all members of the cell. It has been shown in [10] that the throughput of an 802.11 system is tied to the MAC service time seen by each node. In order to determine overload, the algorithm (lines 1 to 3) compares the the cell $C$ with the maximum average service time ($H$) and the cell $D$ with the

RNC-SC (run periodically)

1 $L \leftarrow$ lowest avg. MAC svc time of any cell $D$
2 $H \leftarrow$ highest avg. MAC svc time of any cell $C$
3 **if** $H/L > 1 + \delta$
4 **then**
5 $(s, k) \leftarrow$GET-CLIENT-SECCHANNEL(C)
6 **if** $s \neq \emptyset$
7 **then**
8 Activate AP on channel $k$ if needed
9 Switch client $s$ to channel $k$
10 Run AP-CST at each AP

minimum average service time ($L$). If $H/L > 1 + \delta$, where $\delta$ is a threshold providing hysteresis, then the RNC decides that cell $C$ is overloaded.

Once a cell has been identified as overloaded, we now explain the issues in identifying a potential secondary channel and client to switch.

GET-CLIENT-SECCHANNEL($C$)

$\triangleright$ $I_t[k]$ is max. tolerated interference to Channel k
1 $\forall k\ I_t[k] \leftarrow \infty$
2 **for** each channel $k$
3 **for** each primary cell AP $D$ on channel $k$
4 **for** each station $s$ of $D$
5 **do**
6 $T_1 \leftarrow SS[D, s]/\gamma - SSout_{max}[s, k]$
7 $T_2 \leftarrow SS[s, D]/\gamma - SSout_{max}[D, k]$
8 $I_t[k] \leftarrow min(I_t[k], T_1, T_2)$

$\triangleright$ Reduce Tx power to limit interference to within $I_t[k]$.
9 **for** each secondary channel $k$
10 **do**
11 $r_{AP}[k] \leftarrow SSout_{max}[AP, k]/I_t[k]$
12 **for** $q \in C$
13 $r_q[k] \leftarrow SSout_{max}[q, k]/I_t[k]$

$\triangleright$ Pick station which can communicate at reduced power
$\triangleright$ and has measured the least outside cell interference
14 $minIntf \leftarrow \infty$; $client \leftarrow \emptyset$; $secChannel \leftarrow \emptyset$
15 **for** each secondary channel $k$
16 **for** each station $q$ in $C$
17 **do**
18 **if** (SNR of $q$ at AP $> \gamma$ on channel k) **AND**
19 (SNR of AP at $q > \gamma$ on channel k) **AND**
20 $(minIntf > SSout_{max}[q, k])$)
21 **then**
22 $minIntf \leftarrow SSout_{max}[q, k]$
23 $secChannel \leftarrow k$
24 $client \leftarrow q$

25 return $(client, secChannel)$

*c) Choosing Client and Secondary Channel:* In order to create a secondary cell that has no impact on the primary cell, we need 1) any AP/station of the secondary cell should not interfere at the primary; and 2) the throughput in primary should not decrease because of this change. We call this a *conservative* approach because it does not allow for any inter-cell throughput adjustments; primary cells are largely unaffected by any new secondary cell formation.

A client $q$ of cell $C$ can become a member of a new secondary cell, if $q$ and its secondary AP, can communicate at a transmit power low enough that it does not affect the primary cells operating on that channel. The client and secondary channel are identified in Algorithm GET-CLIENT-SECCHANNEL, which has three main steps

- Compute maximum tolerated interference on each secondary channel k
- Reduce the transmit powers of secondary AP and clients on each secondary channel k
- Choose the client,channel pair such that the client observes minimum interference from outside the cell on that channel

We look at each of these steps in detail below:

*Calculating maximum tolerated interference:* the maximum interference tolerated at a primary AP is the minimum of the interference values tolerated by all its clients and the primary AP itself. Line 6 computes the interference that can be tolerated by station $p$. The interference that can be tolerated by $p$ when its AP alone is transmitting is given by $SS[AP, p]/\gamma$, where $\gamma$ is the capture threshold. However, $p$ already sees signals of strength $SSout_{max}[p]$ from stations which are outside the cell (which could be from a secondary cell that is already functional). Therefore, the remaining interference which $p$ can tolerate from the creation of any *new* secondary cell is given by $SS[AP, p]/\gamma - SSout_{max}[p]$. Likewise, the maximum interference which can be tolerated by the AP depends on the *lowest* received signal strength from any of its clients as computed in line 7.

*Reducing transmission power:* Given the maximum interference tolerated by the primary APs on channel $k$, any new secondary station/AP should produce interference of less than this value. As mentioned earlier, in order to estimate the interference by the secondary cell, the RNC uses values obtained by a periodic site-scan [13] by the clients. The RNC uses the maximum signal strength measured at a station/AP during the passive scan on channel $k$ as the estimate of maximum interference caused by that station/AP at any primary cell on channel $k$ at the normal transmit power.

Suppose station $q$ reports a measurement of $I_q$ and its primary AP $Q$ reports a measurement of $I_Q$ on secondary channel $k$ during the passive scan. If $I_t$ be the maximum tolerated interference on secondary channel $k$, then the secondary AP $S$ and $q$ should set their transmit powers to $P_t * I_t/I_Q$ and $P_t * I_t/I_q$ respectively. This factor of reduction is computed in lines 11 and 13 of algorithm GET-CLIENT-SECCHANNEL respectively. This reduction ensures that the additional interference due to the secondary stations does not exceed the maximum tolerated interference.

*Choosing client and channel with minimum interference:* We would like to choose the client, channel pair such that the client observes minimum interference from outside the cell on that channel. This approach maximizes the throughput of the client and minimizes the interference caused by the client to the outside cells. We ensure that communication is possible between the client, $q$, and the secondary AP at the reduced *transmit* powers with their interference values of $I_q$ and $I_Q$. If so, then $q$ is identified as a member of the secondary cell and the transmission bit rate is set to the highest possible for that transmit power. This is computed in lines 15 to 24 of algorithm GET-CLIENT-SECCHANNEL.

Thus, once a client and channel pair has been identified by the GET-CLIENT-SECCHANNEL routine, the RNC-SC activates the AP on the secondary channel if it was not active previously and switches the client to the secondary channel. The RNC-SC algorithm then calls AP-CST to be executed at each of the APs so that optimum CST values can be recomputed after the change in one of the clients.

Note that RNC-SC only switches at most one client from the overloaded cell at each run. This is again a conservative approach as the interference values and loads of the various cells could change as a result of this switch. Since RNC-SC is called periodically (example, every five seconds), we ensure that the performance improvement takes place at a measured pace.

To summarize, when a cell is overloaded, the RNC switches a client of the cell to a secondary cell, if possible. The switch is possible if the secondary AP and client can operate at low enough transmission power to not interfere at primary cell on the same channel. Further, the primary cells set their CST values appropriately to ignore the secondary cell's activity, while the secondary cell does so to the extent possible at its reduced power. We present a detailed evaluation of the AP-CST and RNC-SC algorithms in the next section.

## V. PERFORMANCE EVALUATION

The topology considered is shown in Figure 9. An area of 1000m x 1000m is divided into 4 cells. The primary channels of these cells are assigned according to current best practice as shown in the figure. Each AP has a coverage range of 250m approximately at 11Mbps. Within each cell, the positions of client stations are sampled from a two-dimensional uniform distribution. All simulations were conducted using the Qualnet simulator. We study the system with the users stationary in their locations. Thus, the load varies due to application characteristics, but not due to user mobility. The maximum bitrate was set to 11Mbps, each simulation run lasted for 100 seconds, and results were averaged over 10 runs. The two-ray propagation model was used with no fading. Fading was disabled as estimating the correct signal strength average values in a realistic fading channel is a non-trivial problem in itself. As shown in [24], it requires an estimation of the correlation between signal strength samples; our emphasis is on estimating maximum gains possible with the use of dynamic CST.

### A. *Homogeneous user/load distribution*

First, we evaluate the performance of the AP-CST algorithm on a system with homogeneous client distribution, i.e., with each cell having the same number of stations and workload.

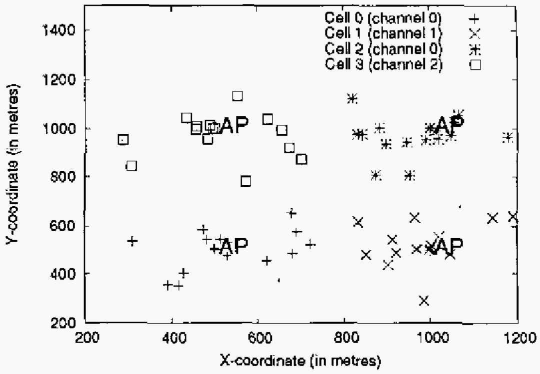


Fig. 9. Distribution of client nodes within the topology

| Cell ID | Old Tput. | Std. Dev.(%) | New Tput. | Std. Dev.(%) | Gain (%) |
|---|---|---|---|---|---|
| 0 | 2.79 | 0.88 | 4.83 | 1.76 | 73.11 |
| 1 | 5.02 | 1.11 | 5.02 | 1.11 | 0 |
| 2 | 2.87 | 1.49 | 4.81 | 1.64 | 67.35 |
| 3 | 5.04 | 0.94 | 5.04 | 0.94 | 0 |

TABLE II

AP-CST RESULTS WITH HOMOGENEOUS USER/LOAD DISTRIBUTION. OVERALL IMPROVEMENT IS 25.31%.

Each cell has 15 clients. The number of clients is chosen to be half the maximum number of clients allowed per cell in current practice (30). The topology so obtained is shown in Figure 9. Each client station has an **HTTP** [2] client with a think time of 1s. That is, the gap between two successive requests for web pages is random with maximum value at most 1 second. In addition, each cell has **one** client doing an **FTP** download. The HTTP and FTP servers are one wired hop away from the APs. TCP ACKs flow upstream and therefore, there is traffic both upstream as well as downstream. AP-CST is executed every 5s.

**Base case** We first consider the system operating at 11Mbps without the use of RTS/CTS or Auto-Rate-Fallback (ARF)[3] The results are shown in Table II. The overall improvement in the entire system's throughput is about 25.30% with the per-cell improvements in cells $C_0$ and $C_2$ being 73.11% and 67.35% respectively. This shows that even Algorithm AP-CST by itself can provide significant improvement in performance.

**Impact of RTS/CTS** Algorithm AP-CST yields significant improvement in throughputs even with the use of RTS/CTS as summarized in Table III. The throughput of all cells is lower when RTS/CTS is used, due to the extra delay in each packet's service time and any impact it may have on TCP workloads. However, the improvement in performance with a dynamic CST is still maintained. *This confirms that an appropriate CST improves performance of RTS/CTS enabled 802.11 systems as well.*

[2]Qualnet uses empirically obtained CDFs (from real traces) for size of HTTP objects retrieved, number of objects per web page, and think time to generate HTTP traffic.

[3]ARF (Auto-Rate-Fallback) [14] is the mechanism by which an 802.11b card switches down to a lower bit rate in the presence of loss.

| Cell ID | Old Tput. | Std. Dev.(%) | New Tput. | Std. Dev.(%) | Gain (%) |
|---|---|---|---|---|---|
| 0 | 1.84 | 0.21 | 3.25 | 0.38 | 76.81 |
| 1 | 3.49 | 1.59 | 3.49 | 1.59 | 0 |
| 2 | 1.77 | 0.74 | 3.17 | 1.00 | 78.86 |
| 3 | 3.43 | 1.21 | 3.43 | 1.21 | 0 |

TABLE III

AP-CST IMPROVEMENT WITH RTS/CTS ENABLED. OVERALL IMPROVEMENT IS 26.7%.

| Cell ID | Old Tput. | Std. Dev.(%) | New Tput. | Std. Dev.(%) | Gain (%) |
|---|---|---|---|---|---|
| 0 | 2.67 | 0.89 | 4.60 | 1.17 | 71.99 |
| 1 | 4.68 | 0.73 | 4.68 | 0.73 | 0 |
| 2 | 2.70 | 0.10 | 4.57 | 1.66 | 69.20 |
| 3 | 4.64 | 0.73 | 4.64 | 0.73 | 0 |

TABLE IV

AP-CST IMPROVEMENT WITH AUTO-RATE FALLBACK. OVERALL IMPROVEMENT IS 25%.

**Impact of ARF** The improvements obtained by AP-CST with ARF enabled in the entire system are shown in Table IV. We note that when ARF is used, the throughput of the entire system falls when one station switches to a lower bit rate due to loss. This is because a station at a lower bit rate gets equal access to the medium in terms of packets transmitted but simply takes longer to transmit each packet. However, *the gains due to a dynamic CST are still maintained to almost the same extent as with no ARF.*

In the rest of the paper, we do not report about use of RTS/CTS and ARF for sake of brevity.

**Quantifying interference** We consider the following question: how much is the penalty due to potential increase in the interference caused by higher CSTs? To do this, we consider two metrics, namely, *collisions* and *retransmissions*. Each collision causes at least one retransmission; one of the colliding frames might survive due to physical layer capture at the receiver. Each retransmission is caused due to either a collision from another station in the same cell or interference from a co-channel cell. Figure 10 and 11 show the number of collisions and retransmissions as a fraction of the total number of packets transmitted respectively. As expected, for the cells $C_1$ and $C_3$ which have no co-channel interference from neighboring cells, the number of collisions and retransmissions does not change at all. These collisions occur due to the inherent randomness in the 802.11 DCF protocol. For the cells $C_0$ and $C_2$ which use the same channel, the number of collisions and retransmissions increases slightly when AP-CST is used. However, note that the increased number of collisions and retransmissions of cells $C_0$ and $C_2$ are comparable to those of $C_1$ and $C_3$. Furthermore, as seen in Tables II to VI, the new throughput in cells $C_0$ and $C_2$ are comparable to those of $C_1$ and $C_3$. *Therefore, there is no significant increase in the number of collisions and retransmissions due to co-channel interference because of optimizing the CST; the increase is mainly due to the increased throughput/activity in the cell.*

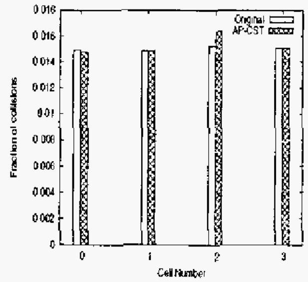


Fig. 10. Collisions

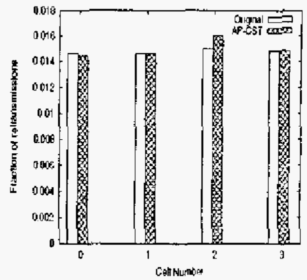


Fig. 11. Retransmissions

**Evaluation of RNC-SC** Next, we consider the evaluation of RNC-SC for homogeneous client distribution. Algorithm RNC-SC is executed every 5 seconds of the simulation. For overloaded cells, the duration of the scan in each secondary channel is set to be 100 milliseconds and the scanning is done every 20 seconds. The results (not shown) indicate that the after the algorithm AP-CST has been executed, RNC-SC does not find any difference in the loads of different cells (since the cells have uniform load) and therefore, there is hardly any improvement over what has been achieved using AP-CST.

### *B. Heterogeneous User/Load Distribution*

In the case of heterogeneous load distributions, the AP-CST algorithm cannot do much beyond decoupling the two co-channel cells; with RNC-SC, since overloaded cells can borrow "extra" secondary channels, it is possible to attain significant performance gains over that of AP-CST.

We now consider the performance of the architecture, when users are distributed heterogeneously, i.e, there is overcrowding in few cells, while the capacity of other cells is not utilized fully. We consider cases when two out of four cells are overloaded (results of cases where one cell is overloaded are subsumed in this, so we omit one cell overloaded cases for brevity). Each station in all cells has an HTTP client as workload. In addition, there are $N_c/10$ FTP download clients, where $N_c$ is the number of stations in 'the cell, in the overloaded cells. We consider three scenarios when the two most loaded cells are:

**I** Co-channel cells $(C_0, C_2)$
**II** Orthogonal cells $(C_1, C_3)$
**III** Both co-channel and orthogonal cells $(C_0, C_3)$.

*d) Scenario-I:* We start with a workload where cell $C_0$ and $C_2$ are the most heavily loaded with 25 clients and 15 clients respectively. Cells $C_1$ and $C_3$ have 10 clients each. The cells have [2,0,1,0] FTP clients respectively.

The improvements in the long-term throughputs (after reaching steady state) of the different cells are summarized in Table V. Cells $C_0$ and $C_2$ have gains of up to 195.28% and 106.99% respectively, while the impact on the primary cells $C_1$ and $C_3$ are marginal. The overall improvement in the throughput across all four cells is 70.31%.

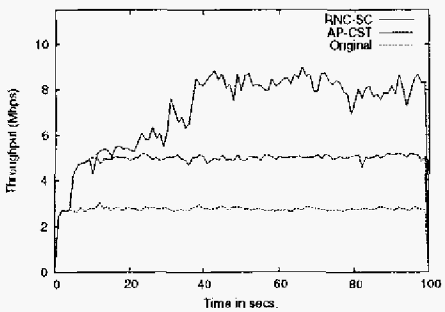


Fig. 12. Cell $C_0$ in Scenario I

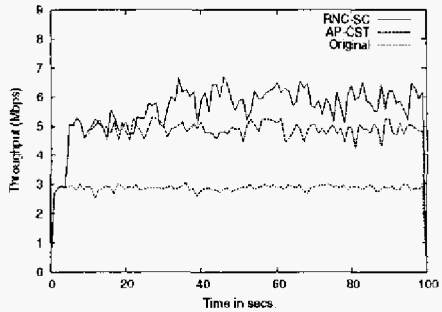


Fig. 13. Cell $C_2$ in Scenario I

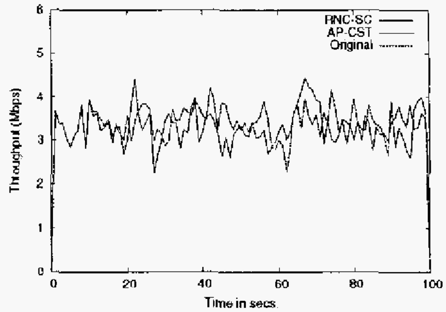


Fig. 14. Cell $C_1$ in Scenario I

| Cell | Old Tput (Mbps) | Dev. (% of Mean) | New Tput (Mbps) | Dev. (% of Mean) | Gain (% of Old) |
|---|---|---|---|---|---|
| 0 | 2.76 | 0.85 | 8.15 | 3.35 | 195.28 |
| 1 | 3.38 | 6.94 | 3.35 | 5.07 | -1.08 |
| 2 | 2.86 | 1.42 | 5.92 | 6.42 | 106.99 |
| 3 | 2.89 | 7.95 | 2.83 | 10.36 | -1.81 |

TABLE V

IMPROVEMENTS ACHIEVED IN SCENARIO-I WITH $C_0$ AND $C_2$ MOST LOADED . OVERALL IMPROVEMENT IS 70.31%.

| Cell | Old Tput (Mbps) | Dev. (% of Mean) | New Tput (Mbps) | Dev. (% of Mean) | Gain (% of Old) |
|---|---|---|---|---|---|
| 0 | 2.61 | 4.26 | 2.67 | 22.57 | 2.28 |
| 1 | 5.00 | 0.33 | 9.67 | 4.49 | 93.40 |
| 2 | 2.68 | 4.93 | 3.30 | 7.42 | 23.21 |
| 3 | 4.95 | 1.20 | 8.21 | 3.54 | 65.85 |

TABLE VI

IMPROVEMENTS ACHIEVED IN SCENARIO-II WITH ORTHOGONAL CHANNEL CELLS BEING MOST LOADED. OVERALL IMPROVEMENT IS 56.5%.

Figure 12 shows the instantaneous throughput of cell $C_0$. The gain in the throughput occurs in identifiable stages. At $t = 5s$, the first instance of AP-CST is executed by cell $C_0$ as it has sensed cell $C_2$ on the same channel. It improves the throughput of the cell to about 5Mbps. *Note that this is the maximum improvement possible with* AP-CST *alone.* The RNC finds the load of cell $C_0$ to be still high after this, and so starts switching stations to secondary cells. This goes until $t = 35s$, when the cell $C_2$ is decided as being the most loaded. Around $t = 35s$, $C_2$'s throughput improves from 5Mbps to about 6Mbps as seen in Figure 13 as a result of secondary cell operation. Once this is done, $C_1$ is again decided as loaded and it climbs as a result of further switches to its maximum value of about 8Mbps, where it stabilizes.

Finally, we consider the instantaneous throughput of primary cell $C_1$ to see the impact of secondary cell creation. Figure 14 shows the throughput obtained by cell $C_1$. For the first 10 seconds, both curves match completely. After that, when RNC-SC is running, the throughput obtained by cell $C_1$ varies slightly from the original. This is because, in some runs, signal measurements from the passive scan are not estimated correctly due to the limitation of the scan duration ($100ms$). This results in creation of a few secondary clients which would not have been allowed in our conservative approach. *This illustrates a limitation of our signal estimation algorithm which we discuss in VI.* However, on the whole, the cell's throughput remains largely unaffected since the change of -1.08% is insignificant compared to the standard deviation.

*e) Scenario II:* Next, we consider a scenario where the stations are split among different cells as [10,25,10,15] with the cells on orthogonal channels, namely, $C_1$ and $C_3$ are the most loaded. Recall that each station has an HTTP client, and the cells have [0,2,0,1] FTP downloads respectively.

Table VI shows the results of improvements in the long-term throughputs of the cells. The overall throughput improves by 56.5%. An interesting observation of these results is that while cell $C_2$ has improved its performance by 23.21% as a result of the first instance of AP-CST, cell $C_0$ has not. This observation is correlated with the large deviation of the throughput of $C_0$.

To explain these observations, we consider the instantaneous throughputs of cells $C_0$ and $C_3$. Figures 15,16, and 17 show the instantaneous throughput of cell $C_0$,$C_1$, and $C_3$ respectively. Algorithm AP-CST provides no improvement for cells $C_1$ and $C_3$ as they are on orthogonal channels. Cells $C_1$ and $C_3$ are identified by RNC-SC as overloaded approximately alternatively and they increase their throughputs until the system stabilizes at around $t = 80s$. However, in some subset of the simulation runs, cell $C_0$ is wrongly identified as loaded. Further, among some members of those runs , $C_0$ wrongly identifies channel 2 as worthy of switching to and switches some station(s) to channel 2. Because $C_3$ "insulates" itself, it does not see any dip in its performance. Cell $C_0$ however is unable to insulate itself (as it only ignores only as much interference as it can at its reduced power), thereby losing out possible gains from AP-CST. The curve for AP-CST in Figure 15 is thus higher than the other two curves after about $t = 25s$. The gains for $C_0$ which were got from the first instance of AP-CST are lost due to these false positives. *Thus, these results illustrate that while there is some impact due to the combination of false positives in load as well as signal measurement (lack of performance improvement in cell $C_0$), overall throughput gains are still substantial.*

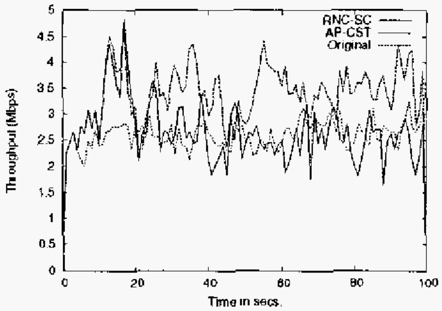


Fig. 15. Cell $C_0$ in Scenario II

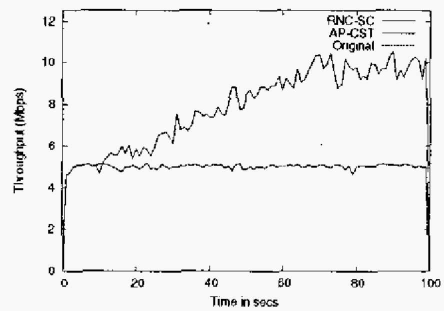


Fig. 16. Cell $C_1$ in Scenario II

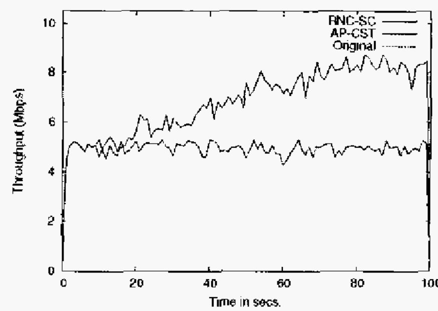


Fig. 17. Cell $C_3$ in Scenario II

| Cell | Old Tput (Mbps) | Dev. (% of Mean) | New Tput (Mbps) | Dev. (% of Mean) | Gain (% of Old) |
|---|---|---|---|---|---|
| 0 | 2.49 | 3.38 | 2.93 | 5.84 | 17.67 |
| 1 | 5.00 | 0.33 | 8.60 | 4.47 | 72.00 |
| 2 | 3.22 | 1.24 | 6.67 | 8.73 | 113.78 |
| 3 | 3.26 | 7.22 | 3.40 | 4.62 | 3.33 |

TABLE VII

IMPROVEMENTS ACHIEVED IN SCENARIO-III WITH CELLS $C_1$ AND $C_2$ BEING MOST LOADED OVERALL IMPROVEMENT IS 48.60%.

*f) Scenario III:* Finally, we consider the following assignment of stations to APs: [10,25,15,10]. The number of FTP clients in cells are [0,2,1,0]. The results are shown in Table VII. The overall improvement in the throughput is 48.60%. In this case, an orthogonal cell $C_1$ and a co-channel cell $C_2$ are both heavily loaded both in terms of number of client stations as well as the workload. The dynamics of cells $C_1$ and $C_2$ are shown in Figures 18 and 19 respectively. As expected, AP-CST helps cell $C_2$ but not $C_1$. After $C_2$ attains about 5Mbps as a result of AP-CST at $t = 5s$, it stays so till about $t = 40s$. In this duration, $C_1$ keeps switching stations to obtain a throughput of 7Mbps and stabilizes at $t = 40s$. After which $C_2$ gains until it stabilizes at around $60s$. Finally, $C_1$ starts gaining at around $t = 60s$, and stabilizes at around $t = 80s$.

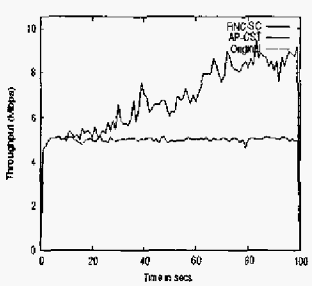


Fig. 18. Cell $C_1$,Scenario III

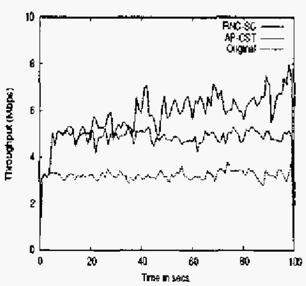


Fig. 19. Cell $C_2$,Scenario III

### C. Large and Arbitrary topologies

While we focused our evaluation on a four-cell scenario due to simulation tractability and the need to understand the subtleties of the algorithms, the architecture can easily be scaled to larger hotspots with 16 cells or more since $\mathcal{ECHOS}$ primarily relies on light weight feedback mechanisms (with parameters bounded by the maximum number of clients being served) and periodic execution (with periods of 5s or more) of non-compute intensive algorithms. Thus, scalability of the RNC is not an issue. However, in a large network, the presence of several co-channel cells, can hinder some cell's ability to optimize its CST or create secondary channels, limiting the overall gains in general. In such scenarios, cells which are likely to be optimized best are those which are at the periphery of the topology.

We also performed other simulations (not shown due to lack of space) with arbitrary topologies (for example, APs are not exactly aligned in a grid pattern, say, due to connectivity constraints) and found similar substantial improvements in performance (e.g. per-cell throughput improvement up to 97%).

### D. Summary

We found that the AP-CST algorithm provided per-cell throughput gains of up to 78% in the co-channel cells when the load was homogeneously distributed among the different cells in a hotspot. These gains were robust with or without the use of RTS/CTS and with or without the use of the multi-rate feature of 802.11. We then considered heterogeneous load distributions in different cells. In this case, as expected, algorithm RNC-SC provided significant improvement over the gains provided by AP-CST. The overall gains of RNC-SC were up to 195%. Furthermore, even in the presence of errors in signal strength or load measurements, the performance of primary cells were largely unaffected by the creation of secondary cells in RNC-SC.

## VI. DISCUSSION

In this section, we discuss potential refinements to our algorithm as well as their limitations.

### A. Refinements

Observe that while each client chooses its optimum CST, the AP chooses the most conservative value of the CST in order to ensure that the farthest client is not affected. This can be improved if the AP maintains the optimum CST for each station separately. The AP can choose the optimum CST for each transmission depending on the destination node. However, the penalties of false positives may be high. For instance, a packet at the head of the output queue of the AP may be retransmitted up to 7 times due to wrong estimations, thereby blocking packet to other clients as well.

We required that the SNR exceed a fixed threshold ($\gamma$) for successful reception. We could possibly allow for *non-conservative* formation of secondary cells by varying $\gamma$, and allowing overall throughput to increase even at the expense of mild loss in primary cells.

### B. Limitations

One limitation of the system is the use of the MAC service time metric as an estimation of the load. In a realistic environment, there is a likelihood that some stations re-transmit all packets several times due to a very poor channel, therby pushing up the average MAC service time. Clearly, this scenario cannot be helped by any scheme which uses another channel in the same operating conditions. This false positive can be handled by detecting if transmissions are more successful at a lower bit-rate.

The next limitation of the scheme is the signal strength measurement. Clearly, signal strength in real WLANs fluctuates. However, several schemes (reference [24] gives a list) which use signal strength to estimate user location in indoor environments have shown how to handle such variations in samples of signal strengths. Further, the scan duration can be adapted by the RNC depending on the load of the primary cells to account for capturing signals from all stations of primary cells.

## VII. Conclusion

Existing settings for CST in 802.11 networks are highly conservative. We advocate that transmitters dynamically chose their CST values depending on their signal strength at the receiver. This paradigm allows us to improve performance in infrastructure mode 802.11 hotspots using two algorithms namely, AP-CST and RNC-SC. The former improves the performance in cells which can sense another cell in the same channel by choosing an appropriate CST value. The latter provides gains over those provided by AP-CST by creating secondary cells which co-exist with primary cells on the same channel without causing any interference/reduction of throughput. Detailed simulation results show that a dynamic CST allows for significant improvement in the performance of hotspot networks. Directions for future work include making the signal and load measurement algorithms more robust to measurement errors.


## Acknowledgements

We thank the anonymous referees for their useful comments and suggestions. Access to the Qualnet network simulator was provided by Scalable Networks Inc. under the Qualnet University Program.


## References


[1] Arup Acharya, Archan Misra, and Sorav Bansal. Maca-p: A mac protocol to improve parallelism in multi-hop wireless networks. In *Proceedings of First IEEE International Conference on Pervasive Computing and Communications*, 2003.

[2] Airespace. http://www.airespace.com.

[3] A. Balachandran, P. Bahl, and G. Voelker. Hot-spot congestion relief in public-area wireless networks. In *Workshop on Mobile Computing Systems and Applications (WMCSA'02)*, June 2002.

[4] Yigal Bejerano and Randeep Bhatia. Mifi: A framework for fairness and qos assurance in current ieee 802.11 networks with multiple access points. In *Proceedings of IEEE Infocom*, 2004.

[5] Soekris boards. http://www.soekris.com.

[6] Cali, F., M. Conti, and E Gregori. Dynamic tuning of the ieee 802.11 protocol to achieve a theoretical performance limit. *IEEE/ACM Transactions on Networking*, 8(6):785–799, December 2000.

[7] M. Cesana, D. Maniezzo, P. Bergamo, and M. Gerla. Interference Aware (IA) MAC: an Enhancement to IEEE802.11b DCF. In *Proceedings of IEEE Vehicular Technology Conference - VTC Fall 2003*, 2003.

[8] Ranveer Chandra, Victor Bahl, and Pradeep Bahl. Multinet: Connecting to multiple ieee 802.11 networks using a single wireless card. In *Proceedings of IEEE Infocom*, 2004.

[9] P.Zerfos et al. DIRAC: A software based wireless router system. In *Proceedings of Mobicom*, 2003.

[10] H.Kim and J.Hou. Improving protocol capacity with model-based frame scheduling in ieee 802.11-operated wlans. In *Proceedings of the 9th annual international conference on Mobile computing and networking (MobiCom)*, 2003.

[11] N. Jain, S.R. Das, and A. Nasipuri. A multichannel csma mac protocol with receiver-based channel selection for multihop wireless networks. In *Proceedings of 9th International Conference on Computer Communications and Networks (IC3N)*, October 2001.

[12] Nitin H. Vaidya Jason Fuemmeler and Venugopal V. Veeravalli. Selecting transmit powers and carrier sense thresholds for csma protocols. In *Technical Report, Univ. of Illinois - Urbana Champaign*, 2004.

[13] IEEE 802.11 Working Group K. http://www.ieee802.org/11/.

[14] A. Kamerman and L. Monteban. Wavelan ii: A high-performance wireless lan for unlicensed band. In *Bell Labs Technical Journal, pages 118-133, Summer 1997*.

[15] Jeffrey Monks, Vaduvur Bharghavan, and Wen mei W. Hwu. A Power Controlled Multiple Access Protocol for Wireless Packet Networks. In *Proceedings of IEEE INFOCOM*, pages 219–228, 2001.

[16] A. Nasipuri and S.R. Das. Multichannel csma with signal power-based channel selection for multihop wireless networks. In *IEEE Vehicular Technology Conference*, September 2000.

[17] Meru Networks. http://www.merunetworks.com.

[18] Qualnet. Network simulator version 3.6. In *http://www.scalable-networks.com*, 2003.

[19] Gartner Research. http://www.gartner.com.

[20] Jungmin So and Nitin Vaidya. A multi-channel mac protocol for ad hoc wireless networks. In *UIUC Technical Report*, January 2003.

[21] TIA/EIA/cdma2000. . *Mobile Station - Base Station Compatibility Standard for Dual-Mode Wideband Spread Spectrum Cellular Systems"*. Washington: Telecommunication Industry Association, 1999.

[22] Kaixin Xu, Mario Gerla, and Sang Bae. How effective is the ieee 802.11 rts/cts handshake in ad hoc networks? In *Proceedings of Globecom*, 2002.

[23] Xue Yang and Nitin Vaidya. On the physical carrier sense in wireless ad hoc networks. In *Technical Report, University of Illinois at Urbana-Champaign*, 2004.

[24] Moustafa Youssef and Ashok Agrawala. Handling Samples Correlation in the Horus System. In *Proceedings of IEEE Infocom*, 2004.