Interoperability of Wi-Fi Hotspots and Cellular Networks*

Dilip Antony Joseph
University of California
Berkeley
dilip@berkeley.edu

B. S. Manoj
Indian Institute of Technology
Madras
bsmanoj@cs.iitm.ernet.in

C. Siva Ram Murthy†
Indian Institute of Technology
Madras
murthy@iitm.ac.in

ABSTRACT

The widespread deployment of Wi-Fi hotspots and wide area cellular networks opens up the exciting possibility of interoperability between these types of networks. Interoperability allows a mobile device to dynamically use the multiple network interfaces available to it so as to maximize user satisfaction and system performance. In this paper, we define three basic user profiles for the network users and demonstrate through simulation studies that dynamic switching on the basis of the user profiles of the mobile devices leads to higher network performance and increased user satisfaction. Careful design of pricing, billing and revenue sharing schemes is necessary to ensure the commercial viability of the multiple service providers involved in an inter-operable network setting. Different pricing and revenue sharing schemes are introduced and analyzed using simulation studies. We also demonstrate how load balancing can improve network performance in an inter-operable network.

Categories and Subject Descriptors

D.2.8 [**Computer-Communication Networks**]: Network Architecture and Design—*wireless communication*

General Terms

Measurement, Performance, Experimentation, Economics

Keywords

Wi-Fi hotspot, packet cellular networks, interoperability, user behaviour, pricing

1. INTRODUCTION

Recent years have seen the widespread deployment of Wi-Fi hotspots in a variety of environments such as airports, railway stations, offices, and other places of commercial interest. A number of companies offer high speed network access services through these hotspots. Cellular networks have also seen an explosive growth in the number of subscribers. For example, there are more than one billion GSM [1] subscribers and more than 50 million CDMA2000 [2] subscribers today. The cellular networks, which traditionally offered only voice calls, now offer data and multimedia services. Mobile devices of the future are likely to support both cellular and Wi-Fi interfaces.

Both cellular and wireless LAN technologies have their pros and cons. The question naturally arises – Is it possible to combine the advantages of both technologies to form a unified system? The answer is YES – Interoperability is the key. Consider the following common scenario involving a typical business executive. While traveling to his office by road every morning, the executive retrieves his email and accesses the Internet via a CDMA PCMCIA card fitted on his laptop. As he walks into his office, while a big file download is still in progress, the system transparently switches to the high speed wireless LAN (WLAN) available in the office building. The wireless LAN offers a much higher bandwidth than the cellular network and is also less expensive (often free) to access. Thus switching to the wireless LAN results in an enhanced Internet surfing experience for the executive as well as in substantial cost savings. Later in the day, if the office WLAN becomes very loaded (for instance, due to a large number of users simultaneously connecting to it), the system may again switch to the CDMA network to get a higher bandwidth, if the *user profile* of the executive demands it. All this switching happens transparent to the user – he enjoys smooth network connectivity at all times in accordance with his requirements.

1.1 Advantages of Interoperability

Interoperability allows a mobile device to dynamically use the multiple network interfaces available to it so as to maximize user satisfaction and system performance. For example, on a laptop fitted with both a CDMA wireless modem and an IEEE 802.11 WLAN card, the Internet can be accessed through either of the interfaces, depending on the user constraints and current network conditions.

Interoperability enables us to combine the advantages of both wireless LANs and wide area cellular networks in providing high speed connectivity inexpensively to a large number of users in a wide coverage area. The Wi-Fi Access Points (APs) provide high bandwidth and fast network access to the users in their coverage areas. The base stations (BSs) of the cellular network ensures network connectivity

---

[†] Author for correspondence.

[*] This work was supported by the Department of Science and Technology, New Delhi, India.

WMASH'04, October 1, 2004, Philadelphia, Pennsylvania, USA.

in a wide region, although at lower speeds. Wi-Fi APs also extend this coverage to places like subways and interiors of buildings, where the signal strength of the cellular network is very weak. Interoperability with wireless LANs also helps in relieving the heavy load on cellular networks, especially in crowded regions. Load balancing between the two different networks can increase the overall system throughput. Moreover, Wi-Fi APs can be set up at a fraction of the cost of installing BSs for the cellular network. An inter-operable system can also support a much larger number of users than any single network. This leads to higher revenues for all the network service providers involved.

1.2 Issues in Interoperability

Achieving interoperability between wireless LANs and cellular networks is a very challenging task. Here we briefly discuss some of the major issues involved in achieving interoperability.

- **When to Switch?**

  The manner in which a mobile device dynamically switches between the multiple network interfaces available to it greatly influences the performance and resource consumption of the system. *User Profile* based switching as a possible solution is explored in Section 3.

- **Smooth Handoffs**

  Maintaining existing network connections while switching between different network interfaces is a difficult task. Smooth handoffs involve issues of diverse addressing schemes, different packet formats and sizes, and packet sequencing across multiple networks.

- **Billing and Revenue Sharing**

  The presence of multiple service providers makes billing and revenue sharing very challenging tasks. It is essential that the schemes are so designed that the commercial viability of all service providers is guaranteed.

- **Security and Authentication**

  A mobile device can connect to multiple networks at the same time. Authentication of the users and security of data transmitted across diverse networks are not easy to achieve.

- **Load Balancing**

  Load balancing between different network architectures requires new metrics to ascertain the load and novel schemes to shift a section of users to a different network, whenever needed.

- **Implementation**

  Implementation of interoperability requires changes in both the network protocols as well as in the protocol stacks of the mobile devices. Maintaining compatibility with the existing systems and protocols while incorporating interoperability is very important.

- **Quality of Service**

  Mobile devices of the future will run applications that have stringent Quality of Service (QoS) requirements. Ensuring QoS in a system supporting handoffs between multiple networks with diverse characteristics is a very challenging problem.

- **Inter-Service Provider Agreements**

  Interoperability calls for co-operation between the different service providers on a large number of issues - for example, in billing and revenue sharing. Conflicting interests often make inter-service provider agreements hard to achieve.

The rest of this paper is organized as follows. In Section 2, we describe the interoperability system architecture and packet routing mechanisms. Section 3 describes the impact of user behaviour on interoperability. Section 4 discusses various pricing schemes and revenue sharing schemes, and also introduces the concept of load balancing by dynamic pricing. Issues in the implementation of interoperability are discussed in Section 5, while Section 6 concludes the paper.

2. SYSTEM ARCHITECTURE

The system under consideration consists of a set of Base Stations (BSs) belonging to a Packet Cellular Network that can provide wide area coverage to the Mobile Stations (MSs), and a set of Wi-Fi Access Points (APs) that provide high speed connectivity to the MSs. Each MS is assumed to support only one wireless interface, that can switch between the packet cellular and Wi-Fi modes of operation. The Wi-Fi APs are assumed to be interconnected with one another, and with the BSs by means of either a wired backbone network, or by high bandwidth point-to-point wireless links. The BSs are connected by means of a high speed wired network.

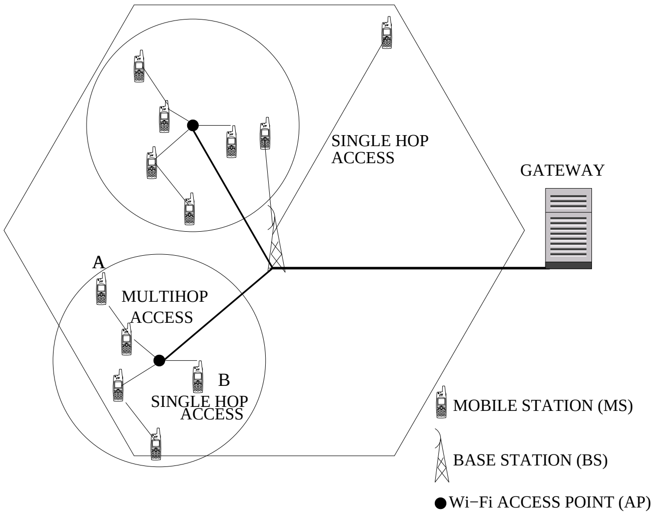

Figure 1: The Interoperability Framework

Figure 1 shows a schematic representation of the system under consideration. The BSs are placed such that the entire terrain is covered, while the Wi-Fi APs are assumed to be randomly distributed throughout the metropolitan area. The packet cellular network in our system is represented by a Single hop Cellular Network (SCN), in which the MSs are in communication with the BS on the control as well as data channels, with a transmission range equal to that of the cell radius. The Wi-Fi hotspots are considered as multi-hop relaying environments similar to the Multi-hop Cellular Networks described in [3] and [4]. The Wi-Fi AP acts as the coordinator for enabling routing and reserving bandwidth for MSs in the hotspot. The MCN architecture as described in [5] assumes a control interface of transmission range equal to that of the cell radius, and a data channel with a transmission range equal to half the cell-radius. Wi-Fi APs with multi-hop relaying is an attractive option in WLANs as it can extend the coverage of a high bandwidth AP to a much larger area. As is evident from the system architecture, each MS has the option of operating either under the control of the BS or under the control of the Wi-Fi AP.

The protocol proposed in [5] works as an infrastructure-aided source routing mechanism [6], that uses the topology information available at the BSs. We now present a brief description of the protocol in [5], and also describe some of the modifications needed to support the new inter-operable architecture. Each BS periodically generates *Beacon* messages that can be received by all mobile stations within its coverage area. An MS chooses to register with a particular BS depending on the received signal strength. It then sends the *RegReq* packet, to which the BS replies with a *RegAck* packet.

Once the registration is complete (*i.e.,* after the MS receives the *RegAck* packet), the MS will periodically generate *Beacon* messages. It updates the BS with information about the set of neighbors that are within its transmission range through *NeighUpdt* packets. Whenever there is a packet to be sent, the source MS originates a *RouteReq* packet to its BS. The BS responds with a *RouteReply* packet containing the shortest path which the MS uses to source route the packet to its destination.

In order to deal with MSs of different user profiles, both the BSs and the Wi-Fi APs periodically generate *Beacons* with transmission ranges R, the cell radius of the cellular network and r, the transmission range of the control channel of the Wi-Fi hotspot (modeled as an MCN) respectively. Each such beacon advertises the per byte transmission cost levied by the AP or BS as well as an estimate of the free bandwidth that will be available to a new MS registering with it. An estimate of the free bandwidth is obtained by dividing the total bandwidth available at the BS or AP by the total number of MSs currently registered with it. This approach assumes that the APs and BSs have a fair packet scheduling scheme running. The crucial difference between the MSs with different user requirements occurs in the registration mechanism.

The network usage that we consider in this paper is essentially one of gateway access, with one of the BSs acting as a gateway to the Internet or as a content server. This means that each MS needs to only find a route to its nearest infrastructure node (either a BS or an AP), which can then connect to the gateway by means of the backbone network. For MSs registered to the Wi-Fi AP, the routing mechanism proceeds in a similar fashion to the MCN routing protocol discussed above. However in the case of multi-hop wireless LANs, the problem of network partitions can arise, especially if the node density around the AP is low. This essentially means that the MS cannot find a multi-hop path over the data channel (of transmission range $r/2$) to its AP. In such a case the AP generates a *PartitionMsg*, to indicate to the MS that it is in a partition, and thus cannot utilize the network. On receiving the *PartitionMsg*, the MS deregisters from the AP and tries to use the nearest BS, so that connectivity is not lost.

3. USER BEHAVIOUR

A mobile station (MS) often finds itself in the coverage of multiple networks at the same time – for example, that of a cellular network and that of a Wi-Fi hotspot. The MS can choose to connect to any of the available networks. The behaviour of the MS is driven by its resource requirements and user interests. For example, an MS engaged in a multimedia transmission will have different requirements from one which is just downloading email from a server. We associate each MS with a user profile [7] that reflects its requirements. This user profile in turn determines how the MS chooses the cellular BS or Wi-Fi AP it connects to. The following are three basic user profiles a mobile station may possess.

1. **CLASS 1 – Bandwidth Conscious User Profile**: The MSs with a Bandwidth Conscious user profile will choose to connect to the BS or Wi-Fi AP which offers the maximum bandwidth. An estimate of the free bandwidth available is sent along with each beacon packet periodically originated by the APs and the BSs. A CLASS 1 MS on receiving such a beacon determines to switch to the new BS or Wi-Fi AP if the bandwidth advertised is greater than the free bandwidth estimate at its currently registered BS or Wi-Fi AP by a threshold value. MSs with high bandwidth requirements (like those engaged in a multimedia download) possess this type of user profile.

2. **CLASS 2 – Cost Conscious User Profile**: An MS with a Cost Conscious user profile always chooses to connect to the network with the lowest transmission cost per byte. Each BS and Wi-Fi AP advertises its associated transmission cost in the beacons sent by it. A CLASS 2 MS will switch to a new network only if the cost of the new network is less than that of its currently registered network by a threshold value. An MS engaged in non-real time file downloads can possess this user profile.

3. **CLASS 3 – Glitch Conscious User Profile**: A Glitch Conscious MS has glitch free connectivity as its priority. We define a glitch as an interruption in the transmission or connectivity which occurs when an MS moves from one network to another. Thus an MS with this user profile tries to minimize the number of handoffs it undergoes between different networks to achieve the smoothest possible transmission. This is done by remaining connected with the cellular network, which has a larger coverage area, at all possible times. An MS engaged in a voice call may use this profile.

In all the three different user profile classes, we assume that maintaining connectivity is of utmost importance to the MS. In order to maintain connectivity, an MS may connect to a network whose parameters go against its user profile. For example, a Cost Conscious MS may connect to a higher cost network when it falls outside the coverage area of the low cost network to which it is currently connected.

The user profile of an MS affects MS's behaviour, the resource consumption of the network and the traffic patterns, as it moves across the terrain. This behaviour is illustrated in Figure 2, in which an MS moves from point A to point E along the dotted line shown. The total bandwidth available at the base station BS1 is 1 Mbps while that at the two access points AP1 and AP2, are 11 Mbps each. In the scenario depicted here, we assume that the free bandwidth available at AP1 is much less than that at either BS1 or AP2 due to a large number of MSs registered to AP1. The free bandwidth available at AP2 is greater than that at BS1. Also, the per byte transmission cost associated with BS1 is considered to be four times that of either AP1 or AP2. We now describe the behaviour of the MS for each of the three different user profile classes it may posses.

- **CLASS 1 MS:** The bandwidth conscious CLASS 1 MS always tries to be registered to the network offering the maximum free bandwidth. It can be seen from Figure 2 that the MS registers with the sole AP accessible to it at the beginning of its journey (Point A). At point B, the MS comes under the transmission range of BS1 also. Since BS1 has more free bandwidth than AP1, the MS will switch over to BS1 and will remain registered with it till the point D. On entering the range of AP2 at D, the MS switches over to AP2, although it is still in the range of the BS. This is because AP2 offers a higher amount of free bandwidth than BS1. The MS remains with AP2 till the end of its journey.
- **CLASS 2 MS:** The cost conscious CLASS 2 MS tries to register with the least cost network at all times. After starting its journey from point A, the MS remains registered to AP1 till the point C. It must be noted here that the MS does not switch over to BS1 after it enters BS1's transmission range at point B. This is because the transmission cost associated with BS1 is higher than that associated with AP1. At point C, the MS goes out of the range of AP1 and is thus forced to register with the higher cost BS1 in order to maintain connectivity. On reaching the point D, the MS registers with the lower cost AP2 and remains with it till the end of its journey.
- **CLASS 3 MS:** The glitch conscious CLASS 3 MS tries to minimize the number of glitches in its connection by registering with the larger range cellular network at all possible times. The MS remains registered with AP1 between points A and B. However, once it enters the range of BS1 at point B, it switches over to the BS and remains with it till the end of the journey at point E.

3.1 Impact of User Behaviour

We have studied the impact of user profiles on the network throughput, per byte access costs and number of glitches suffered by the mobile nodes through extensive simulations in GloMoSim[8]. Table 1 gives the default parameters used across all simulation experiments described in this paper, unless specifically indicated otherwise. Random waypoint mobility model was used across all simulations.

Figure 3 shows the variation in Packet Delivery Ratio (PDR) of the three classes of mobile nodes as the mobility is varied from 2 m/s to 20 m/s. We observe the surprising result that, although the bandwidth conscious nodes attain the highest PDR at low mobility, at high values of mobility the glitch conscious nodes obtain a PDR higher than that of even the bandwidth conscious users. This can be attributed to the fact the glitch conscious users remain registered to the larger coverage cellular BSs as against the bandwidth conscious users, which suffer a very large number of glitches while remaining registered to the smaller coverage but high bandwidth APs at high mobility.

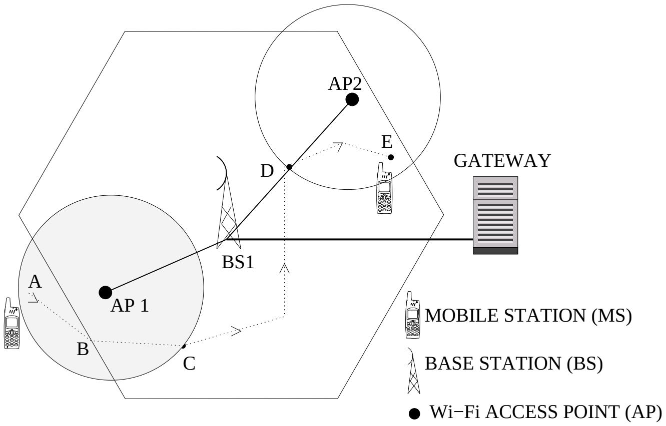

Figure 2: Behaviour of the mobile stations with different user profiles as they move across the terrain

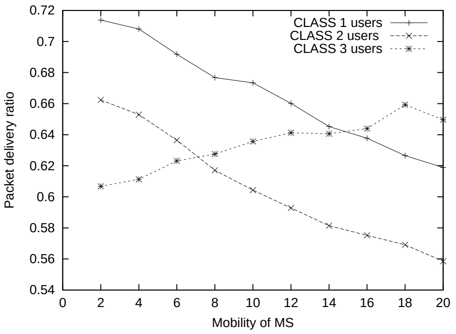

Figure 3: Packet Delivery Ratio versus Mobility of the MSs

Figure 4 shows that the glitch conscious users incur the highest per byte transmission costs. We can also observe from Figures 3 and 4 that the bandwidth conscious users incur costs comparable to or even lower than that of the cost conscious users. This leads to the important result that it is possible for the bandwidth conscious nodes to attain high PDRs at very low costs. This result follows from the fact that the Wi-Fi APs offer high bandwidth at a low price.

As shown in Figure 5, increasing mobility is also accompanied by an increasing number of glitches for all three classes of MSs. We observe that the glitch conscious nodes suffer the least number of glitches. Bandwidth conscious nodes switch every time a beacon advertising a free bandwidth higher than that of its currently registered BS or AP is received. This results in these type of nodes suffering a much larger number of glitches than the cost conscious nodes, which switch between APs and BSs only on the basis of cost. It is possible to reduce the number of glitches suffered by bandwidth conscious users by fixing a *bandwidth switch threshold*, as shown by Figure 6. However, simulations also

Table 1: Default Simulation Parameters

Parameter	Value	Parameter	Value
Terrain X range	4020m	Beacon period	1s
Terrain Y range	5220m	Bandwidth BS (Control)	1 Mbps
Number of cells	11	Bandwidth BS (Data)	1 Mbps
Cell radius	1km	Bandwidth AP (Control)	1 Mbps
Transmission range BS	1km	Bandwidth AP (Data)	11 Mbps
Transmission range AP (Control)	250m	Transmission range AP (Data)	125m
Transmission cost per byte (BS)	4 units	Number of MSs	600
Transmission cost per byte (AP)	1 unit	Mobility of the MSs	10 m/s
Mean inter-packet arrival time	2000	Number of APs	40

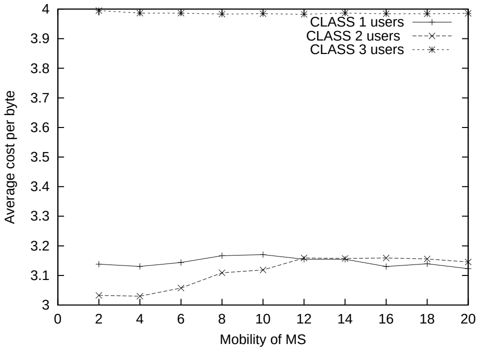

Figure 4: Average Per Byte Access Cost vs Mobility of the MSs

indicate that the PDR or CLASS 1 MSs does not improve in spite of the lower number of glitches suffered. The gain in PDR due to lower number of glitches is offset by the disadvantage of missing out high free bandwidth APs or BSs due to large bandwidth switch thresholds.

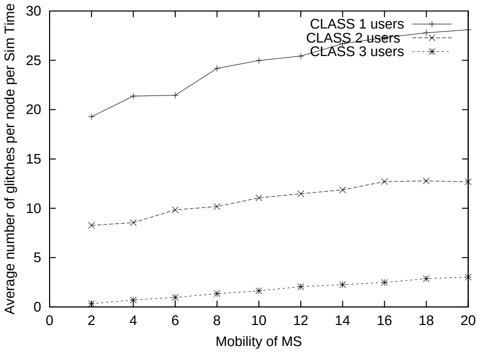

Figure 5: Average number of glitches suffered by an MS vs Mobility of the MSs

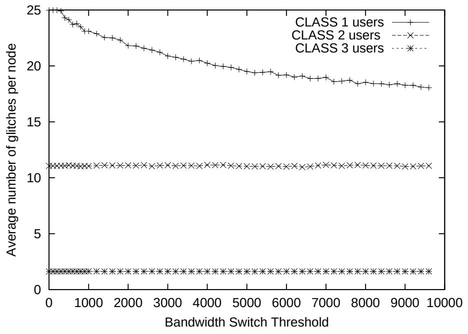

Figure 6: Average number of glitches suffered by an MS vs Bandwidth Switch Threshold

Figure 7 show that increasing the number of Wi-Fi APs in the system leads to higher PDR until a limit, after which the performance of the system degrades due to increased congestion. This is because all the Wi-Fi APs operate in the same frequency channel. Frequency planning among the APs located close together may reduce the degradation in system performance. When the number of mobile nodes present in the system is low, the highest PDR is achieved by the glitch conscious users due to the minimal load at the cellular BSs. As the number of nodes is increased, the PDRs of all three types of nodes decrease, but the highest PDR is attained by the bandwidth conscious nodes. This result is captured by Figure 8.

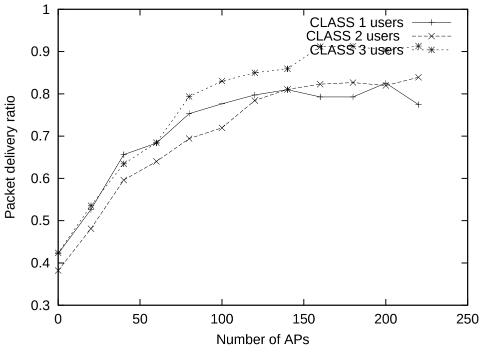

Figure 7: Packet Delivery Ratio vs Number of APs

The Wi-Fi APs are usually single hop environments. When the APs are modeled as multi-hop relaying environments, it is found (Figure 9) that glitch conscious nodes attain the highest PDR. This is due to the large number of network partitions (lack of a multi-hop path between the node and AP) suffered by the bandwidth conscious and cost conscious users when they are connected to the multi-hop Wi-Fi APs.

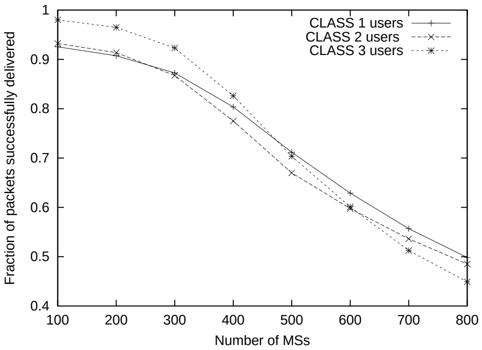

Figure 8: Packet Delivery Ratio vs Number of MSs

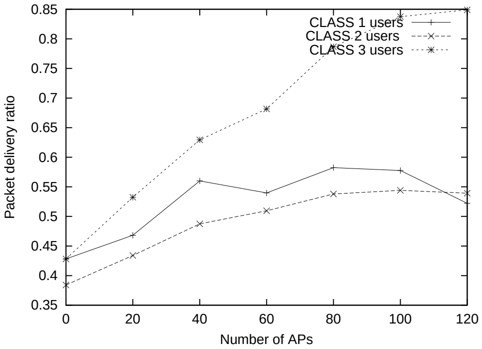

Figure 9: Packet Delivery Ratio vs Number of APs (APs in multi hop mode)

4. REVENUE SHARING

An inter-operable network setting involves interaction between multiple service providers. It is essential that the commercial viability of all the services providers involved in the system is guaranteed. Thus pricing and revenue sharing are very important aspects of an inter-operable network.

4.1 Pricing Schemes

Flat-rate pricing and *volume-based pricing* are two major pricing schemes that can be employed. Flat rate scheme is the currently existing and popular scheme. In flat-rate pricing, a user is permitted to utilize the network services for a specified period of time at a fixed price without any restrictions on the bandwidth consumed. The volume-based pricing approach charges a user based on the amount of data which he/she transacted over the network. In addition to both these schemes, business establishments can provide Wi-Fi network services for free, as a value addition to customers visiting their premises for core business activities.

A flat rate scheme is attractive for the network users (MSs) at low network loads. We can see from Figure 10 that average per byte cost incurred by the MSs increases as the number of MSs in the system is increased. In the presence of a large number of MSs, the number of packets that are successfully sent or received by the MSs decreases as a result of the increased congestion in the network. This causes the increase in the average per byte cost incurred by the MSs.

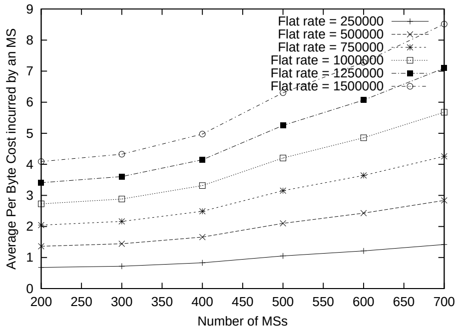

Figure 10: Average per byte cost incurred by an MS versus the number of MSs in the system at various flat rate prices.

A higher flat rate implies higher revenues for the SP as well as a higher average per byte cost for the MSs. We can see from Figure 11 that volume based pricing results in a higher revenue than flat rate schemes when the number of MSs in the system is low. But in this case, the average per byte cost incurred by the MSs in the volume based scheme is higher than that in the corresponding flat rate scheme. However at particular values of load (for example with 450 MSs in the system), flat rate schemes (for example, flat rate $= 10^7$) were found to generate both higher revenues for the service providers, as well as charges a lower average per byte cost to the MSs, than volume based schemes.

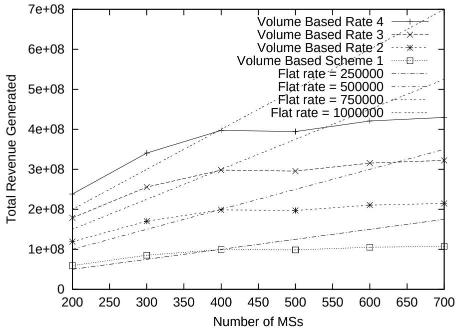

Figure 11: Comparison of total network revenue under various flat rate and volume based schemes (per byte rate = 4, 3, 2, and 1)

4.2 Revenue Sharing

Revenue sharing models describe the way in which money paid by the network users is split among the various service providers (SPs). The revenue sharing scheme used in a particular system directly influences the revenue obtained by the different service providers, and in turn determines their profitability and commercial viability. In a *Fixed-fraction sharing model*, the total revenue to be shared among the service providers is fixed apriori. In a *Volume-based sharing model*, the revenue obtained by a particular service provider depends on the volume of data transacted by that service provider.

In a simple volume based revenue sharing scheme, the WAN SP and Billing Agency treat the Wi-Fi AP as an intermediate node that participates in the forwarding of data. The Wi-Fi AP is reimbursed the amount β by the WAN SP for every byte of data it has forwarded. This scheme is illustrated in Figure 12. The value of β has to be decided in such a way that the WAN SP's revenue does not fall below a minimum threshold, and at the same time, the Wi-Fi SPs also get a fair share of the generated income. Since this model pays the APs on the basis of the traffic that they have transmitted instead of equal sharing among the APs, it can also be used when each AP is operated by a different Wi-Fi SP.

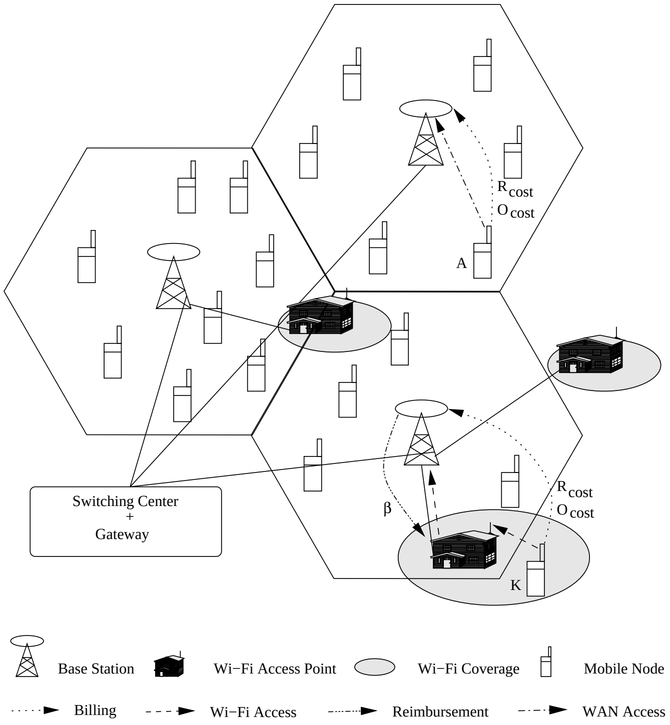

Figure 12: Illustration of a simple revenue-sharing scheme

We compare the following three schemes for fixing the value of β:

- **Constant Beta**

  This is the trivial scheme in which the Wi-Fi AP receives the constant amount β for each byte of data forwarded by it.

- **Level-based Beta**

  In this scheme, the value of β varies in a step-like fashion with the total number of bytes forwarded. For example, $\beta = b_1$ for the first x bytes of data forwarded, $\beta = b_2$ for the next y bytes and $\beta = b_3$ for the remaining.

- **Continuous Function Beta** This scheme is a variation of the Level-based Beta Scheme in which the value of β varies continuously for each byte of forwarded data, *i.e.,* $\beta = f(t)$ for the t^{th} byte of data forwarded. $f(t)$ is commonly taken to be a negative exponential function.

Figure 13 shows illustrates the above described β schemes. Beta schemes 1 and 2 represent the Constant β schemes with the value of the constant equal to 2.5 units and 1.5 units respectively. Scheme 3 is Level-based, while Scheme 4 represents the Continuous Function scheme $f(t) = 3.9e^{-10^{-5}t}$, where t is the total number of bytes transacted through the Wi-Fi AP.

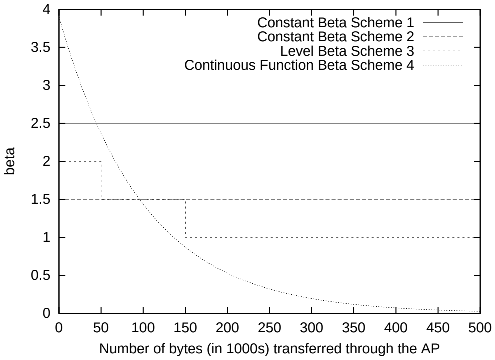

Figure 13: Variation of β with total number of bytes transferred in the various β schemes

Figure 14 shows that, among all the four schemes, Scheme 1 generates the maximum revenue for the APs. However, this scheme also results in the lowest revenue for the BSs, as seen in Figure 15. We can also observe that schemes 2, 3, and 4 generate approximately the same amount of revenue for the BSs. However in Scheme 2, at low number of MSs in the system, the revenue reimbursed to the APs is very low. This revenue may not be sufficient to keep the operation of the AP commercially viable. The advantage of the level based and continuous function schemes is that the APs are reimbursed sufficient revenue to maintain commercial viability even at low loads without decreasing the BSs' revenue by a large amount, as is evident from Figure 14.

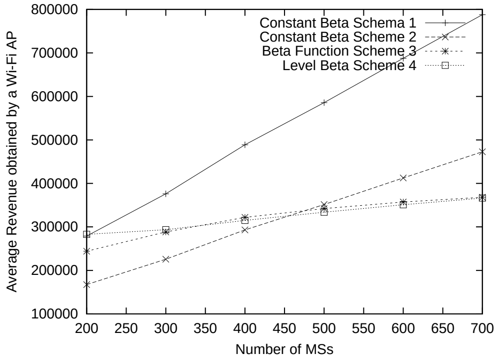

Figure 14: Effect of the different β schemes on the average revenue earned by an AP

4.3 Balancing Load by Dynamic Pricing

Differentiated network access costs provide an opportunity to balance the load across the various APs and BSs. Dynamically increasing the cost associated with a BS or an

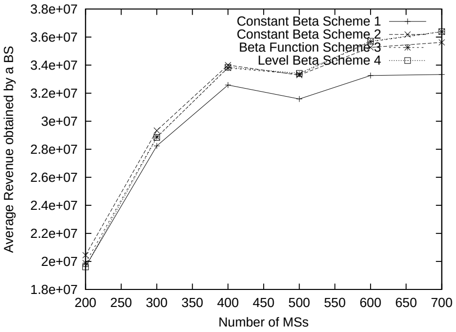

Figure 15: Effect of the different β schemes on the average revenue earned by a BS

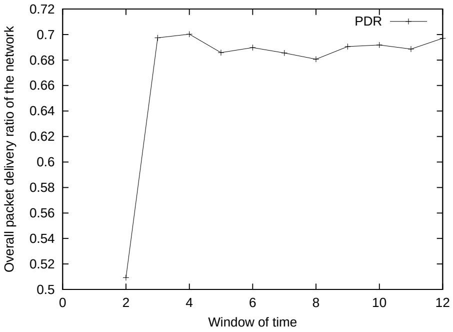

Figure 16: Variation in Packet Delivery Ratio with simulation time

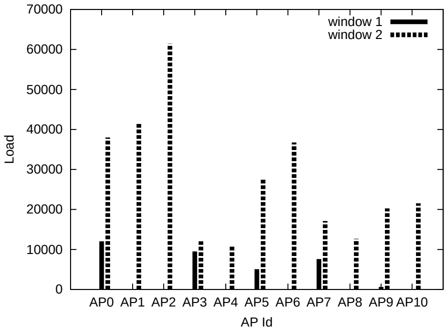

Figure 17: Impact of local load balancing on the load at APs surrounding a particular BS

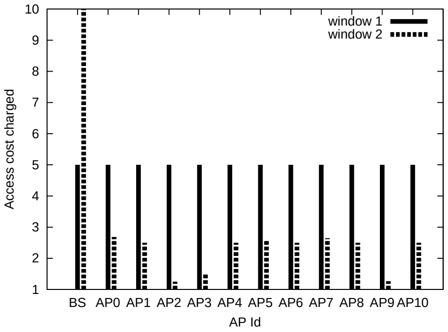

Figure 18: Impact of local load balancing on the access cost charged by APs surrounding a particular BS

AP will cause MSs with the cost conscious user profile defined in Section 3 to switch away to lightly loaded (and hence lower cost) APs or BSs in their vicinity. This balancing of load results in a higher network throughput. This scheme also has the added advantage of decreasing the average per byte access cost incurred by the cost conscious MSs.

Global Average Balancing and *Local Average Balancing* are two simple load balancing schemes. In the Global Average Balancing scheme, the cost associated with an AP or BS is changed depending on the deviation of its load from the average load across all APs and BSs in the system. In the case of Local Average Balancing, the average load is calculated only in the vicinity of the AP or BS whose cost is to be updated.

Simulation studies have shown that dynamic pricing based load balancing leads to a higher Packet Delivery Ratio (PDR) for the network. Initially all BSs and APs charge the same access cost of say 5 units. Local average load balancing is periodically applied at the end of windows of 30 seconds duration. Figure 16 shows the marked jump in PDR at time window 2 where balancing is first applied. Thereafter, the PDR shows little variation as load balancing is in effect. Figure 17 shows that the load at the APs surrounding a particular BS has increased after load balancing is applied. Load balancing has resulted in all APs within the particular cell to share a part of the load, which was initially almost wholly concentrated at the BS. The percentage of the total load in the cell handled by the BS reduced from over 98% to approximately 85% after load balancing. Due to the great difference in coverage areas, the load on the BS still remains much higher than that of any of the surrounding APs. Figure 18 shows the change in access cost charged by the APs and BSs after the application of load balancing. We can see that the BS has the highest access cost due to the higher load encountered by it.

5. IMPLEMENTATION ISSUES

Implementing smooth interoperability between wireless LANs and wide area cellular networks is a challenging task. It calls for modifications in the network stacks of the mobile devices' operating systems as well as requires additional features in the WLAN and the cellular network protocols. The applications running on the mobile node may or may not be aware of the dynamic switching between the various network interfaces taking place at the lower layers of the network stack. Switching decisions are based on the mobile node's *User Profile* and the current network conditions. This also calls for devising methods to characterize and measure the current network conditions. The implementation of interoperability in the mobile devices may be done at multiple levels of the network stack – application, transport, network layers. Each layer has its advantages and disadvantages in terms of ease of implementation, speed, efficiency, customizability and extendability.

5.1 Modifications to Support User Profile Based Interoperability in WLANs and Cellular Networks

Almost all currently deployed WLANs are based on IEEE 802.11. Existing wide area cellular networks use GSM or CDMA technologies. All these technologies consist of elaborate beaconing and signaling mechanisms to support the transmission of voice and data traffic over the networks. However as the protocols were never designed with interoperability in mind, they lack certain features which are essential to support user profile based interoperability. Some of the important pieces of information required for interoperability are discussed below.

5.1.1 *Current Available Bandwidth*

User profile based switching requires knowledge of the throughput attainable in the different networks at any instant. The current network load, the total bandwidth available at each BS or AP, the bandwidth partitioning algorithm, which are required to calculate the current available bandwidth, are all details internal to a particular network and are not available to a user of the network. Periodic beacons containing the free bandwidth information should be broadcasted by the BSs (in CDMA, GSM) and the Access Points (802.11). The dynamic network switching modules in the mobile devices' network stacks extract the bandwidth information from these beacons and make the appropriate switching decisions.

5.1.2 *Access Cost Information*

To support the *Cost Conscious User Profile*, the network switching module needs to know the access charges associated with the different networks. When the cost is constant across all BSs and APs of a particular network, the fixed value may be embedded as it is into the switching module. In the case of dynamic pricing – *i.e.,* the access cost associated with an BS/AP changes with geographic location, time, network load, etc – the access cost information needs to be conveyed to the mobile nodes through periodically broadcasted beacons, as in the case of free bandwidth information mentioned in the previous section.

5.1.3 *Network Coverage Information*

The dynamic switching module uses knowledge about the coverage areas of different BSs/APs while making switching decisions for the *Glitch Conscious* users. As in the case of variable access costs, variable coverage areas may be advertised through periodic beacons.

5.2 Design of a Dynamic Network Switching Module

In this section, we describe the design and implementation of a user profile based dynamic network switching module. The function of this module is to take a decision on which network interface is to be used currently. The module reads in the user profile data from its configuration files and periodically queries the various network interfaces to obtain their current status. The structure of the module is shown in Figure 19. This module can be plugged in at multiple layers of the operating system's network stack. The decision to use a particular interface is made by the module, but the mechanism to implement switching will be different in different layers.

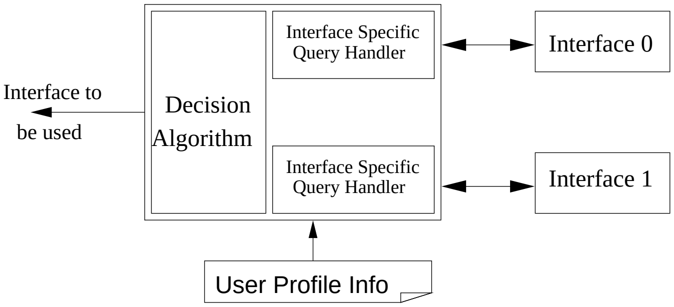

Figure 19: Structure of the User Profile based Dynamic Network Switching Module

5.3 Incorporating Interoperability into the Network Protocol Stack

Interoperability can be incorporated at multiple levels in the network protocol stack of the mobile device's operating system, as described below.

5.3.1 *MAC Layer*

A-MAC, proposed as part of the AdaptNet protocol suite [9], is a two-layered MAC which can support interoperability. The master sublayer of A-MAC forwards data to the multiple network interfaces on the basis of the *virtual cube* concept. The virtual cube based decision making module can be replaced or supplemented by the dynamic switching module described in Section 5.2 to support user profile based interoperability.

5.3.2 *Network Layer*

Maintaining a single static IP address for the mobile device across all network interfaces is the main challenge in implementing interoperability at the network layer. A single static IP address for a mobile device can be achieved through Mobile IP [10]. The mobile device is considered to have moved away to a foreign network whenever it switches to a new network interface. This solution requires the Wi-Fi APs and the cellular BSs to support the Home Agents and Foreign Agents of the Mobile IP protocol.

5.3.3 *Transport Layer*

Implementation of interoperability at the transport layer involves the major challenge of maintaining TCP connections across switching between multiple interfaces. A TCP connection is characterized by the 4-tuple *(Source IP, Source Port, Destination IP, Destination Port)*. As different interfaces have different IP addresses, switching the network interface will break the TCP connection. Redirectable Sockets, *RedSocks*, introduced in [11] is a possible solution to this problem. *pTCP* [12], which achieves bandwidth aggregation by striping data across multiple TCP connections, is another protocol by which smooth interoperability can be attained at the transport layer.

5.3.4 *Application Layer*

Applications may be written such that they are aware of the multiple network interfaces, and may incorporate application specific considerations into the switching decisions. All calls to open a socket in the application must be modified so as to attach the socket to the interface specified by the switching module rather than to the default interface. In Linux, this involves appropriately setting the SO_BINDTODEVICE socket option. While application specific optimizations are possible, the disadvantage of implementing interoperability at the application layer is that all applications using interoperability are to be modified or rewritten from scratch.

5.4 *interoproxy* - An Application Layer Implementation

interoproxy, is an application layer implementation of interoperability. This interoperable HTTP proxy uses the dynamic network switching module for choosing the interface over which the HTTP requests are to be made, on the basis of the user profile and the current network conditions. An innovative HTTP request replay mechanism is also added to the proxy to take care of requests lost during switching from one interface to another. We must note here that *interoproxy* runs locally on the same machine as the Internet browser and is used even when a direct connection to the Internet is available.

interoproxy was implemented by extending *tinyproxy*[13], a lightweight HTTP proxy licensed under the GNU Public License (GPL), to support user profile based dynamic switching between multiple network interfaces. Modifications to the source code were limited to only the places where sockets were opened to relay the web browser's request to the Internet or to an upstream proxy. The tinyproxy `opensock` function was modified to bind the socket to the interface specified by the interface selection logic using the Linux `SO_BINDTODEVICE` socket option.

5.5 Smooth Interoperability through Request Replay

Smooth interoperability requires that the user remains unaware of the dynamic switching taking place between the multiple network interfaces. Consider the following situation. A user requests a web page through his web browser. *interoproxy* opens a TCP connection to the web server over the CDMA interface. Before the HTTP reply was received, the CDMA interface got disconnected and the proxy switched to the WLAN interface. In such a case, the proxy will automatically replay the requests sent out on the broken interface for which no reply was received, on the new interface. Thus the user is unaware of the switching and gets the desired web page. In some cases, the switching may take place even though the current interface is not completely disconnected. In such a case, the proxy may decide to wait for replies to already placed requests on the old interface or it may choose to terminate the active HTTP sessions and replay the requests on the newly active interface.

6. CONCLUSION

Interoperability between wireless LANs and wide area cellular networks is a very challenging task. Mobile devices of the future, equipped with multiple network interfaces, will dynamically switch between the interfaces on the basis of their user profile.

Simulation studies have shown that user profile based switching leads to higher network performance and increased user satisfaction. Pricing and revenue sharing are extremely important in an interoperable scenario as the commercial viability of all the service providers involved must be assured. Various pricing schemes and revenue sharing schemes were analyzed in this paper. It was also shown that load balancing between the cellular BSs and Wi-Fi APs on the basis of dynamic pricing leads to an improvement in the overall network performance.

7. REFERENCES

[1] GSM World, `http://www.gsmworld.com`.

[2] CDMA Development Group, `http://www.cdg.org`.

[3] V. Sekar, B. S. Manoj, and C. Siva Ram Murthy, "Routing for a Single Interface MCN Architecture and Pricing Schemes for Data Traffic in Multihop Cellular Networks" in *Proc. IEEE ICC 2003*, pp. 969-973, May 2003.

[4] Y. D. Lin and Y. C. Hsu, "Multihop Cellular: A New Architecture for Wireless Communications," in *Proc. IEEE INFOCOM 2000*, pp. 1273-1282, March 2000.

[5] R Ananthapadmanabha, B. S. Manoj, and C. Siva Ram Murthy, "Multihop Cellular Networks: The Architecture and Routing Protocols", in *Proc. IEEE PIMRC 2001*, pp. 78-83, October 2001.

[6] D. B. Johnson and D. A. Maltz, "Dynamic Source Routing in Ad hoc Wireless Networks", *Mobile Computing, Kluwer Academic Publishers*, vol. 353, pp. 153-181, 1996.

[7] Dilip Antony Joseph, B.S. Manoj, and C. Siva Ram Murthy, "The Interoperability of Wi-Fi Hotspots and Packet Cellular Networks and the Impact of User Behaviour", to appear in *Proc. IEEE PIMRC 2004*, September 2004.

[8] X. Zeng, R. Bagrodia, and M. Gerla, "GloMoSim: A Library for Parallel Simulation of Large-scale Wireless Networks," in *Proc. PADS-98*, Banff, Canada, May 1998.

[9] I. Akyildiz, Y. Altunbasak, F. Fekri, and R. Sivakumar, "AdaptNet: An Adaptive Protocol Suite for the Next-Generation Wireless Internet", IEEE Communications Magazine, vol. 42, no. 3, pp. 128-136, March 2004.

[10] C. Perkins, "Mobile IP", IEEE Communications Magazine, vol. 35, no. 5, pp. 84-86, March 1997.

[11] M. Haungs, R. Pandey, E. Barr, J. F. Barnes, "A Fast Connection-Time Redirection Mechanism for Internet Application Scalability", in *Proc. HiPC 2002*, pp. 209-218, December 2002.

[12] H. Y. Hsieh and R. Sivakumar, "A Transport Layer Approach for Achieving Aggregate Bandwidths on Multi-homed Mobile Hosts". in *Proc. ACM MOBICOM 2002*, pp. 83-94, September 2002.

[13] tinyproxy - A Lightweight HTTP proxy, `http://tinyproxy.sourceforge.net`