NEXT GENERATION WIRELESS NETWORKS

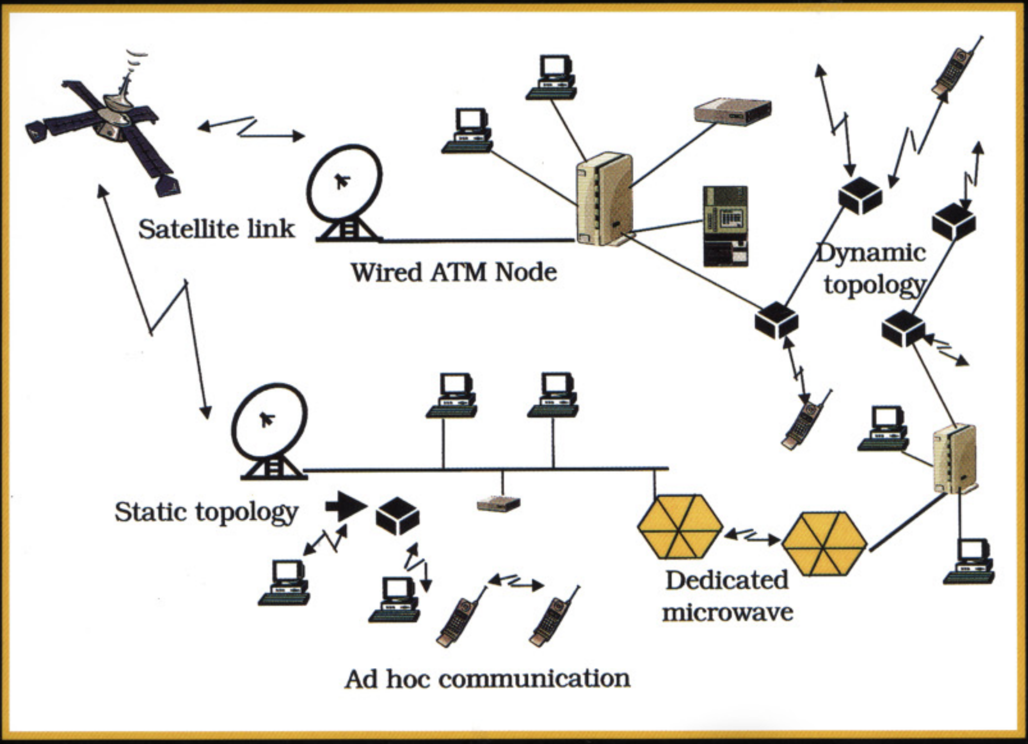

SIRIN TEKINAY
(EDITOR)

NEXT GENERATION WIRELESS NETWORKS

THE KLUWER INTERNATIONAL SERIES
IN ENGINEERING AND COMPUTER SCIENCE

NEXT GENERATION WIRELESS NETWORKS

edited by

Sirin Tekinay
New Jersey Institute of Technology

KLUWER ACADEMIC PUBLISHERS
NEW YORK, BOSTON, DORDRECHT, LONDON, MOSCOW

eBook ISBN: 0-306-47310-0
Print ISBN: 0-792-37240-9

Created in the United States of America

Visit Kluwer Online at: http://www.kluweronline.com
and Kluwer's eBookstore at: http://www.ebooks.kluweronline.com

Contents

Introduction

This book is a collection of extended versions of the papers presented at the Symposium on Next Generation Wireless Networks, May 26, 2000, New Jersey Institute of Technology, Newark, NJ. Each chapter includes, in addition to technical contributions, a tutorial of the corresponding area. It has been a privilege to bring together these contributions from researchers on the leading edge of the field. The papers were submitted in response to a call for papers aiming to concentrate on the applications and services for the "next generation," deliberately omitting the numeric reference so that the authors' vision of the future would not be limited by the definitive requirements of a particular set of standards. The book, as a result, reflects the top-down approach by focusing on enabling technologies for the applications and services that are the defining essentials for future wireless networks. This approach strikes a balance between the academia and the industry by addressing new wireless network architectures enabling mobility and location enhanced applications and services that will give wireless systems the competitive edge over others.

The main theme of the book is the advent of wireless networks as an irreplaceable means of global communication as opposed to a mere substitute for, or a competitor of, wireline networks. Geolocation emerges as the facilitator of mobility and location sensitive services. The fields of geolocation and wireless communications have been forced to merge, following the Federal Commission of Communications' (FCC) ruling that obliges wireless providers with emergency caller geolocation. This initial driving force has quickly been augmented by the already existing and increasing popularity of positioning and navigation systems using the Global Positioning System (GPS), in addition to the wireless providers' zeal to add value to the geolocation capability. The result is the currently experienced evolution of wireless networks where mobile location is a natural aid to network management, to a variety of applications and value added services. At this time, the path of the evolution is not clearly focused, neither is the winning set of applications obvious. What we do know is that next generation wireless networks will continue to change the way we live.

The first part of the book contains tutorials on three network architectures that aim to achieve this vision. The first chapter, by Ana Lúcia lacono and Christopher Rose of WINLAB, Rutgers University, presents the concept of "Infostations," that arise from the tradeoff between the size of the radio coverage area of a single transceiver and the feasible information rate. Infostations favor the latter, making use of mobility of users in redeeming the selective patchy coverage pattern. The "anytime, anywhere" motto of PCS is replaced by "manytime, manywhere" access in Infostations. The second chapter by Abbas Jamalipour of the University of Sydney describes the role of satellites in broadband wireless access. In the third chapter, Patricia Morreale of Stevens Institute of Technology describes the “Infocity” concept, which is based on the integration of wireless and wireline networks in order to provide the envisaged high- speed ubiquitous access to information and communication.

The second part of the book includes contributions that describe the state of the art wireless geolocation systems and trends. The first chapter is co-authored by scientists from Bell Labs of Lucent Technologies. Bob Richton, Giovanni Vanucci, and Steven Wilkus depict the widely accepted standard solution to the wireless geolocation problem; i.e., “assisted GPS.” This concept links wireless networks with GPS in order to reach the accuracy and availability requirements for the user geolocation information. The wireless network assumes a supporting role in geolocation, in order to aid the end user equipment through its prescribed communication with GPS. The second contribution, also authored by a Bell Labs scientist, Oguz Sunay, provides a tutorial on all alternatives for wireless geolocation and details the evaluation procedures that are currently under research by standards bodies and relevant work groups. In the third contribution, Malcolm Macnaughton, Craig Scott and Chris Drane, researchers from the University of Technology, Sydney, present the chronicle of research efforts and empirical data collection on the geolocation capability of the existing wireless infrastructure, during which they truly bring theory and practice together in efforts to sharpen the geolocation capability of the wireless system.

The third part presents contributions that demonstrate the use of location information in next generation wireless network applications and services. The first contribution, by Mostafa Bassiouni and Wei Cui of the University Central Florida focuses on the use of real time geolocation measurements in improving mobile connectivity and enhancing the effectiveness of fault recovery in configurable cellular networks. The last contribution, by Guenther Popischil, Alexander Schneider, and Ernst Bonek of Technischen Universität Wien, portrays the creation of the “killer application,” for next generation wireless networks.

I am proud to have put together this volume comprising of chapters by contributors who are among the elite that are making the future happen. I thank Dr. Oguz Sunay for his invaluable, tireless efforts in ensuring the technical flow and cohesiveness of the book. I would also like to thank Alex Greene, the Publisher of this book from Kluwer Academic Publishers, for his patient, capable help. Finally, I’d like to express my gratitude to my brilliant research associate Mr. Amer Chatovich, whose careful, meticulous pursuance has made this project possible.

Sirin Tekinay

Chapter 1

INFOSTATIONS: NEW PERSPECTIVES ON WIRELESS DATA NETWORKS

Ana Lúcia Iacono
Christopher Rose
WINLAB - Wireless Information Network Laboratory
Department of Electrical and Computer Engineering
Rutgers, The State University of New Jersey

Abstract We discuss file delivery issues for a new approach to inexpensive, high rate wireless data called *Infostations.* As opposed to ubiquitous coverage, infostations offer geographically intermittent coverage at high speed (1Mbps to 1Gbps) since data, as compared to voice, can often tolerate significant delay. The infostations paradigm flips the usual "slow-radio/fast-network" scenario upside down and offers intriguing new design problems for wireless data networks. Collectively, we at WINLAB believe that the infostations scenario, especially with the emergence of the World Wide Web as both a communications medium and defacto standard is one way to obtain low cost wireless data. And perhaps controversially, we offer arguments that currently proposed extensions to cellular systems (such as the coming Third Generation) will not be able to offer data as inexpensively. In this chapter we describe the infostations concept and then concentrate on issues above the physical layer. Specifically, we worry about delay bounds on information delivery for variety of simple user mobility scenarios and infostation geometries. We then provide heuristic algorithms which closely approach these bounds.

Keywords: wireless communications, mobile computing, wireless data, wireless internet, scheduling algorithms, delay bounds, mobility management

1. INTRODUCTION

Over the past 10 years, wireless voice communication has grown from a rarity to a necessity. In contrast, wireless data services at rates and price sufficient to generate equal excitement remain elusive. In response, the wireless industry has proposed third generation systems with rates in the hundreds of kilobit per second range. However, the dominant traffic on such systems will probably be voice at least initially, and here lies a "Catch-22" first observed by Roy Yates here at WINLAB [1].

Consider that the bit rate currently associated with voice communications is on the order of 10 kbps and let us use this voice channel rate as our unit of measure. This channel costs v cents per minute. It therefore costs approximately $13v$ cents to transmit a one megabyte file – prohibitively expensive at current rates. In addition, what is particularly interesting is that this basic fact does not change with the introduction of higher rate services as long as voice is the dominant traffic. One megabyte of data *always* costs $13v$ cents since the basic voice channel rate is unlikely to change drastically for both economic and legacy reasons. Thus, unless normal voice communications becomes essentially free, it seems that wireless data will never be inexpensive when provided using a cellular architecture.

This conundrum causes us to re-examine the cellular paradigm. Specifically, cellular wireless was built to carry voice traffic for people accustomed to the reliability and ubiquity of fixed telephone service. Thus, the goal of the cellular industry was coverage anytime and anywhere. However, to provide large coverage the system must be designed so that users both near and far from the access point (a base station) achieve some minimum quality of service. From a systems perspective however, it would be more efficient to serve users closer to the base station at higher rate, be done with them and then serve users farther away. However, for voice systems the implication is intermittent coverage which is incompatible with continuous interactive traffic such as voice.

In contrast, data can tolerate delay and the system throughput could be increased by offering rates commensurate with achievable signal to interference ratios. Add to this that customers are often in motion the basic (somewhat surprising) infostations design precept emerges for single non-dispersive, non-directional channels:

Infostations should not be shared between users

That is, at any point in time, only one user should be attached to an infostation. This basic idea has roots in information theory and water-filling of channels in space, time and frequency (see [2] for a development on multiple user dispersive channels). If we consider different frequency or spatial sub-channels, then the precept still holds if each sub-channel is considered to be an infostation unto itself and users attempt to use the "infostation(s)" with the best channel (s), even though these infostations might be co-located.

A possibly non-obvious consequence of such spatio-temporal-frequency water-filling results in another defining characteristic of the infostations paradigm. For users in (ergodic) motion, the places at which transmissions should occur are where the channel quality is above some threshold – a result first shown by Joan Borras [3, 4] and based in part on work by Andrea Goldsmith [5] on fading channels. This implies that a user traveling with uniform velocity in an isotropic environment should transmit or receive only when it is close to an infostation, and from this the notion that

Infostation coverage areas are spatially discontinuous

emerges naturally.

Thus, we define infostations as a wireless communication system characterized by sequential user access with discontinuous coverage areas and high data rate transmissions. As opposed to the moderate rate ubiquitous coverage in cellular systems, infostations offer high speed discontinuous coverage which may be accessed by users in transiently close proximity to an infostation, and in fact can maximize system capacity. Furthermore, the removal of the need to coordinate channels among multiple users and over the system as a whole should lead to simple inexpensive realizations. And owing to the bursty nature of data communications and its tolerance of moderate delay, the infostation scenario with its inherently lower associated costs might be an attractive alternative to the classical concept of anytime anywhere communications networks.

1.1 EXAMPLES

Although not specifically an infostation, consider a system introduced by Apple, called Airport [6]: a base station that costs $299 and wireless networking cards that cost $99 enable up to 10 computers to share a 11 Mbits/second Internet connection at distances up to 150 feet. As Peter Lewis describes in his article in The New York Times [7]:

"That is so important, and it has such potential to change the way we use computers and information appliances around the house, that I'm compelled to repeat it in a different way: I'm sitting outside the house on the deck, with an iBook on my lap, enjoying a glorious autumn day, reading the current e-news, checking e-mail … There are no wires,

cables or extension cords in sight. As the stars come out, I simply stroll back into the house and continue working from the sofa in the living room."

The salient feature of this narrative is a perceived desire for manytime manywhere web access as opposed to the more traditional anytime anywhere access we expect for voice services. Note in particular that Lewis did not suggest he was *using* the computer during his journey from deck to sofa.

There are a variety of possible infostation system architectures. For example, many infostations may be owned by a single company and they may be clustered and connected to cluster controllers according to their location, creating a hierarchical architecture, as shown in Figure 1.1. This is somewhat analogous to the large telephone company cellular systems where many base stations are connected to mobile telephone switching offices through dedicated high speed lines.

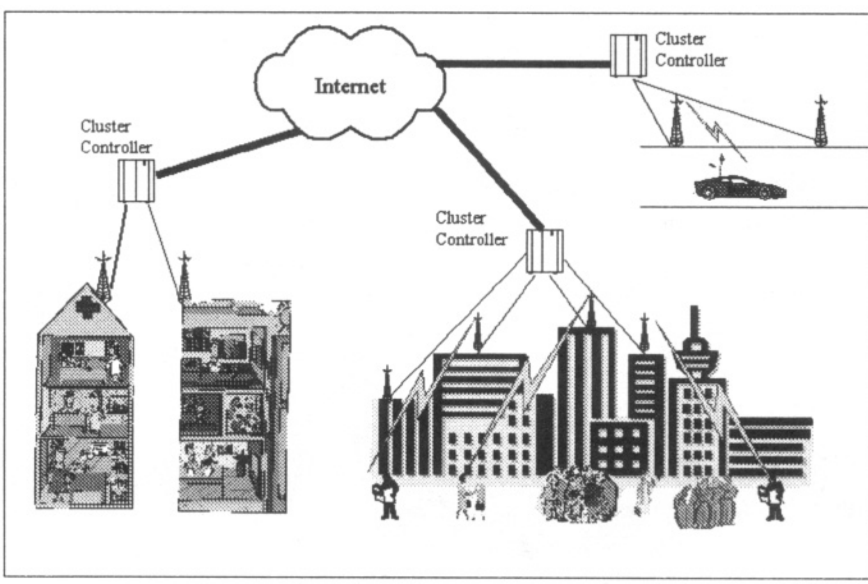

Figure 1.1 Cluster of Infostations

Another possible scenario might have small businesses such as convenience stores carry infostation service as a sideline – analogous to lottery sales agents. This architecture is shown in Figure 1.2. To be economically attractive, the start up cost to such a "Mom and Pop" operator should be low. There could also be a mixture of the two, as in a franchise setting where infostation operators leased the infrastructure from the founding company.

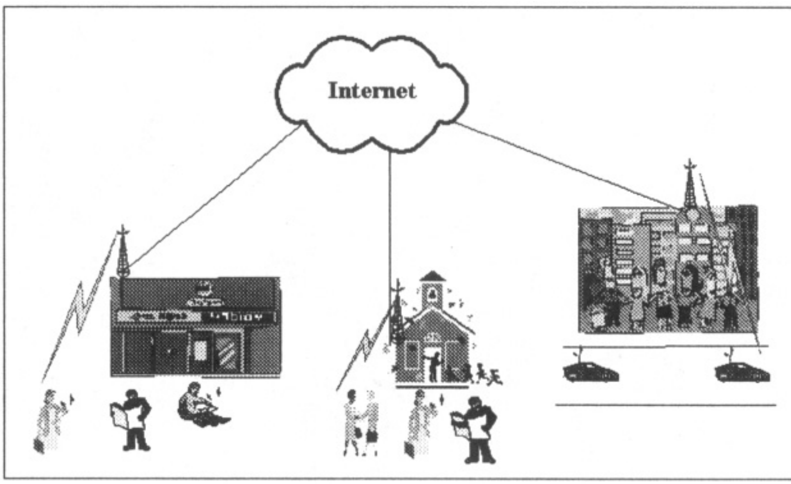

Figure 1.2 Independent Infostations

The network could also be isolated from the Internet and could be used for local communications, as in an office building or home. This is shown in Figure 1.3. Yet another architecture is to integrate infostations with a ubiquitous, low data-rate system (e.g. CDPD or other [8]) and use them as bandwidth boosters. Figure 1.4 shows one example of a hybrid infostation network. According to where infostations are placed, the user mobility can be characterized by three situations [9]: mobile users moving with high speed, such as in a highway, characterize what is called a "drive-through" scenario; users with medium speed, such as in a sidewalk or a mall, characterize a "walk-through" scenario; finally stationary users, such as in an airport lounge or a classroom, characterize "sit-through" scenario.

At WINLAB we have been studying several different problems related to an infostation network. A study of the infostation system performance in terms of capacity, throughput and delay was presented in [4] where various models and different power allocation, symbol rate adaptation

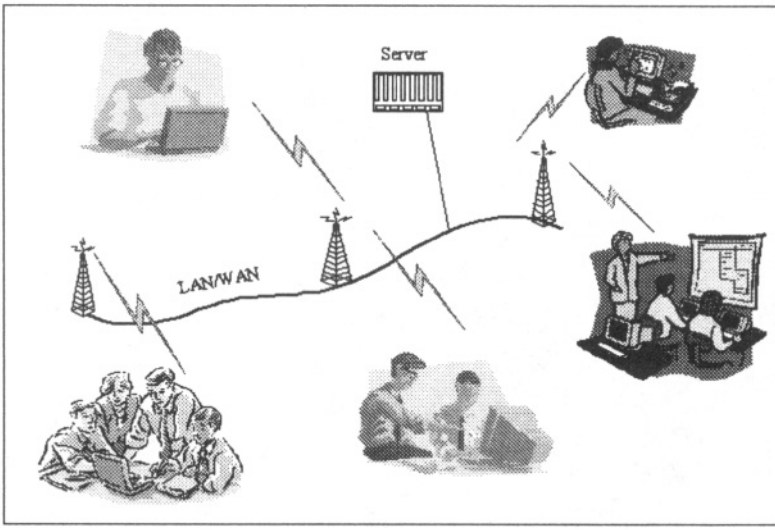

Figure 1.3 Isolated Infostation System

and modulation schemes are presented. A medium access scheme called WINMAC has been proposed in [10] to support efficient packet communications between an infostation and mobile terminals. This protocol adapts to the radio channel condition and achieves enhanced communication reliability through packet retransmission and data rate adjustment. Due to the fact that one of the main services that infostation will provide is Internet access, another area of interest is the design of a link layer protocol to transmit IP packets efficiently via the wireless link. An error control scheme for the Radio Link Protocol is proposed in [11]. The scheme uses multicopy and error threshold detection to improve the system performance. Infostation operation issues such as registration, authentication and billing are addressed in [12]. Some radio design issues are examined in [13]. There is also a variety of other work both at WINLAB and elsewhere ranging from physical layer issues up through applications [3, 12, 14, 15, 16, 4, 17, 18, 19, 1, 8].

1.2 USER MOBILITY AND INFOSTATIONS

One might wish to place an infostations system in an airport lounge, in a conference room or in a small office at an affordable price. One

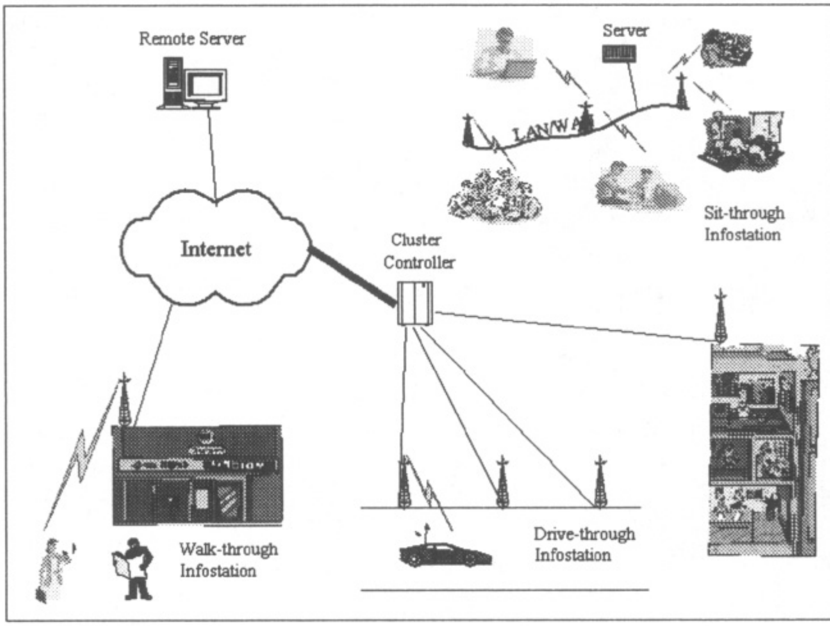

Figure 1.4 The Infostation Network.

common characteristic of these situations is relatively low user mobility. Although the coverage area is small, the system is designed based on the fact that the user will stay in the coverage area during the time of the connection and in fact, using beam steering techniques one might "move the infostation" as opposed to moving the user. However, from the perspective of the fixed network, users are in relatively fixed locations.

However, when designing system where users roam, then user mobility must be considered since users may visit several infostations during a single connection. For example, a user might roam over a shopping mall with stores offering local infostation services and a user would not stay connected to a single infostation while shopping. Likewise on a highway with infostations at regular intervals users might traverse great distances (from the fixed network perspective) between infostation contacts.

Now consider that data communication, such as messaging systems or web applications, is inherently asymmetric with much greater volume occurring on the downlink from network to user. Under this scenario, if the information is available at the infostation, then the main issue is to send it to the mobile as rapidly as possible. Since the data rate is high, as long as the user is in the coverage area, this can be done in a few seconds. However, if the information is not available at the infostation and has to be transfered from a server, then the information has to pass through the fixed network before reaching the mobile. Thus, in the "drive-through" scenario, since the coverage area is small, the time which the mobile spends in coverage at a given infostation may not be sufficient to transfer the information from the server to the infostation.

This is a situation peculiar to infostations where the radio rate is assumed much higher than the fixed network rate. Given an inexpensive high-speed radio, there are a number possible reasons for this inversion of the status quo. For economy, one could have a low cost relatively low rate connection (i.e. commodity telephone modems) to each infostation. Alternately, even if one connects the infostations with high speed links, some types of services (i.e. HTTP request) have typical transmission rates of the order of Kbits/second. The server transmission rate and network congestion play an important role in determining the speed of the connection. Another scenario would have fixed network links servicing some primary traffic with the infostation as an add-on service sharing these links. Regardless, in all these cases although the radio rate is high, the user would have to restart a request at the next infostation in the path, and this process will increase the delivery delay, especially if only a small fraction of time is spent in coverage by any one user.

1.3 PROBLEM OVERVIEW, MOTIVATION

The obvious solution to this radio/fixed-net mismatch is to cache or prefetch information at the infostations. As an example, an intelligent prefetching algorithm which attempts to predict *what* the user will need was proposed in [14] as a solution to a location dependent application (map request). The algorithm uses location and speed information to select which of a set of maps should be prefetched. Based on location, time or user dependency, different types of applications would need different schemes for prefetching. However, suppose the information needed is *known* and can be of any sort such as a web page, a map, or personal e-mail. Then, the issue becomes how to partition the information, and then *when* and *where* to send the packets over the fixed network so that they arrive at the user with minimal overall delay.

Thus, consider a system where the infostations are connected as a *cluster* in a hierarchy where there is a higher level with a *cluster controller*, as shown in Figure 1.5. The cluster controller is the entity that has information on all requests that were made in that cluster and how many users are being served at every infostation in that cluster. Note that a cluster would be a natural way of connecting different infostations in the same geographical area, but the cluster controller does not have to be necessarily in the same geographical area.

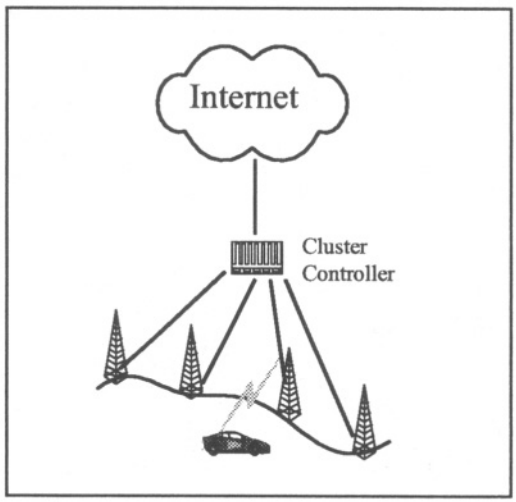

Figure 1.5 The Drive-through Scenario.

The cluster controller can coordinate the delivery of some packets to the next infostation in the mobile path, so that they are locally available at that infostation when the mobile user arrives in its coverage area. If the path is not known then the cluster controller can send the packets to infostations that are most likely to be in the mobile user path. Therefore, during the time the user is going between two infostations, the system can download the information to the next coverage area, reducing the delivery delay, as shown in Figure 1.6.

The optimization problem is then, given some parameters and system configuration, to deliver the information from its current location(s) to the mobile user in a minimum amount of time. The important parameters that have to be taken into account are:

- the overall amount of information that is requested, or file size;

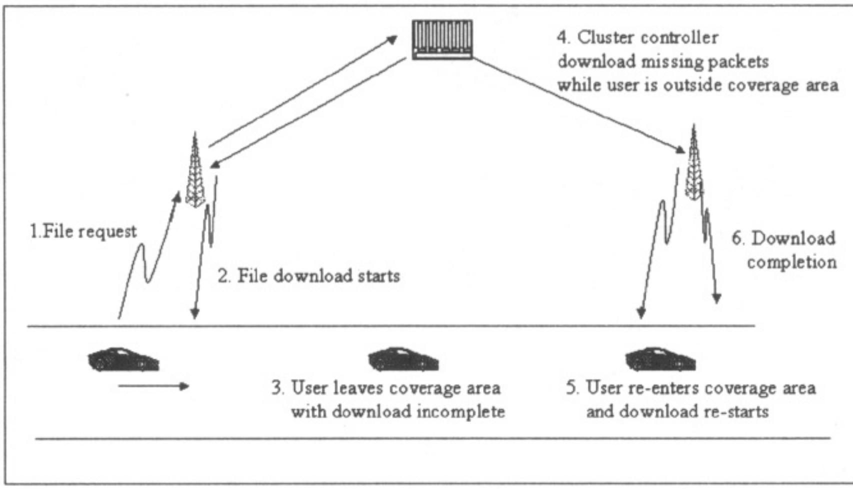

Figure 1.6 The Problem Approach.

- the location of the file, which can be stored at the infostation, at the cluster controller, at the home server (Internet) or distributed over a number of locations;
- the data rate of the wired and wireless network;
- the number of infostations at the cluster;
- the infostations' location;
- the user mobility model.

To better understand the approach used here, consider the single user case where there is a cluster with a given number of infostations, *M*. Let R_i be the rate between the cluster controller and the Internet, R_b the rate of each link between the cluster controller and the infostations and R_r the radio rate. Assume that the user requests a file, which is then divided into packets, and each packet can be sent to different infostations.

Note that if the file is stored in "The Internet" then the network will be able to download the file to the user at the lowest link speed of the network. To take advantage of the fast radio, the cluster controller will prefetch file packets to infostations in the user path. Note that if radio data rate (R_r) were low, then every request should be re-initiated at every infostation. That is, with a slow radio, there is little use in prefetching information to the infostations. Therefore we are interested in the case where $R_r > R_b$ and $R_r > R_i$. In this case the prefetching approach is helpful and only the radio rate will restrict the *maximum* amount of information that should be prefetched at a given infostation since there is a limit on how much can be downloaded to the user in the coverage area.

Given that the radio data rate is large, the specific delivery problem ranges from the trivial to the difficult. Consider the case when information is stored at some server on the Internet. Since we assume $R_r \gg R_i, R_b$, then the fixed network is the limiting factor. In the case when $R_i \leq R_b$ then the cluster controller can broadcast *all* the packets received to *all* the infostations, as shown in Figure 1.7. All the infostations will have the same information that the cluster controller has, as a copy network. As the user passes through the infostations the radio can then download as many packets as possible to the user and discard packets that were already received.

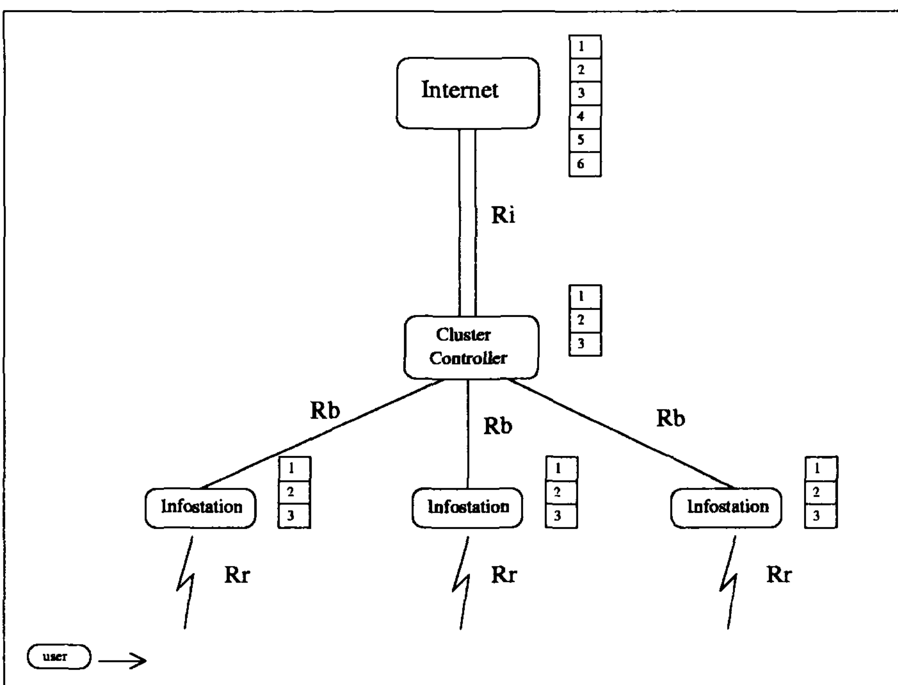

Figure 1.7 An example of a copy network.