

Exploiting Parallel Networks Using Dynamic Channel Scheduling

(Invited Paper)

Lara B. Deek, Kevin C. Almeroth, Mike P. Wittie
Department of Computer Science
UC Santa Barbara
Santa Barbara, CA 93101
{laradeek, almeroth, mwittie}@cs.ucsb.edu

Khaled A. Harras
Department of Computer Science
Carnegie Mellon University
Qatar
kharras@cs.cmu.edu

ABSTRACT

Many researchers have been focusing on the outcomes and consequences of the rapid increase and proliferation of mobile wireless technologies. If it is not already the case, it will soon be rare for a user to be in a situation where absolutely no network connection exists. In fact, through numerous devices, users will soon expect to be connected in all places at all times. Through the great variety and increase in the capabilities of these devices, it is not a surprise to find a single user with many connection opportunities. As a result, we believe that the next major evolution of wireless mobile networks will be in the exploitation of multiple network connections in parallel. Due to network heterogeneity, the major challenge in such situations is to determine the way that these networks can be utilized to better serve different network applications. In this work, we propose a dynamic channel scheduling mechanism that adapts to the state of the available channels to provide more efficient usage of network connectivity. We do so by observing channel throughput, creating a set of channel usage combinations, and then choosing the most efficient combination. We evaluate an implementation of the proposed mechanism using emulation. Our results show that under realistic conditions our dynamic approach greatly improves cost delay metrics, and the overall user-perceived performance compared to a more static approach.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Store and forward networks; C.2.2 [Computer-Communication Networks]: Applications; C.2.5 [Computer-Communication Networks]: Access schemes

General Terms

Performance, Experimentation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WICON '08, November 17-19, 2008, Maui, Hawaii, USA.
Copyright 2008 ICST 978-963-9799-36-3.

1. INTRODUCTION

The increase in communication demands and the popularity of the Internet have prompted the development and expansion of wireless mobile technologies that provide connectivity in increasingly diverse environments and under highly variable conditions. Today's users not only expect reliable network connectivity, but also demand the highest level of performance regardless of their location or underlying network conditions. A promising way to meet user expectations and to provide greater user-satisfaction is to exploit the availability of multiple network connections in parallel. The ParaNets network architecture is an early effort to use multiple heterogeneous networks in unison [1]. The system we propose in this paper builds on ParaNets to transparently unify available network connections into a single Internet service. Specifically in this paper, our goal is to provide the best possible performance subject to the cost of using individual networks and in response to varying network conditions.

Heterogeneous parallel networks, such as cellular, satellite, and Wi-Fi, are characterized by different bandwidth, delay, and monetary cost. These characteristics make each of these networks best-suited for specific types of data transmissions. For example, it may not be frugal to use high-cost satellite bandwidth for a movie download when a free Wi-Fi connection is available. Similarly, a user may not want to search for a free access point to send a short email when also already connected to a cellular data network. When multiple heterogeneous networks such as these are available, the major challenge is to determine the efficient use of each channel with respect to user performance expectations. The system we propose in this paper dynamically determines efficient usage of available parallel networks based on traffic requirements as well as the state of these networks to provide a unified Internet service.

We build our system over the ParaNets-Enabled Data Bundling System for Intermittent Connectivity (DBS-IC) architecture, originally developed to improve communication over challenged networks [1]. ParaNets-Enabled DBS-IC takes advantage of the availability of multiple heterogeneous networks in parallel to create a perception of constant connectivity in challenged networks. ParaNets-Enabled DBS-IC statically distributes data over multiple channels based on the expected performance of the channels' underlying technologies, as well as traffic bandwidth and cost requirements. While the static data allocation schemes achieve

improvements in transmission delay and per bit cost, they suffer when underlying channel performance fluctuates. Our goal in this paper, therefore, is to address the inflexibility of static channel allocation strategies.

The approach presented in this paper is a dynamic channel scheduling mechanism suitable for use as the ParaNets-Enabled DBS-IC allocation strategy. Instead of a static allocation strategy as originally proposed, we present a technique that transmits varying amounts of data over all available channels based on actively measured channel conditions. The amount of data allocated to particular channels is referred to as a *data-to-channel* allocation.

We use the effective throughput of each of the parallel network connections as a determining factor to identify the best allocation strategy. Throughput is the average net bit rate delivered to the application layer and includes the effects of lower-layer factors such as packet loss, transmission delay, and channel congestion. In response to the throughput performance offered on each network, our solution dynamically adjusts the ParaNets-Enabled DBS-IC data-to-channel allocation strategy to achieve the highest possible goodput within the user-specified data transmission cost.

We emulate an implementation of our system on Emulab [2]. We configure different channels within the Emulab framework to act as satellite, cellular, and Wi-Fi networks according to expected throughput, delay, and intermittent connectivity characteristics. We evaluate the effectiveness of dynamic data-to-channel allocation strategies in response to varying channel available bandwidth and availability specific to each emulated network technology. We also compare the performance of our dynamic data-to-channel allocation strategy with that of its static predecessor under varying network conditions and cost requirements. Our results show that the dynamic scheme outperforms the static scheme in all cases. At low costs and extreme channel conditions, the difference between the static and dynamic approach is small; however, as cost increases and under less extreme channel conditions, the dynamic approach exhibits a notable performance improvement.

The remainder of this paper is organized as follows. In Section 2, we discuss the related work. Section 3 presents our system architecture and its operation. In Section 4, we describe our evaluation environment. In Section 5, we present our results. In Section 6, we discuss our future work. The paper is concluded in Section 7.

2. RELATED WORK

We begin this section with work that has addressed challenged network environments and systems that have been designed to deal with such challenged environments. We then refer to research that has dealt with communication over heterogeneous networks, namely satellite, cellular, and Wi-Fi networks. Finally, we discuss work that has addressed allocation strategies to improve performance over communication channels.

A challenged network was first defined by Kevin Fall as a network that experiences mobility, frequent disruption of network availability, high round trip time, and high packet drop rate [3]. Disruption Tolerant Networks (DTNs) are a form of challenged networks. DTNs deal with hostile environments that are characterized by the lack of a continuous connection, and the inability to establish an end-to-end

routing path thereby causing long delays¹. Researchers have studied DTNs with a major focus on routing issues in such extreme environments [4]. While our system focuses on delivering bundled data from a wired stationary device to a mobile node, the same bundling and opportunistic delivery concepts are applicable to “pure” DTNs.

A recent focus of several DTN projects has been to address and study transport layer issues in extreme networking environments as well as focusing on forwarding techniques and routing algorithms. Our paper takes a different approach from DTN’s by taking advantage of any available network and developing a dynamic approach to efficiently and intelligently utilize alternative network paths in the data transmission process. Instead of considering only the target DTN network, we consider all available networks and use a combination of available network connections in parallel.

ParaNets is a recently proposed architecture which utilizes end-to-end networks in parallel with an intermittent wireless connection to best serve a challenged network [1]. By taking advantage of other networks, such as cellular and satellite networks, in addition to an opportunistic Wi-Fi connection, connectivity and communication can be made more reliable and efficient. In ParaNets, the cellular and satellite networks are treated as channels that can be used depending on the size and type of the message. These “alternate” networks are particularly well suited for control information. ParaNets simulation results on top of DTN show that the use of alternative channels of communication when the Wi-Fi connection is broken significantly improves the performance of the challenged network, and hence, user-perceived performance [1].

The ParaNets-Enabled DBS-IC system on which we build our dynamic channel scheduling strategy is an architecture that takes advantage of intermittent connection opportunities as well as infrastructure networks in parallel [5]. ParaNets-Enabled DBS-IC adopts DBS-IC as the challenged network architecture. DBS-IC predicts which data the user will request and prefetches the data before the user explicitly requests it, thereby attempting to create an image of constant connectivity [6]. Additionally, DBS-IC proactively takes advantage of periods of connectivity by grouping user-requested data into a packet called a *bundle* and delivering it to the user incrementally in the form of prioritized *mini-bundles*. ParaNets-Enabled DBS-IC allows DBS-IC communication over multiple heterogeneous networks in parallel. However, it does not address methods to efficiently deliver data over ParaNets—the focus of this paper.

Research on the integration of heterogeneous networks - namely cellular, satellite, and Wi-Fi networks - has focused mostly on inter-system handoffs and on handling inherent TCP/IP unsuitability to challenged environments [7, 8]. CAP and APHOHN are examples of architectures built for such purposes [9, 10]. In these environments, however, the networks are not utilized to best serve the intermittently-connected network. Other research has discussed channel allocation schemes and methods to improve overall network performance; yet, this work has largely concentrated on dealing with single channels [11, 12]. In our work, we devise appropriate allocation strategies between multiple channels at a time; furthermore, we incorporate application layer parameters such as user preferences, size and type of data, or

¹<http://www.dtnrg.com/>

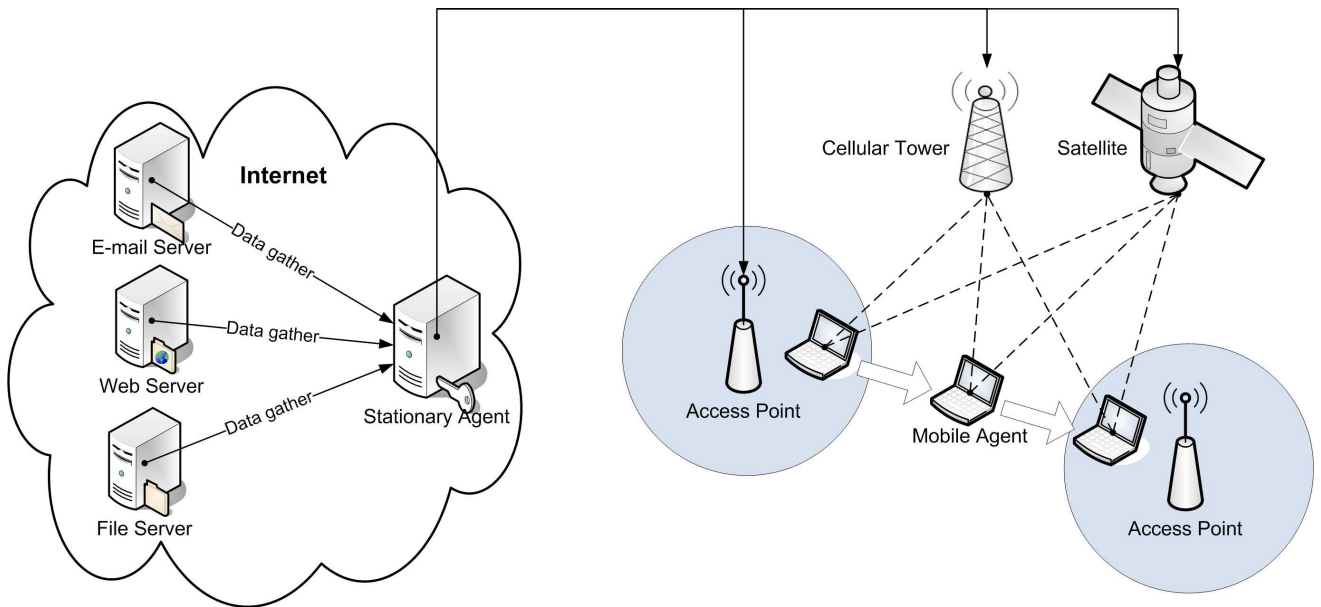


Figure 1: The architecture of a ParaNets-Enabled DBS-IC system.

cost to provide realistic and suitable recommendations for dynamic channel scheduling.

This paper distinguishes itself from previous work in that it examines dynamic methods to efficiently exploit parallel networks. The dynamic portion of our system is based on parameters such as the conditions and state of the underlying network channels as well as user-defined transmission cost constraints. We identify and quantify the benefits of ParaNets and, in particular, how effective and important its use is in challenged network environments.

3. SYSTEM ARCHITECTURE

In this section, we present an overview of our proposed approach. We first describe the system architecture of ParaNets-Enabled DBS-IC and its static data-to-channel allocation mechanisms. We then explain the operation of the dynamic channel scheduling algorithm proposed in this paper.

3.1 ParaNets-Enabled DBS-IC

Our proposed system is built over ParaNets-Enabled DBS-IC depicted in Figure 1. The *stationary agent* maintains Internet connectivity and acts as a caching proxy for a *mobile agent*. The role of the stationary agent is to prefetch and deliver the information requested by the mobile agent. The mobile agent is located on a mobile device in motion between access points. Since access points provide coverage in a limited area, mobile agent mobility results in intermittent Wi-Fi connectivity. In addition to Wi-Fi connections, the mobile agent may also be in range of satellite and cellular networks which provide wide area connectivity. The details and characteristics of the Wi-Fi, cellular, and satellite connections we emulate are detailed in Section 4.

The stationary agent prefetches data for each mobile agent based on (1) a predictive method whereby the stationary agent keeps a history of the information the user has requested to view in past sessions, and (2) a method whereby the mobile agent explicitly informs the stationary agent of

the web, e-mail, or file data it would like to access. The stationary agent then contacts the corresponding application servers and gathers the requested information.

The stationary agent bundles the gathered data into a set of mini-bundles to be divided for transmission onto the available channels according to our proposed data-to-channel allocation strategy. An alternative to using mini-bundles is the aggregation of data into a single bundle. However, mini-bundles allow for prioritized transmission of content grouped into high priority mini-bundles and for straightforward re-assembly of data received over different channels. Mini-bundle also simplify the dynamic evaluation of candidate data-to-channel allocation strategies as described in Section 3.2.

The static data-to-channel allocation strategy is determined as follows. The static strategy calculates the throughput and cost values of all the possible data-to-channel combinations based on ideal network performance, where performance reflects network throughput under ideal channel conditions. Users specify the maximum cost in cents/MB that they are willing to incur to receive data: the *user-defined cost*. The static mechanism will then calculate all possible data-to-channel mappings using the user-defined cost as a threshold. The range of points that the static strategy will scan are those that incur the user specified maximum cost or a lower cost in close proximity to that maximum cost. Within this range of points, the static scheme will select the transmission strategy that achieves the highest throughput. By repeating this process for different user-defined costs, the static strategy will create a set of allocation combinations that it will use to build a data-to-channel mapping.

Figure 2 shows one possible channel-share-to-cost mapping based on individual channels' usage price and expected bandwidth where the x-axis specifies the user-defined cost. We note that the behavior of the static mapping strategy shown in Figure 2 varies. For example, the use of the cellular channel fluctuates whereas that of the satellite behaves like a step function. This behavior reflects the selective ap-




Figure 2: Static data-to-channel allocation mapping.

proach that the static method follows to determine the data-to-channel allocation mapping. Methods for channel usage are based on finding and selecting the data-to-channel allocation combination that achieves the highest throughput within a particular cost range that is specified by the user-defined cost. In other words, there are no fundamental network conditions for channel usage. The static mechanism assumes that the respective channels will behave ideally in all situations, which is far from realistic.

Based on a simple survey of currently available service plans, we use a cellular data plan of 0.50 cents/MB at 14.4 Mbps², and a satellite plan of 0.65 cents/MB at 10 Mbps³. This cost model reflects the higher per bit cost of cellular and satellite delivery as well as their greater expected availability. Users willing to pay more for use of these networks generally experience lower download times as a larger share of their data is sent over more frequently available cellular and satellite networks. On the other hand, users willing to pay less, will take greater advantage of the opportunistic Wi-Fi connections, but may need to wait for such connections to become available. In such situations, any time critical mini-bundles are transmitted using cellular or satellite channels. Based on the user-defined cost for data transmissions, the stationary agent looks up the share of mini-bundles to be sent on each channel in the static mapping and transmits data accordingly.

While the static mapping between the cost and channel share has been shown to have shorter download times at lower cost when a single connection is used, the solution suffers when the performance of the underlying network changes. For example, the throughput of the Wi-Fi connection can vary with respect to the load or availability of the access point. Similarly, cellular connection throughput can depend on signal strength [13]. To deal with changes in network conditions, a dynamic mechanism to determine channel data share is needed. Our extension to ParaNets-Enabled DBS-IC monitors available bandwidth on each channel and adjusts the channel-share-to-cost mapping dynamically and then uses that mapping to choose the best data-to-channel allocation strategy.

²<http://www.t-mobile.co.uk/>

³<http://www.wildblue.com/>




Figure 3: Combined channel goodput of data-to-channel allocation mappings.

3.2 Dynamic Data-to-Channel Allocation

The earlier version of ParaNets-Enabled DBS-IC adopted a static approach of utilizing the parallel networks. We now propose a mechanism to adjust the data-to-channel allocation strategy to changing network conditions using a dynamic channel scheduling algorithm.

The static data-to-channel allocation strategy of ParaNets-Enabled DBS-IC works well when network bandwidth adheres to advertised performance. However, in real network deployments the offered bandwidth will vary with changes in channel load, signal strength, and intermittent connectivity. Subsequently, the performance of any static allocation strategy will degrade when available bandwidth diverges from the amount of bandwidth assumed when the static calculation was performed. To assure efficient operation of ParaNets-Enabled DBS-IC, the data-to-channel allocation strategy needs to dynamically adjust in response to changes in network conditions.

Given that network access is assumed to have non-zero cost, the goal of our dynamic data-to-channel allocation strategy, therefore, is to find the allocation with the highest goodput within a given cost limit. To select the appropriate data-to-channel allocation mapping, we evaluate the goodput performance of candidate allocation mappings with respect to their cost. The set of candidate allocation mappings is the set of all possible ways of dividing the queued mini-bundles across the available channels. Thus, one particular allocation mapping for ten mini-bundles might be to send five mini-bundles on the Wi-Fi connection, three on the cellular connection, and two on the satellite connection. To evaluate the relative performance of candidate allocation mappings, we maintain a running average of channel goodput based on past transmissions. Channel throughput reflects the state and condition of the underlying network such as network load, packet delay, packet loss, or temporary loss of connectivity. By considering channel throughput, we are considering network factors relevant to bulk data transfer performance.

To choose a data-to-channel allocation mapping, we use the measured channel available bandwidth and calculate the expected combined goodput of each candidate mapping with respect to cost. The results of a sample calculation are presented in Figure 3. Data in Figure 3 exhibits a bell-shaped

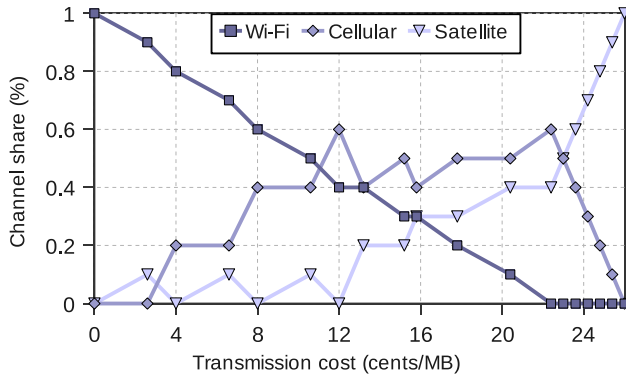


Figure 4: Dynamic data-to-channel allocation mapping.

plot of goodput versus transmission cost. At zero cost, only the the Wi-Fi channel is used. As the cost limit increases, the mapping gradually shifts transmission of data to the more expensive cellular and satellite networks. This shift in data-to-channel mappings increases network throughput as networks are exploited in parallel in order to speed the transmission of data. This trend, however, is inverted beyond a point where maximum goodput is achieved and whereby additional use of the alternative low-bandwidth channels will not increase throughput, but only increase the cost.

The data-to-channel allocation mapping is chosen as follows. Based on the distribution of goodput combined with cost, the dynamic mechanism chooses the best data-to-channel allocation mapping as the one with with the highest goodput. The dynamic mechanism recurses by finding the next best point located to the left of the previous value. The set of these points are then connected together. The resulting line, shown in Figure 3, forms the upper left boundary of the bell-shaped curve. It follows that all the data-to-channel allocation mappings below and to the right of this set of points represent a lower goodput-to-cost ratio and should not be chosen. Out of the set of selected points, the dynamic channel scheduling algorithm chooses the best points below the per megabyte cost a user is willing to incur.

Figure 4 shows the dynamic data-to-channel mappings for the set of selected points from Figure 3. For each of the selected points in Figure 3, Figure 4 shows the fraction of mini-bundles that are sent using each of the three available network connections. For example, for the lowest transmission cost, all bundles are sent using the Wi-Fi connection; for the transmission cost of “12”, 60% of the mini-bundles are sent using the cellular network, 40% are sent using the Wi-Fi connection, and none are sent using the satellite network. For completeness, we calculate and include in Figure 4 the set of points to the right of the highest goodput point.

The results of Figure 4 complement the findings of Figure 3: to increase combined throughput, parallel networks must be exploited for data delivery. A consequence of the increased throughput, however, is increased cost. As cost increases (moving along the x-axis), the data-to-channel mapping shifts from complete use of the Wi-Fi channel to more use of the cellular and satellite networks.

The set of points in the data-to-channel mappings presented in Figure 4 varies with varying network conditions.

Changes in network conditions affect the location of points on the graph. The result, therefore, is a change in the resulting data-to-channel mapping and a different set of selected points. For example, if the Wi-Fi channel conditions deteriorate, the goodput will decrease for combinations with Wi-Fi transmission, and the dynamic data-to-channel mapping will determine the efficient points that will express a decrease in Wi-Fi channel usage at each cost value.

The mapping in Figure 4 represents the allocation mappings selected even when channel bandwidth changes dynamically. The dynamic scheduling strategy used to produce this mapping differs compared to the static mapping in Figure 2 where the allocation strategies are based only on the advertised channel bandwidth values. The *static strategy* builds its data-to-channel allocation mapping based on ideal conditions for network performance. The static mechanism selects transmission combinations that represent the highest throughput within the range between the user defined cost and the costs in its lower proximity. The *dynamic strategy*, on the other hand, selects the efficient points based on continuous network measurement, as shown in Figure 3. Further, when the cost metric exceeds the optimum point at a cost value of 16, the static approach would continue to operate in the high-cost region whereas the dynamic approach will remain below this threshold. The dynamic approach will not operate beyond this threshold point because it would incur additional cost but with decreased throughput. This result is shown in Figure 3 by the corresponding drop in the second half of the bell-shaped curve. These data points indicate that at costs higher than the threshold point, the dynamic strategy always performs better.

We observe that the point distribution of the dynamic strategy in Figure 4 is more well-behaved in comparison to that of the static strategy in Figure 2. At low costs, both schemes behave similarly. However, as the allowable cost limits are increased, the dynamic approach follows a more continuous pattern that is representative of transitions between channel usage that are based on realistic network conditions. This continuity implies that a change in the use of a particular network will be based on knowledge of the actual network conditions. For example, if the satellite connection is weak, channel usage over the satellite network will be incremented only after a particular increase in the allowable cost and not earlier. The static mapping, however, is not aware of actual network state and its mappings are based on favorable, unloaded conditions; therefore, it does not always exhibit effective balancing across all of the underlying channels. Ultimately, with an increase in network channel conditions, the differences between both approaches becomes more apparent as the dynamic approach is more adaptive and offers efficient data-to-channel mappings. A detailed comparison between the static strategy and the dynamic strategy is presented in Section 5.

The advantage of our dynamic strategy is the data-to-channel allocation mappings that adjust to changing network conditions. Additionally, unlike in the static case, our dynamic strategy always chooses values left of the threshold point in the bell-shaped curve, whereas the static strategy might choose a high-priced solution, while a lower priced, higher combined goodput allocation strategy is likely available. When a user’s service cost is fixed, the best strategy is the point with the highest aggregate goodput. However, when a user’s service cost is usage-based, the user will likely

be interested in choosing a lower performance and price point, thereby reducing use of more expensive channels by still achieving reasonable throughput and response time.

4. EVALUATION ENVIRONMENT

In order to evaluate the performance of our system, we have created an emulation environment that enable us to evaluate the performance of the ParaNets-Enabled DBS-IC static and dynamic data-to-channel allocation strategies. The main configuration parameters for our evaluation are presented in Table 1. We choose the parameters that best represent the behavior of the underlying cellular (3G), satellite (LEO), and Wi-Fi (802.11g) communication channels. We evaluate our system under a range of channel conditions and present results from one scenario generally representative of the trends we have observed.

In real-world deployments, cellular and Wi-Fi channels may be restricted by high traffic loads due to high demand for connectivity, whereas the satellite channels typically remain well-provisioned. We therefore vary the load over the cellular and Wi-Fi channels but maintain unloaded conditions for the satellite network. We choose a load of 50% to represent low load network conditions, 70% for moderate congestion, and 90% for extreme congestion. We excluded values that were outside of this range as values of less than 50% do not perform any differently than a system with 50% load. Furthermore, values greater than 90% introduced additional performance artifacts not related to the focus of our study.

To perform our evaluation, we developed an implementation of our system and tested it over an emulated network. We chose to implement our system, as opposed to simulating it, and tested it over emulated network conditions to achieve higher fidelity of physical and transport layer performance. Our network emulations were conducted over Emulab [2], and we used the results obtained in previous studies to accurately and realistically configure the emulation environment [7, 14, 15, 16].

Our network testbed was configured by an Emulab NS-2⁴ file. The stationary agent ran on a machine with a stable connection to the Internet. It was configured to have five 10/100 Mbps full-duplex switched Ethernet connections. The mobile agent as located on an intermittently connected machine with five Ethernet network cards. The mobile and stationary agents were connected through three of these interfaces; each interface mimics the characteristics of either the satellite, cellular, or Wi-Fi connection. We use the commands provided by Emulab to adjust the characteristics of the three network connections as listed in Table 1.

To extend the period of measurement and accurately characterize network throughput we chose as the basis for our experiments the task of transferring 40 MB. While 40 MB is a rather large amount of data, it was useful to demonstrate the likely behavior of the system averaged over a longer period of time. Based on our results, we believe our scheme to perform equally well when used with smaller amounts of prefetched data and mini-bundle sizes.

In our evaluation, we assumed continuous connectivity for the cellular and satellite networks; therefore, we did not model interruptions caused by horizontal handoffs within either of these two communication channels. However, hor-

izontal handoffs in the Wi-Fi network were taken into consideration. We modeled Wi-Fi disconnectivity following the suburban driving model discussed by Gass et al. [16].

In our simulations we configured data rate and delay characteristics representative of respective network channel technologies. For Wi-Fi connections we used IEEE 802.11g standard. 802.11 is widely deployed in home and municipal networks and its connection handoff characteristics are well understood [16, 17]. For the cellular connection, we considered third-generation (3G) technology recently becoming more widely deployed [7, 15]. For the satellite connection, we emulated the characteristics of Low Earth Orbit Satellites (LEOs) commonly used for commercial Internet service [14].

The metrics we use to evaluate our system are goodput and goodput-to-cost ratio.

- Combined Goodput is the rate at which data is received at the mobile agent from the combined transmissions of the underlying communication channels. It measures the average bit rate delivered to the mobile agent's application layer exclusive of all protocol overhead and data packet retransmissions. It provides a higher-level measurement of the efficiency of our dynamic strategy. Since we deal with multimedia file transfers, combined goodput corresponds to the achieved file transfer rate over the three available parallel networks. We use this metric to measure the user-perceived performance to a particular data-to-channel allocation strategy.
- Goodput-to-cost measures system performance with respect to per transmission unit cost. It maps goodput to user cost. The higher the goodput, the greater the transmission efficiency at a particular cost. We use this ratio to evaluate the transmission choices available to the ISP to transmit data with the lowest cost.

The goal of our work is to evaluate our system with respect to goodput and goodput-to-cost ratio. In order to do so, we chose a configuration environment that enables us to study this behavior, as detailed in this section. In the next section, we present the set of results that demonstrate the performance of the ParaNets-Enabled DBS-IC static and dynamic data-to-channel allocation strategies. We focus on the changes in the performance of the evaluated strategies in response to changing network conditions. We chose values for network parameters that use realistic assumptions for the underlying network conditions as explained earlier.

5. RESULTS

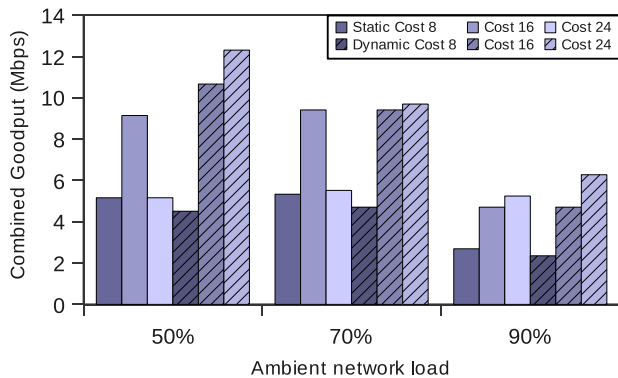
In this section, we present results sufficient to demonstrate the key findings of our work. The following results were obtained by varying network load between 50%, 70%, and 90% for the cellular and Wi-Fi channels while keeping the load at 0% on the satellite channel. We focus on the changes in the performance of the proposed schemes in response to changing network conditions. We choose these values as they reflect a realistic framework for the underlying network conditions as explained previously.

Figure 5 compares the performance of the static and the dynamic strategies with respect to the user-specified maximum transmission cost. Performance is measured with respect to combined goodput, which specifies the rate at which

⁴<http://www.isi.edu/nsnam/ns/>

Table 1: Emulation configuration.

| | Parameter | Value Range | Nominal Value |
|-------------|---------------------------------|--|-----------------|
| Application | Prefetched data | 10 MB - 20 MB - 40 MB - 100 MB | 40 MB |
| | Number of mini-bundles | 10 | 10 |
| Wi-Fi | Data rate | 54 Mbps | 54 Mbps |
| | Round trip time | (60 ± 10) ms | (60 ± 10) ms |
| | Ambient load | 0% - 50% - 90% | 50% |
| | Intermittent connectivity Model | 15s on, 30s off; 20s on, 10s off; 40s on, 15s off; 120s on, 20s off | 20s on, 10s off |
| Cellular | Data rate | 14.4 Mbps | 14.4 Mbps |
| | Round trip time | (100 ± 25) ms | (100 ± 25) ms |
| | Ambient load | 0% - 50% - 90% | 50% |
| | Cost | 0.50 ¢/MB | 0.50 ¢/MB |
| Satellite | Data rate | 10 Mbps | 10 Mbps |
| | Round trip time | (200 ± 50) ms | (200 ± 50) ms |
| | Ambient load | 0% | 0% |
| | Cost | 0.65 ¢/MB | 0.65 ¢/MB |

**Figure 5: Comparative performance of dynamic and static data-to-channel allocation mechanisms.**

the user is receiving data over all the parallel channels simultaneously. For each load value, we measure the average combined goodput according to the data-to-channel allocation strategy of both the static and dynamic scheduling strategies. The higher the combined goodput, the faster the user is receiving data for a given cost value and the better the achieved performance. We observe that the dynamic strategy outperforms the static strategy regardless of cost constraints. While the static strategy is sending data based on a fixed transmission schedule, the dynamic strategy determines the most efficient data-to-channel allocation mapping based on channel conditions and cost requirements. The dynamic strategy makes more accurate decisions for channel usage and outperforms the static strategy by adapting its transmission pattern to the varying network conditions.

We also observe in Figure 5 that, as the cost limit increases, the difference between the static and dynamic strategies becomes more pronounced. At low costs, both schemes achieve almost the same level of performance. This trend is due to the limited number of possible channel combinations available when a small cost value is imposed by the user.

Increases in performance at higher allowable cost values can be attributed to the fact that the use of the highly available cellular and satellite networks for data transmission incur greater costs, but mitigate the negative effects of Wi-Fi intermittent connectivity. Alternatively, the Wi-Fi communication network is free, but periodically unavailable. At low costs, the dynamic strategy shifts the majority of mini-bundles to be transmitted using the Wi-Fi channel, even if the wait time is long (because the channel is currently experiencing an outage). This choice is similar to the static allocation strategy, which for low cost values, essentially makes the same choice—to use the free Wi-Fi channel regardless of the time it takes to deliver all of the mini-bundles.

As we increase the load on the Wi-Fi and cellular channels from 50% to 70% and finally to 90%, the difference in the average goodput between both methods becomes less prominent. As a result of increasing the load on the cellular and Wi-Fi channels, the range of possible data-to-channel allocation strategies is narrowed to using the well-provisioned satellite network for the majority of data transmissions. As the conditions of the parallel networks are made to represent a congested network, the majority of mini-bundles will be shifted to a single channel for effective data transmission. In such a case, the difference in performance between the static and the dynamic methods becomes less pronounced.

Figure 6 shows the goodput-to-cost ratio for lightly-loaded channel conditions graphed on a log scale. We plot the goodput-to-cost ratio with respect to the cost that the user is willing to incur. This graph allows us to evaluate system performance with respect to per-transmission-unit cost. We observe that the goodput-to-cost ratio decreases rapidly with increases in transmission cost. This trend can be correlated with the bell-shaped behavior of the aggregate goodput-to-cost measurements as depicted in Figure 3. At low costs, data is predominantly sent over the cheap Wi-Fi channel. Although goodput is low when low cost is a requirement, the low cost values will lead to an apparent tradeoff between the goodput achieved and the benefits received with respect to cost. When the user is willing to pay more, the costly alternative cellular and satellite connections will be

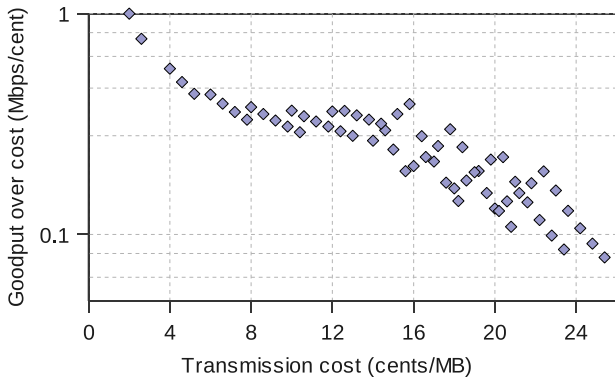


Figure 6: Data-to-channel allocation efficiency with respect to transmission cost.

more heavily utilized to reduce delay; the goodput-to-cost tradeoff, however, becomes less favorable.

The steep decrease in transmission cost with respect to system goodput suggests the possibility of cost savings by Internet Service Providers (ISPs). Evaluating such behavior is important in that it allows us to identify the benefits our system can provide to ISPs as well as to individual users. Through a comparison of user goodput versus the goodput-to-cost ratio, the ISP can balance user satisfaction and savings by sending data to the user using a channel allocation strategy with slightly lower performance but much lower cost. We plan to focus on this tradeoff in future work.

6. FUTURE WORK

The dynamic scheduling strategy we propose in this paper adopts a passive approach for measuring channel throughput. By sending data over a particular channel, our system is also calculating the average throughput of delivering that data and deciding on the next data-to-channel allocation mapping accordingly. Such measurements can become stale after a period of not using a particular channel. We plan to investigate the adoption of a more active approach that probes the network if a particular channel goes without being used for a long period of time.

In this work, we present one dynamic approach for channel allocation strategies. An alternative dynamic algorithm we propose would dedicate queues to each of the available network channels. The data to be sent would wait in each of these queues; when a channel is connected, we allocate a mini-bundle of data to that channel's queue. In general, such an approach would allow even more dynamic decisions to be made. The tradeoff is that bad decisions could be made in cases where a mini-bundle is allocated to a low bandwidth, slow channel. In this scenario, completion of the entire bundle transfer would be dependent on the last mini-bundle now being sent over the slow channel. The decision would have been a bad one if a high speed channel could have completed the job much more quickly, but was not available when the original allocation to the slow queue was made. A more intelligent approach can be achieved, however, if we have some knowledge of the conditions of the network for each queue. A study of more elaborate dynamic/hybrid approaches that

can combine two or more sources of information is left for future work.

Related work on TCP performance for satellite networks, particularly LEO links, and cellular networks, particularly 3G links, have identified problems and a wide variety of solutions. A survey performed on the usage of different transport protocols in IP-based networks with satellites has identified the reasons why TCP has difficulty with satellite links [18]; other work evaluates how well end-to-end transport connections perform in a satellite environment and summarize how latency and asymmetry can impair the performance of TCP [19, 20]. Similar studies have been performed on 3G connections and have arrived to similar conclusions concerning the problems TCP faces with 3G cellular links [21]. As part of our future work, we intend to study the impact transport layer protocols on the performance of our system, and possibly identify alternative transport layer protocols that would work seamlessly within our system.

7. CONCLUSIONS

In this paper, we have presented a dynamic approach to efficiently exploit the availability of heterogeneous networks in parallel. Our goal has been to improve user-perceived performance by providing an efficient and accurate recommendation for channel usage. Efficient usage of available channels involves reducing transmission time to the minimum possible for a given user-defined cost constraint. Our proposed dynamic strategy accomplishes this goal by adapting the data-to-channel allocation mappings to the underlying network conditions and, in so doing, providing more accurate recommendations.

We build our system over a previously proposed architecture, ParaNets-Enabled DBS-IC, that takes advantage of the availability of multiple networks in parallel in order to create a perception of constant connectivity. We extend the static allocation strategy of ParaNets-enabled DBS-IC and propose adapting our channel usage scheme to the underlying network state and conditions. We compare our adaptive strategy against this static approach under different scenarios. Our results showed that the improvement in performance varied with cost and channel conditions; the lower the cost or the more extreme the channel conditions, the less improvement the dynamic scheme shows. However, in all cases, the dynamic approach outperforms the static strategy and shows significant improvement.

We present our dynamic strategy with an emphasis on the improvement it would achieve for user-perceived performance. Through the analysis of our results, we verify that our system can also be studied with a concern for ISP-perceived performance. Through a comparison of user goodput versus the goodput-to-cost ratio, an ISP can balance user satisfaction and savings by sending data to the user using a channel allocation strategy with slightly lower performance, but much lower cost.

A dynamic approach for utilizing parallel networks is a solid step towards improving the performance of data transmissions in environments with heterogeneous connection opportunities. However, there are several areas open to future work.

8. REFERENCES

- [1] K. Harras, M. Wittie, K. Almeroth, and E. Belding. ParaneTs: A parallel network architecture for challenged networks. In *HOTMOBILE '07: Proceedings of the Eighth IEEE Workshop on Mobile Computing Systems and Applications*, Nov. 2007.
- [2] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An integrated experimental environment for distributed systems and networks. *SIGOPS Operating System Review*, 36(SI):255–270, 2002.
- [3] K. Fall. A delay-tolerant network architecture for challenged internets. In *SIGCOMM '03: Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, Feb. 2003.
- [4] A. Lindgren, A. Doria, and O. Schelén. Probabilistic routing in intermittently connected networks. *SIGMOBILE Mobile Computer Communications Review*, 7(3):19–20, 2003.
- [5] L. B. Deek, S. Thoubian, S. Jamijian, K. Harras, and H. Artaïl. Exploiting parallel networks in intermittently-connected mobile environments. In *WiMob '08: Proceedings of Fourth International Conference on Wireless and Mobile Computing, Networking and Communications*, Oct. 2008.
- [6] C. Holman, K. A. Harras, K. C. Almeroth, and A. Lam. A proactive data bundling system for intermittent mobile connections. In *SECON '06: Proceedings of 3rd Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks*, Sep. 2006.
- [7] S. E. Kim, J. A. Copeland, C. Chen, L. Wang, L. Liu, and W. Wu. Interworking between WLANs and 3G networks: TCP challenges. In *WCNC '04: Proceedings of 2004 IEEE Wireless Communications and Networking Conference*, Mar. 2004.
- [8] M. Buddhikot, G. Chandranmenon, S. Han, Y. W. Lee, S. Miller, and L. Salgarelli. Integration of 802.11 and third-generation wireless data networks. In *INFOCOM '03: Proceedings of Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies*, Mar. 2003.
- [9] K. W. Lee, Y. B. Ko, and T. Nandagopal. Load mitigation in cellular data networks by peer data sharing over wlan channels. *Computer Networks*, 47(2):271–286, 2005.
- [10] D. Kliazovich, F. Granelli, G. Pau, and M. Gerla. APOHN: subnetwork layering to improve TCP performance over heterogeneous paths. In *NGI '06: Proceedings of 2nd Conference on Next Generation Internet Design and Engineering*, Apr. 2006.
- [11] J. Yang. An efficient fault-tolerant distributed channel allocation algorithm for cellular networks. *IEEE Transactions on Mobile Computing*, 4(6):578–587, 2005.
- [12] J. H. Lin, C. R. Dow, Y. H. Li, C. M. Lin, and T. C. Huang. An efficient channel allocation scheme in cell overlapping systems. In *CCNC '05: Proceedings of Second IEEE Conference on Consumer Communications and Networking*, Jan. 2005.
- [13] M. P. Wittie, B. Stone-Gross, K. C. Almeroth, and E. M. Belding. MIST: Cellular data network measurement for mobile applications. In *BROADNETS '07: In Proceedings of Fourth International Conference on Broadband Communications, Networks and Systems*, Sep. 2007.
- [14] S. Subramanian, S. Sivakumar, W. J. Phillips, and W. Robertson. Investigating TCP performance issues in satellite networks. In *CNSR '05: In Proceedings of the 3rd Annual Conference on Communication Networks and Services Research*, May 2005.
- [15] M. Frodigh, S. Parkvall, C. Roobol, P. Johansson, and P. Larsson. Future-generation wireless networks. *IEEE Personal Communications*, 8(5):10–17, Oct. 2001.
- [16] J. Scott R. Gass and C. Diot. Measurements of in-motion 802.11 networking. In *WMCSA '06: In Proceedings of Seventh IEEE Workshop on Mobile Computing Systems and Applications*, Apr. 2006.
- [17] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan. Interactive WiFi connectivity for moving vehicles. *Sigcomm Computer Communications Review*, 38(4):427–438, 2008.
- [18] A. Houyou, R. Holzer, H. Meer, and M. Heindl. Performance of transport layer protocols in LEO pico-satellite constellations. Technical report, University of Passau, 2005.
- [19] C. Partridge and T. J. Shepard. TCP/IP performance over satellite links. *IEEE Network*, 11(5):44–49, 1997.
- [20] T. R. Henderson and Y. H. Katz. Transport protocols for internet-compatible satellite networks. *IEEE Journal on Selected Areas in Communications*, 17:326–344, 1999.
- [21] I. Koukoutsidis. TCP over 3G links problems and solutions. Technical report, Fourth - ICS, 2005.