iShare: Exploiting Opportunistic Ad hoc Connections for Improving Data Download of Cellular Users

Long Vu, Klara Nahrstedt
Department of Computer Science, University of Illinois
Email:{longvu2,klara}@illinois.edu

Ivica Rimac, Volker Hilt, Markus Hofmann
Bell Labs/Alcatel-Lucent, NJ07733, USA
Email: {rimac,volkerh,hofmann}@bell-labs.com

***Abstract*—This paper[1] presents an efficient sharing protocol (iShare) that blends different wireless interfaces of the mobile device for content dissemination service. With iShare, mobile users download content from a source via the cellular link and at the same time form an ad hoc mesh network for peer-to-peer exchange of content data. The mesh remains robust to network dynamics, minimizes ad hoc communication overhead, and parallelizes the downloading process among mesh members. In order to counter selfish behavior and balance the download among mesh members, we apply a practical "tit-for-tat" incentive mechanism, which exploits proximity and mutual content interest of mobile users. We simulate, evaluate the performance of iShare and compare it to other content dissemination schemes using cellular broadcast channels, cellular unicast channels, and tree-based protocols. The obtained results show that iShare significantly outperforms alternative approaches. The results also confirm that iShare enables users to continuously obtain data via ad hoc connection during the cellular handoff period, and provides multi-homing download for groups spanning adjacent cellular cells.**

I. INTRODUCTION

The proliferation of mobile devices and swift advance of wireless technology lead to the emergence of pervasive computing paradigm. Modern mobile devices today, besides their traditional low-bandwidth, long-range communication such as cellular or satellite, already come equipped with additional wireless interfaces such as IEEE 802.11 (wifi) and Bluetooth. These interfaces can be leveraged to use unlicensed spectrum for high-bandwidth, short-range (i.e., ad hoc) communication without any additional investment in the network infrastructure. As a result, the short-range communication becomes promising to exchange data among mobile devices and reduce the load on the long-range links.

Some recent research projects have focused on combining cellular and ad hoc connectivity on cellular device [7], [6], [10], [8]. Most of the published approaches are based on either the construction of a tree of ad hoc nodes rooted at high-data-rate proxy nodes connected to the cellular network or the full collaboration among peers. The construction and maintenance of a tree structure incurs a high communication overhead under network dynamics while the assumption of full collaboration among peers are not always realistic.

Besides the aforementioned technical drawbacks of existing approaches, they also lack realistic and applicable incentive mechanism for user collaboration, e.g., by assuming general unselfishness and total cooperativeness among nodes [4], [9], [5]. In reality, however, users might not necessarily turn on their ad hoc interfaces to forward data to other users for free. In contrast, people tend to collaborate and are more willing to use ad hoc interface to exchange data if they share mutual content interest.

In this paper, we first present scenarios where mobile users share mutual content interest, which motivate them collaborate in exchanging content over the P2P channel. Then, we present an efficient and practical sharing protocol (iShare) that leverages cellular and P2P interfaces on mobile devices to provide improved content dissemination service. The iShare protocol is based on the mesh formation using ad hoc connectivity among co-located mobile users, which makes the solution very robust to network dynamics and minimizes ad hoc communication overhead. Our design also includes a lightweight and practical incentive mechanism based on "tit-for-tat" reciprocation that helps to enforce cooperation among members, counter selfish behavior, and balance the load among ad hoc mesh members. By utilizing promiscuous and broadcast modes of the ad hoc channel, our incentive mechanism is particularly suited and effective in mobile networks.

In the following sections, we first present motivating scenarios where the mobile users download similar content and the system model in Section II. We then present the design of the iShare protocol in Section III. Next, Section IV evaluates iShare protocol by simulation in NS2. Finally, we conclude the paper in Section V.

II. SYSTEM MODEL

A. Motivating Scenarios

There are numerous scenarios where co-located mobile users download similar content.

The first class of scenarios exists when co-located people stay at the same location and download the location-dependent content. For example, the fans attend a football

---

[1] This work is supported by Boeing Trusted Software Center at the Information Trust Institute, University of Illinois.

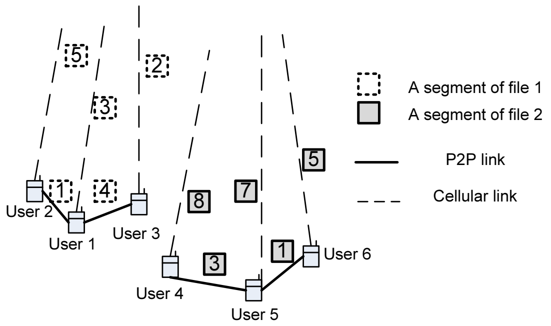

Figure 1. Network Model

match in a stadium and use their cell phones to download the same video clip of a replay scene from the content server. The fans, while downloading the clip via cellular link, can exchange downloaded data via the ad hoc interface of their cell phones to speed up the download.

The second class of scenarios can be found when co-located people request a global content. For example, when waiting for their buses in the morning at a bus station, cellular users may want to download the breaking news video clip about an emergency event occurring the night before in Chicago downtown from the content server. Since it may take from 5 to 20 minutes to download the clip from the server via the cellular link, the ad hoc interface of these co-located mobile devices can be used to exchange downloaded chunks of data in order to accelerate their downloads.

The third class of scenarios exists when a group of people is moving toward the same place. For example, when the audiences walk toward the outdoor concert area, they may download short video clips of artist profiles or concert program from a content server to their cell phones. Again, the ad hoc connection becomes useful to exchange downloaded data and accelerate the download process.

In the above scenarios, co-located users may not download the content at *exactly* the same time. However, since the content in these scenarios is spatial-temporal (i.e., location and time dependent), the downloading periods of co-located users are highly *overlapping*. Our iShare protocol exploits the co-location, mutual content interest of mobile users, and the overlap of downloading periods to form an ad hoc mesh of mobile users to exchange their downloaded data.

B. Data Model and Network Model

In our context, mobile users download a file from the content server. We assume that the file is similar to the Bittorrent file, which has a unique *file id* and consists of multiple equal-sized segments. Each segment also has a unique *segment id* to distinguish it from other segments of the same file.

Figure 1 shows our network model where co-located mobile users (or mobile nodes) download the same file. We assume that each mobile node has two communication links: cellular link and P2P link (e.g., wifi). In Figure 1, users 1,2,3 download file 1 from server 1 via their cellular links (likewise, users 4,5,6 download file 2 from server 2). Co-located nodes download the same file and form a mesh structure, in which each mesh member may download different segments of the file in parallel and exchange downloaded segments via P2P links. For example, node n_1 may download segment 3, node n_2 may download segment 5, and they exchange segment 1. In this paper, the terms P2P link, ad hoc link, and wifi connection are used interchangeably.

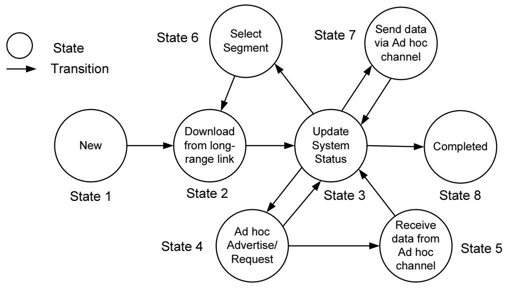

Figure 2. An iShare node's protocol state machine

III. THE ISHARE PROTOCOL

A. Overview of iShare

Figure 2 shows the protocol state machine of an iShare node n. When the user starts his file download, his wireless device n stays in state 1 and n downloads the list of segment ids of the desired file from the cellular link. Then, n moves to state 2 and downloads one random data segment from the cellular link again. After that, n stays in states 2,3,6 and continuously downloads its missing segments via the cellular link, one random missing segment at a time. At the same time, n attends an ad hoc mesh of iShare nodes, advertises its available segments, requests missing segments, and receives segments from its neighbors via the ad hoc link (states 3,4,5). For a fair collaboration with other nodes, n applies the "tit-for-tat" incentive mechanism to send segments to neighbors via the unicast ad hoc link (state 7). Whenever a missing segment is obtained by either link, n switches to state 3 and updates the system status. When the desired file is obtained entirely, n switches to the completed state.

B. Bootstrapping iShare

This section focuses on states 1,2,3 in Figure 2. When a mobile user starts requesting a file, his device (i.e., the iShare node n) is in the New state. n first obtains the metadata of the file such as *file id* and the list of segment ids from the content server and downloads a random segment s of the file via the cellular link. Receiving s, n stays in state 3 where n puts s into its memory and updates its currently missing/available segments. The next question is whether n turns its ad hoc interface on to find iShare neighbors since

0	1	0	0	1	1	0	1	0	0	0	1	1	0	1	1
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

Figure 3. The HELLO message format

n needs to save energy. There are two solutions. First, we can use a server from the infrastructure to track the location of nodes using GPS or Wifi access points [2] and return the closest iShare neighbor n_1 to n if n_1 is downloading the same content as n (by asking the file server). If n_1 is inside n's ad hoc range, they can collaborate via the ad hoc connection. This method is feasible because mobile devices today are equipped GPS devices and localization using Wifi access points can be applicable for indoor environment. For the second approach, after n starts its download, n can periodically check the neighborhood for iShare neighbors who are downloading the same content. This method is energy-consuming. In the future, mobile phones can use GPS IIF for an accurate localization both indoor and outdoor [1].

C. Ad hoc data exchange

After downloading the first segment from the cellular link, the iShare node n starts using both cellular and P2P links. This section focuses on states 3,4,5 in Figure 2.

1) Content advertisement and request: For an iShare node n, the ad hoc channel is used to advertise its available segments and request its missing segments. Particularly, n periodically broadcasts a HELLO message, which is in bit vector format as shown in Figure 3. In this figure, the index of the bit, starting from 1 to 16, is the *segment id*; for example, the 12^{th} index denotes the 12^{th} segment. The HELLO message represents the latest downloaded segments of n. Notice that the length of a HELLO message is the number of segments of the file. The HELLO message can be used as both segment advertisement and segment request, where 1 represents one downloaded segment and 0 represents a missing segment in n's memory. Thus, the HELLO message efficiently reduces ad hoc network contention.

2) The mesh structure and data exchange: When the node n keeps its ad hoc interface on, n attends an ad hoc mesh of iShare nodes. The mesh has following characteristics. First, the mesh structure is formed automatically since the co-located nodes are within the ad hoc communication range. This incurs little construction/maintenance cost since nodes only need to keep the one-hop neighbor list [2]. Second, any two one-hop mesh neighbors can exchange data whenever they are within the communication range. This one-hop communication adapts the network dynamics and fits the tit-for-tat incentive mechanism very well (See Section III-E).

D. Downloading data from cellular link

The iShare nodes exchange downloaded data via ad hoc connections. At the same time, they continuously download segments from the cellular link. This section presents how iShare nodes download segments from cellular link and focuses on states 2,3,6 in Figure 2.

2	0	0	4	1	3	0	0	0	3	0	4	3	0	0	0
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

Figure 4. The aggregated HELLO message format

To reduce the download from the cellular link, iShare nodes utilize HELLO messages to download the best segments. In particular, node n decides to download its missing segment s if s is the least available segment in n's neighborhood. To do so, n aggregates all the latest HELLO messages received from its one-hop neighbors to create an aggregated HELLO message as shown in Figure 4. In this figure, each square can be a byte instead of a bit like the HELLO message. Notice that for each neighbor n_1 of n, n only keeps the latest HELLO message received from n_1 for the most updated available segments of n_1. In Figure 4, a square represents a segment s with the number of available copies of s in n's neighborhood. Node n downloads a missing segment whose number of copies is least founded in the aggregated HELLO message. If there exist more than one missing segments with equal number of available copies, n downloads one at random. For example, if n's sent HELLO message is in Figure 3 and n's aggregated HELLO message is in Figure 4, then n may download segment 3 from the cellular link since segment 3 is missing at n and n's neighbors. Whenever n finishes downloading a segment s, n is in the state 3 in Figure 2. Here, n inserts s into its memory and continues downloading its missing segments from the cellular link.

The random segment selection presented above parallelizes the download among mesh members. Particularly, mesh members concurrently download different segments and exchange these segments via ad hoc connections as shown in Figure 1. The aggregated HELLO message thus minimizes redundant downloads from the content server and reduces the load on the cellular link.

E. The tit-for-tat incentive mechanism

The mesh ad hoc structure efficiently parallelizes the download and reduces the load on the content server. However, for an efficient iShare protocol, we need to design an incentive mechanism to motivate the collaboration of iShare nodes. Particularly, we focus on states 3 and 7 in Figure 2.

1) The tit-for-tat period: For two iShare nodes n_1 and n_2, tit-for-tat means if n_1 gives c segments (needed by n_2) to n_2 then n_2 will give c segments (needed by n_1) to n_1. Applying tit-for-tat, iShare nodes divide time into equal-sized periods, called tit-for-tat period (TTP). The TTP is then used as follows. Given two one-hop mesh neighbors n_1 and n_2, n_1 uses the current TTP to receive segments from n_2 so that n_1 can send n_1's segments back to n_2

---

[2] HELLO message is the overhead of data advertisement, not the mesh structure

in the next TTP. Also, n_1 sends segments to n_2 in the current TTP so that n_2 can send segments to n_1 in the next TTP. Notice that the length of a TTP is longer than that of the HELLO message broadcast period since nodes need to update available segments to perform tit-for-tat.

During a TTP, a node n counts the number of segments received from its neighbors. Given two neighbors n_1 and n_2, during a TTP, if n_2 sends 15 segments (needed by n_1) to n_1, then n_1 has counter $c_2 = 15$, corresponding to n_2. At the end of the TTP, n_1 is at the state 7, if n_1 has more than 15 segments that n_2 needs (known from n_2's HELLO message), n_1 only sends 15 random segments to n_2 via *unicast ad hoc connection*. If n_1 has less than 15 segments needed by n_2, n_1 sends them all to n_2. Here, the unicast connection is used to obtain the fair collaboration between neighbors and ensure a reliable ad hoc data exchange. In the next section, we present how to bootstrap and adapt the tit-for-tat under network dynamics.

2) Applying the tit-for-tat: The tit-for-tat mechanism presented above encourages iShare nodes to collaborate. However, it may not work effectively if the node neighborhood changes frequently, since a new pair of one-hop neighbors needs to start tit-for-tat from scratch. Thus, we present two techniques to bootstrap and adapt the tit-for-tat under network dynamics.

First, we turn on the promiscuous mode of the wifi interface so that the iShare node n can potentially overhear messages, which are destined to n's neighbors in the above unicast communication of the tit-for-tat. By doing so, n opportunistically receives more data from the ad hoc channel. Of course, when the network is dense or congested, the overheard messages might be dropped and n misses the chance. Second, during a TTP, n broadcasts in the ad hoc channel a small number of its segments whose available copies are least in n's neighborhood.

Using the promiscuous mode and broadcast mechanism, iShare nodes improve the “tit” step of the tit-for-tat so that they exchange more segments in the “tat” step. These two techniques allow the new neighbors to exchange data effectively under network dynamics, without restarting the tit-for-tat from scratch. These techniques also enable new nodes to join the downloading group smoothly since they are given several segments for free. However, to exchange data with the old nodes effectively, new nodes need to download new segments via the cellular link. Otherwise, they become “iShare selfish nodes” and their downloading times might be longer as shown in Figure 7(a).

F. Completed state

When n finishes downloading the entire file, n switches to the completed state. Here, there are two options. If the iShare node n is rational, n leaves the mesh ad hoc network and stops all ad hoc communications. If the iShare node n is collaborative, n may stay for a certain period to support other downloading nodes.

IV. EVALUATION

We evaluate the performance of iShare where cellular nodes download a file via cellular link and at the same time they exchange downloaded data via the wifi interface. We use Network Simulator 2 (NS2) to simulate cellular cells and mobile nodes with the settings in Table I. Here, the segment size is 4KB since from our simulation we observe that smaller segment incurs longer HELLO message while bigger segment causes more ad hoc collision. A node has two interfaces: cellular link and IEEE 802.11b ad hoc link. We use RTS/CTS for unicast ad hoc communication. For the cellular technology, we use 1xEV-DO (Evolution-Data Only) with a peak data rate of 2.4Mbps. We implement the Proportional Fair Scheduler of the cellular network [3] in NS2. We evaluate two metrics: “average downloading time” and “average number of downloaded segments”. The former is the average (AVG) period for an iShare node to finish downloading a file. The latter is the average number of segments a non-iShare node (or background node) can download via *only* the cellular unicast link for a given period. In our context, the background node is the node which downloads the content from the cellular link but it does not use iShare protocol.

In our plots, broadcast channel means the base station broadcasts the file to all downloading nodes with a fixed rate of 208.4 Kbps and the broadcast channel takes 25% of cell bandwidth. Cellular unicast channel means downloading nodes only use cellular unicast link (without ad hoc link) to download. In our simulation, we set the HELLO broadcast period 2 seconds and TTP 7 seconds. The number of broadcast (tit) segments is 1% of the file size. The default configurations of our plots are: 15 iShare nodes, node speed is 5 (m/s), 30 background nodes, file size is 3000 4KB-segments. Node transmission range is 125 (m) since in reality the transmission range of wifi interface is much less than its theoretical range (e.g., 250 m). We run each simulation 10 times and plot the mean. More detail of the experiment results can be found in our technical report [11].

Field	**Value/Unit**
Segment size	4KB
File size	[1000...6000] segments
Node ad hoc transmission range	125(m)
Base station radius	750(m)
Mobility model	Random Way Point
Node speed (Mobility-NS2)	[1,3,5,7,11] (m/s)
Pause time (Mobility-NS2)	5 (seconds)

Table I
SIMULATION SETTINGS

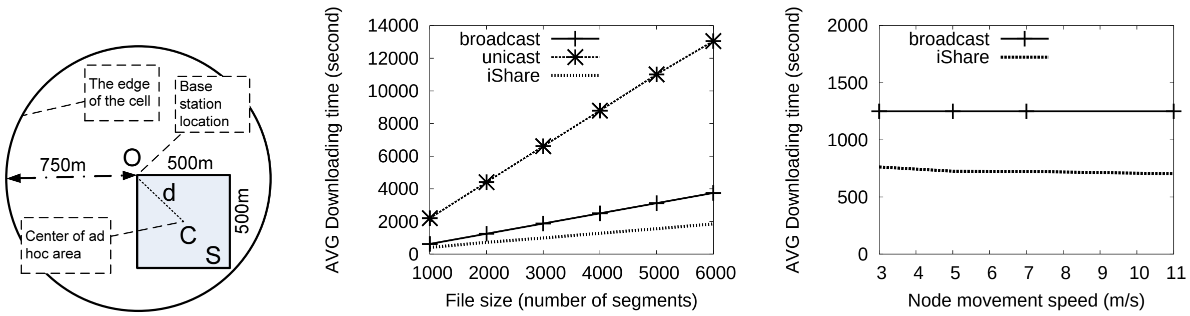

(a) Experiment Area (b) Impact of File size (c) Network Dynamics

Figure 5. Performance of iShare nodes in a cell

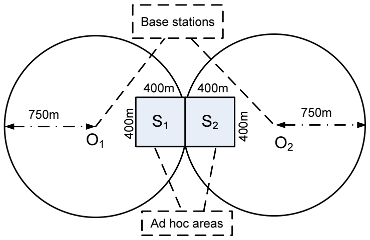

Figure 6. Experiment area of one spanning group

A. The Single Downloading Group

Here, the base station is at O and nodes are generated at random in the square S as shown in Figure 5(a).

1) Impact of file size and Network Dynamics: Figure 5(b) shows that the average downloading time linearly increases when the file size increases. iShare outperforms both broadcast and unicast channels since ad hoc connection accelerates the downloads of iShare nodes. Figure 5(c) shows when node speed varies, iShare remains quite stable (even slightly better for more dynamic networks as nodes can exchange more data with more neighbors) and always better than broadcast channel. This confirms the robustness of the mesh structure created by iShare to network dynamics.

2) Performance of iShare selfish nodes: In this simulation, we have 15 iShare nodes and we vary the number of iShare selfish nodes from 1 to 15. In our context, iShare selfish nodes download data via unicast cellular links and at the same time only overhears (due to tit-for-tat) data from wifi channel without sharing their data. Figure 7(a) shows that when more iShare selfish nodes exist, their downloading time increases noticeably. Meanwhile, the downloading time of iShare nodes only increases slightly due to the higher load on the cell tower resulting from higher number of selfish nodes in the cellular cell. In other words, iShare selfish nodes suffer from their own existences or iShare limits the selfishness.

3) A tree-based protocol vs. iShare: We compares the performance of iShare and a tree-based protocol of co-operative downloading nodes [7], which is implemented as follows. Downloading nodes elect the proxies whose distances to the base station are shortest (implying the best cellular channel downloading rate). These proxies download segments from the base station and broadcasts the segments to tree members, which replay the segments through the tree of ad hoc nodes. We make sure the tree is always connected. Here, iShare and the tree-based protocol both use the same simulation settings. Figure 7(b) shows that iShare protocol consistently outperforms the tree-based protocol. In this figure, k denotes number of concurrent proxies of the tree. When $k = 1$, the tree-based protocol performs much worse than cellular broadcast channel. When k increases, the tree of k roots performs noticeably better, although always worse than iShare.

4) A Spanning Group: Here, we first create a group g_1 of 10 iShare nodes within the square S_1 in Figure 6. Then, we create a group g_2 of 20 iShare nodes within two squares S_1 and S_2 spanning two adjacent cells. We call g_2 a spanning group, which has a significant number of low-downloading-rate nodes (nodes are close and at the edge the cell). Nodes in g_1 and g_2 download segments from their current base stations and exchange segments via the ad hoc channel. Here, node speed is 7 m/s and each cell has 30 background users. Figure 7(c) shows that g_2 consistently outperforms g_1 and broadcast channel. This result has several implications. First, iShare provides an efficient method to reduce downloading time, especially for low-data-rate nodes at the boundary of the cell. Second, iShare nodes can continuously obtain data via the ad hoc channel during their cellular handoff periods. Finally, iShare offers the multi-homing download for a spanning group, where group members download content from different/adjacent base stations and exchange segments via ad hoc connections to improve downloading throughput.

B. Multiple Downloading Groups

We use the settings in Table I and Figure 5(a). We have 3 groups within S, each group has 10 nodes and downloads a different file. Thus, there is no ad hoc inter-group communication. We assume cell bandwidth for the broadcast channel

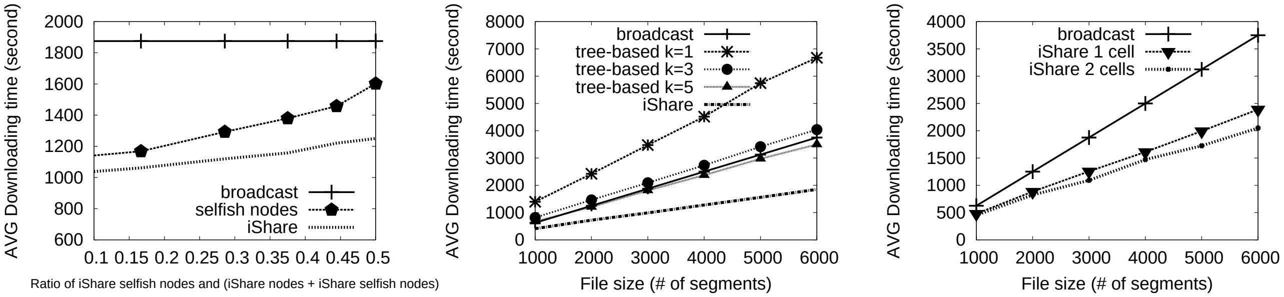

(a) Performance of iShare selfish nodes

(b) A tree-based protocol vs. iShare

(c) One group spanning two adjacent cells

Figure 7. Performance of selfish nodes, comparison of iShare and tree-based protocols, and performance of a spanning group

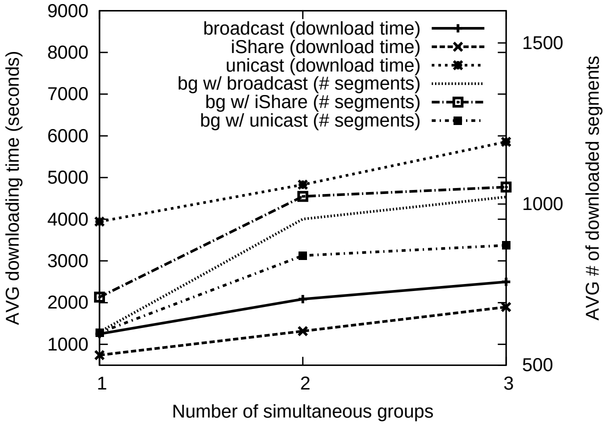

Figure 8. Performance of multiple groups in one cell

is 25%, 15%, 12.5% corresponding to 1, 2, 3 simultaneous groups; or 25%, 30%, 37.5% aggregated bandwidth for the broadcast channel. Here, the file size is 3000 segments, node speed is 5 (m/s), number of background users is 30. Figure 8 shows that the downloading time of iShare nodes increases when more groups exists since iShare nodes suffer from a higher contention in the ad hoc channel. Moreover, due to the promiscuous mode, iShare nodes receive redundant messages from nodes in other groups, which may collide with the desired overheard messages from nodes in the same group. As a result, the performance of tit-for-tat degrades. However, iShare performs consistently better than broadcast channel. Figure 8 also shows background nodes download more data (the right y-axis) from cellular link if mobile nodes use iShare protocol (i.e., bg w/ iShare). This result suggests the use of multiple channels for multiple ad hoc groups in the same cellular cell.

V. CONCLUSIONS

We have developed an efficient sharing protocol (iShare) that combines different wireless interfaces on mobile devices to improve content dissemination services. The results from simulation in NS2 show that iShare significantly outperforms alternative schemes based on cellular broadcast channels, cellular unicast channels, or tree-based protocols. Furthermore, the results confirmed that "tit-for-tat" mechanism succeeds in countering selfishness user behavior and adapting very well to network dynamics. Finally, the obtained results showed the multi-homing download feature for groups spanning over adjacent cellular cells.

REFERENCES

[1] GPS IIF. http://www.boeing.com/defense-space/space/gps/index.html.

[2] Skyhook. http://www.skyhookwireless.com.

[3] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi. CDMA/HDR: A bandwidth-efficient high-speed wireless data service for normadic users. *IEEE Communications Magazine*, 2000.

[4] R. Bhatia, L. E. Li, H. Luo, and R. Ramjee. ICAM: Integrated cellular and ad-hoc multicast. *IEEE Transactions on Mobile Computing*, 5(8):1004–1015, August 2006.

[5] H. Hanano, Y. Murata, N. Shibata, K. Yasumoto, and M. Ito. Video ads dissemination through wifi-cellular hybrid networks. In *Proceedings of Percom*, 2009.

[6] H.-Y. Hsieh and R. Sivakumar. On using peer-to-peer communication in cellular wireless data networks. *IEEE Trans. on Mobile Computing*, pages 57–72, 2004.

[7] H. Luo, R. Ramjee, P. Sinha, L. Li, and S. Lu. The design and evaluation of unified cellular and ad hoc networks. *IEEE Transactions on Mobile Computing*, 2007.

[8] T. K. Madsen, Q. Zhang, F. Fitzek, and M. Katz. Exploiting cooperation for performance enhancement and high data rates. *Journal of Communications*, 2009.

[9] K. S. and M. M. Efficient mobile access to internet data via a wireless peer-to-peer network. In *Pervasive Computing*, 2004.

[10] M. Stiemerling and S. Kiesel. Cooperative p2p video streaming for mobile peers. In *Proceedings of IEEE ICCCN*, 2010.

[11] L. Vu, I. Rimac, V. Hilt, M. Hofmann, and K. Nahrstedt. iShare: exploiting opportunistic ad hoc connections for improving data download of cellular users. In *UIUC Technical Report*, 2009.