Energy-Aware Web Caching for Mobile Terminals

Françoise Sailhan, Valérie Issarny

► To cite this version:

Françoise Sailhan, Valérie Issarny. Energy-Aware Web Caching for Mobile Terminals. International Conference on Distributed Computing Systems: ICDCS 2002, 2002, Vienna, Austria. pp.820-825. inria-00415133

HAL Id: inria-00415133

<https://hal.inria.fr/inria-00415133>

Submitted on 10 Sep 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Energy-aware Web Caching for Mobile Terminals

Françoise Sailhan, Valérie Issarny

Inria -Rocquencourt

Domaine de Voluceau, Rocquencourt, BP 105, 78153 Le Chesnay Cédex, France

Francoise.Sailhan, Valerie.Issarny@inria.fr

Abstract

Terminal's latency, connectivity, energy and memory are the main characteristics of today's mobile environments whose performance may be improved by caching. In this paper, we present an adaptive scheme for mobile Web data caching, which accounts for congestion of the wireless network and energy limitation of mobile terminals. Our main design objective is to minimize the energy cost of peer-to-peer communication among mobile terminals so as to allow for unexpensive Web access when a fixed access point is not available in the communication range of the mobile terminal. We propose a collaborative cache management strategy among mobile terminals interacting via an ad-hoc network. We further provide evaluation of the proposed solution in terms of energy consumption on mobile devices.

1. Introduction

The last decade has seen the rapid convergence of two pervasive technologies: wireless communication and the Internet. The resulting mobile Internet a priori enables users to easily access information anytime, anywhere. However, we have not reached the point where anywhere, anytime Internet access is actually offered. This paper addresses the above issue, concentrating more specifically on Web caching in a mobile environment to allow for Web access, without requiring availability of an infrastructure in the nearby environment. There exists two different ways of configuring a mobile network: *infrastructure-based* and *ad-hoc-based*. The former type of network structure is the most prominent, as it is in particular used in both Wireless LANs (e.g., IEEE 802.11) and global wireless networks (e.g., GSM, GPRS, UMTS)¹. An infrastructure-based wireless network uses fixed network access points (known as *base stations*) with which mobile terminals interact for communicating, i.e., a base station forwards messages that are sent/received by mobile terminals. One limitation of the infrastructure-based configuration is that base stations con-

stitute bottlenecks. In addition, it requires that any mobile terminal be in the communication range of a base station. However, this comes at a high cost for network providers, and is only supported if either communication is charged (i.e., global networking) or there is the will to ease access to information technology (e.g., local networking in buildings). The ad-hoc-based network structure alleviates this problem by enabling mobile terminals to cooperatively form a dynamic and temporary network without any pre-existing infrastructure. Hence, it is a very cheap solution. In general, ad-hoc and infrastructure-based networking should be seen as complementary rather than as competitive. Ad-hoc networking is much convenient for accessing informations available in the local area, and possibly reaching a WLAN base station, which comes at no cost for users. Ultimately, the user may decide to pay for communication using wireless global networking facility, if the connectivity using the WLAN happens to be bad. This paper concentrates on improving the Web latency using a WLAN, exploiting both ad-hoc and infrastructure-based capabilities of the network.

The issue that we are addressing is on setting up an ad-hoc network of mobile terminals that cooperate to exchange Web pages, hence enabling Web access at no financial cost for mobile users. In that context, it is crucial to account for the specifics of mobile terminals. The capacity of batteries goes up slowly and all the powerful components that will be soon available (e.g., LCD screens, 3D graphics, high performance processor) reduce battery life. In particular, communication is one of the major sources of energy consumption [6]. Thus, it is mandatory to devise adequate solutions to energy saving on the mobile terminals, for all the constituents of the mobile environment, i.e., hardware, network operating system, and distributed software.

This paper introduces such a distributed application software, which implements ad-hoc cooperative Web caching among mobile terminals. The proposed solution aims at improving the Web latency on mobile terminals while optimizing associated energy consumption. The solution accounts for both the capacities of mobile terminals and the network features; it comprises: (i) a cooperative caching protocol among mobile terminals that builds upon an exist-

¹This network structure is also referred to as the *base station mode* in the IEEE 802.11 WLAN, and the *BSS* in GPRS.

ing ad-hoc network protocol, and (ii) a local caching strategy for the mobile terminal. The next section discusses related work, addressing background in the area of cooperative Web caching and Web access from mobile terminals. Section 3 then introduces the proposed ad-hoc cooperative Web caching protocol, and is followed in Section 4 by the presentation of the local caching implemented on the terminals. Section 5 provides an evaluation of our proposal, giving the energy consumption associated with cooperative caching. Finally, Section 6 concludes with a summary of our contribution.

2. Related Work

The ever growing popularity of the Web and the resulting poor latency for users have given rise to huge effort on improving the Web latency, which mainly lies in the introduction of Web caching protocols. Due to network topology, the idea of making network caches cooperate has emerged. The hierarchical approach is pioneering and lies in introducing a cache on every network node, the system's hierarchical structure coming from the national networks' hierarchical organization. A cache locates a missing requested object by issuing a request to the cache at the hierarchy's upper level. The process is iterated until either the object is found or the root cache is reached, which may ultimately involve contacting the object's server. The object is then copied in all the caches contacted as it is returned to the client. A transversal system enriches the hierarchy by integrating a set of sibling caches that are close in terms of latency time, at each hierarchical level. On a cache miss, a cache not only contacts its ancestral cache but also its siblings. A number of transversal cooperative caching protocols have been proposed in order to minimize the number of messages that are exchanged among sibling caches to retrieve cached objects [2]. Proposed solutions all amount to maintain a partial knowledge of the objects that are cached on siblings [4, 12].

In the context of mobility, proxy caches are used not only to cache and retrieve documents but also to manage user mobility. These proxies are access points to the Internet for mobile terminals. Proxies implement functionalities dedicated to the transfer of data to mobile terminals, such as compression, filtering, format conversion [11][1].

To the best of our knowledge, Web caching for mobile terminals has only been studied in the context of proxy caches, and hence for infrastructure-based mobile networks. As raised earlier, complementing such solutions with ad-hoc-based ones will allow for both enhanced connectivity and Web access at low cost. In that context, the caches of the mobile terminals cooperate in a way similar to proxy caches in transversal cooperative cache systems.

3. Cooperative Caching in Ad-hoc Networks

Ad-hoc routing protocols are implemented over a base WLAN (typically, IEEE 802.11) and manage the routing of messages among mobile terminals. These protocols differ in the way they manage the routing table (§ 3.1). Using the ad-hoc routing protocol that offers the best trade-offs in terms of energy consumption and response time, we propose a specialization of the protocol that is specifically aimed at handling remote access to Web pages (§ 3.2). We then introduce our ad-hoc cooperative caching protocol, which has primarily been designed to minimize energy consumption (§ 3.3).

3.1. Ad-hoc Networking

The main issue to be addressed in the design of an ad-hoc routing protocol is to compute an optimal *communication path* between two mobile terminals. This computation minimizes the number of control messages exchanged among mobile terminals in order to avoid network congestion.

There exist two types of ad-hoc protocols: *proactive* and *reactive*. Proactive protocols (e.g., OLSR [3]) update their routing table periodically. Reactive protocols (e.g., AODV [13], DSR [10]) do not take any initiative for finding a route to a destination, before the information is needed, and thus a priori reduce the network load due to the traffic of control messages. ZRP [7] is a hybrid protocol that combines the reactive and proactive modes. The design rationale of ZRP is that it is considered advantageous to accurately know the *neighbours* of any mobile terminal (i.e., mobile terminals that are accessible in a fixed number of hops), since they are close. Hence, communicating with neighbours is less expensive and neighbours are most likely to take part in the routing of the messages sent from the terminal. As a result, ZRP implements: (i) a proactive protocol for communication with mobile terminals in the neighbourhood, and (ii) a reactive protocol for communication with the other terminals. With respect to a given mobile terminal, its neighbourhood is referred to as its *zone*.

We thus use ZRP over IEEE 802.11, as the base ad-hoc protocol for realizing ad-hoc cooperative caching among mobile terminals. Mobile terminals belonging to the zone of a given terminal then form a cooperative cache system for this terminal since the cost for communicating with them is low both in terms of energy consumption and message exchanges [8]. However, cooperative caching must not be restricted to the mobile terminals belonging to the zone: low-cost reachability of a base station must be accounted for as well as knowledge of a terminal that does not belong to the zone but that is likely to store a requested Web document given commonalities in performed Web accesses.

3.2. Ad-hoc Communication for Web Caching

A mobile terminal may get Web data that are not cached locally through two communication paths: (i) using the

infrastructure-based mode, the terminal may interact with the nearby base station, which forwards the request to the Web, (ii) using the ad-hoc-based mode, the terminal may request for the data to the mobile terminals in its communication range (i.e., accessible in one hop in a base WLAN or in a number of hops using some ad-hoc network protocol). The ad-hoc-based mode must be enabled for the case where a base station is not reachable in one hop. In this case, a base station can still be reachable in a number of hops, thanks to mobile terminals forwarding the requests. For instance, UMTS supports such a mechanism to extend the range of the infrastructure². Let N be the number of hops that are necessary to access a base station, then any mobile terminal that is at a distance greater or equal to N is not contacted to get a document. Figure 1 depicts the case where the mobile terminal A reaches the base station D in 3 hops, using the mobile terminals B and C for routing the request. Then, if either a mobile terminal in the zone of A (e.g., B belonging to the path leading to the base station or any other terminal in the zone) or a known mobile terminal located outside the zone but at a lower distance than the base station D (e.g., C that is in the path leading to D or E that does not belong to the path) holds the requested document in its cache, it returns it to A . Otherwise, the request reaches D , and D forwards it to the Web. We get the following ad-

Figure 1. Getting Web data

hoc communication protocol over ZRP, to retrieve a remote Web object W , with respect to a given mobile terminal A (see Figure 1):

In-zone communication:

If a base station is in the zone of A , then A requests for W through the base station only. Otherwise, A broadcasts the request message for W to the mobile terminals in the zone of A , incurring a low energy cost since the routing table contains the necessary information.

Peer-to-peer communication: If W is not cached by any of the mobile terminals in the zone of A , then a peer-to-peer communication scheme is realized with mobile terminals that are known to share interests with A (see § 3.3) and that are at a distance that is less than the one between A and the nearest base station. Mobile terminals outside the zone of A are basically known through two ways: (i) they belong to the path used to reach the nearest base station, (ii) they were previously either in the zone or in the path used to reach the

base station. The request for W is ultimately forwarded to the nearest base station.

Based on the above, the communication and energy costs associated with getting a Web object is kept to a minimum: (i) broadcast is within a zone and is thus unexpensive by construction of ZRP, (ii) peer-to-peer communication occurs with mobile terminals that are both the most likely to store a requested object and closer than a base station.

3.3. Ad-hoc Cooperative Caching

Web data are distributed among the mobile terminals according to Web accesses performed by their user. Without a proxy-type architecture that centralizes requests, local statistics are relied on for a mobile terminal A to identify mobile terminals that are likely to store a Web object requested on A . Such statistics are maintained on A through a *terminal profile* for every mobile terminal with which A interacts. The terminal profile is characterized by a value that counts the number of times the corresponding mobile terminal either is known to cache an object requested by A or requested for an object to A that A had in its cache. This value is used to identify the mobile terminals with which peer-to-peer communication is undertaken (see previous subsection). The list of known mobile terminals outside the zone and that are at a distance less than a base station are weighed according to the value of $\mathcal{F} = \text{terminal profile} \div \text{hops}$ where the number of hops, *hops*, is obtained from the routing table. Mobile terminals for which the value of $\mathcal{F}$ is the greatest are first contacted and the process is iterated until a *hit* message is received or there is no more mobile terminals eligible for the request. In addition to the management of *terminal profiles* to identify mobile terminals that share common interests, we must account for the heterogeneity of the terminals' capacity (i.e., battery, processing, storage, communication). For instance, for two mobile terminals that are equally likely to store a requested object, it is better to contact the one that has the greatest capacity.

A mobile terminal that receives the request for W and caches it increments its local value of A 's terminal profile. If the terminal is further willing to cooperate (e.g., absence of energy safeguarding or of security policy enforcement), it returns a *hit* message, which embeds:

- (i) *TTL* that gives the Time To Leave field of the document.
- (ii) *Capacity* that characterizes the capacity of the terminal to handle requests, whose value is in the range $[0..1]$, 1 denoting the highest capacity³.

For every *hit* message that it receives, A increments the *terminal profile* of the sender. Among the mobile terminals that replied by a *hit* message, A selects the terminal from which W should be obtained, that is the one that maximizes

³Currently, we use the percentage of the energy budget that is left to set the value of *Capacity*. It is part of our future work to investigate a more accurate way of computing *Capacity*, in particular accounting for the various resources.

²<http://www.3gpp.org>

the following function⁴: $\mathcal{R} = \text{Capacity} \times (\lambda \times \text{TTL} + \mu \times \text{hops})$ where the value of λ and μ are set so as to favour communication with the closest nodes.

To minimize both network load and energy consumption, we do not use *miss* messages for mobile terminals to notify that they do not cache a requested object. Hence, we need to use timeouts to detect the absence of a requested object. The value of the timeout is set according to the greatest number of hops that are involved to interact with the mobile terminals to which the object is requested, together with the current network load. Upon expiration of the timeout, if *hit* messages have been received, the Web object will be requested to the mobile terminal that maximizes $\mathcal{R}$. Otherwise, the next iteration of the cooperative caching protocol is processed (i.e., from the broadcast step to the peer-to-peer iterative steps). For the case where a *hit* message is received after timeout expiration, while the object is still not retrieved, the message is accounted for in the current step of the protocol.

4. Local Caching

Our ad-hoc cooperative caching protocol is complemented by a local caching strategy that is adaptive according to the current capacity of the terminal (i.e., available energy and network connectivity). We weigh every cached document according to both its probability of being accessed in the future and the energy cost associated with getting remotely the document. Then, documents with the lowest weighs are those that are removed from the cache when the cache gets full. The document weigh is computed according to the following criteria:

Popularity. The *Popularity* value serves approximating the probability of future access, both on the terminal and from remote terminals, as enabled by the cooperative caching protocol. The probability is approximated according to the number of times the document has been requested since it has been cached.

AccessCost. The *AccessCost* value gives an estimate of the energy cost associated with getting the document remotely if it is to be removed from the cache. This cost varies depending on whether: a base station is accessible in the zone of the terminal, the document is cached on a mobile in the zone of the terminal, communication out of the zone is required to retrieve the document. The value of the access cost is computed according to the energy consumption associated with intra-zone and inter-zone communication (see § 5). It is further assumed if the document was obtained from a terminal that is still in the zone, as identified using the routing table.

Coherency. A document is valid for a limited lifetime, which is known using the *TTL* field. However, when the

⁴For the case where the selected terminal is no longer accessible, e.g., due to energy safeguarding, the request will be sent to the next eligible terminal and will be so until the page is received.

energy remaining on the terminal is low, it is better to favour energy saving over the accuracy of the document. Hence, the value of *Coherency* is equal to $\nu \times \text{TTL}$ where ν increases as the available energy decreases.

We get the following function to compute the document weigh:

$$W = \alpha \times \text{Popularity} + \beta \times \text{AccessCost} + \gamma \times \text{Coherency} + \delta \times \text{Size}.$$

The values of α , β , γ and δ are set so as make the values of *AccessCost*, *Popularity*, *Size*, and *Coherency*, decreasingly prominent factors for deciding whether a document should be kept in the cache. The accurate definition of α , β , γ , and δ is part of our future work, where we are in particular interested in a definition that is adaptive according to the evolution of the mobile environment. Notice that our $\mathcal{W}$ function offers similarities with hybrid replacement algorithms on stationary hosts that were proposed in the literature (e.g., [9]), our function differs in that energy saving is a prominent criterion.

5. Analysis

Control messages generated by ZRP and messages induced by ad-hoc cooperative caching affect the network traffic and energy consumption on mobile terminals. Performance of ZRP has been evaluated in [8, 7] using event-driven simulation; this evaluation is gauged by considering the control traffic generated by ZRP, which is reported in terms of numbers of ID fields transmitted, and the time taken for route discovery. Knowing the traffic received by a mobile terminal, the energy consumption for mobile terminals is easily evaluated. For this reason, in the following, we evaluate energy consumption associated with ad-hoc cooperative caching, as the sum of the energy consumption induced for the various mobile terminals that are involved (both in and outside the zone) in the cooperation, which adds to the energy cost induced by ZRP. We do not consider the computation cost (i.e., local cache management) since it is negligible compared to the energy cost of communication, and it is induced by any local cache management. We further use the following wireless interface for our evaluation: 2.4Ghz DSSS lucent IEEE 802.11 WLAN 2Mbps.

Focusing on the energy cost associated with communication, the cost associated with the emission of one message is the sum of the cost of: (i) Emission (resp. reception) for the sender (resp. destination) node, (ii) reception and emission for nodes forwarding the message, (iii) reception for non-destination nodes (i.e., terminals that receive messages due to their location, although they are not involved in the message routing).

The IEEE 802.11 protocol provides the following collision avoidance mechanism for point-to-point traffic. Prior to any point-to-point transmission, the sender broadcasts a RTS (Request To Send) control message, which specifies the destination node and the data size (for duration). The sender then waits for a CTS (Clear To Send) message from the des-

Mobile	Energy consumption ($\mu W.sec$)	m ($\mu W.sec/byte$)	p ($\mu W.sec/byte$)
Sender X	$\varepsilon_{send} = m_{send} * size + p_{send}$	$m_{send} = 1.9$	$p_{send} = 454$
Destination A	$\varepsilon_{dest} = m_{dest} * size + p_{dest}$	$m_{dest} = 0.5$	$p_{dest} = 356$
Non-destination nodes			
in range of sender X & destination A	$\varepsilon_{AX} = m_{AX} * size + p_{AX}$	$m_{AX} = -0.22$	$p_{AX} = 210$
in range of sender X	$\varepsilon_X = m_X * size + p_X$	$m_X = -0.04$	$p_X = 90$
in range of destination A	$\varepsilon_A = m_A * size + p_A$	$m_A = 0$	$p_A = 119$

Table 1. Energy consumption on nodes for point-to point communication

mination node. Once it receives the CTS, the sender sends the data message. Finally, the destination node sends an ACK message upon the reception of the data message. Therefore, the energy consumed by any mobile terminal for sending, receiving or discarding a message is given by the linear equation [5]: $\varepsilon = m * size + p$; where $size$ is the message size, and m (resp. p) denotes the incremental (resp. fixed) energy cost associated with the message. Table 1 gives the energy cost, relative to the size of the message, for the destination nodes (i.e., the actual sender and destination nodes and the forwarding nodes that act as both sender and destination nodes) and also for non-destination nodes. Non-destination nodes in the range of the sender receive RTS messages and thus enter a reduced energy consumption mode during data emission; this leads to have a negative value for m_{AX} and m_X since the energy consumption is less than the one in the idle mode. Finally, non-destination nodes in the range of the destination node but not the sender do not receive the RTS message, and thus cannot enter in the reduced energy consumption mode; this leads to have the incremental cost m_A equal to 0.

Let us now assume a network of 500 mobile terminals whose communication range is of about 250m, that is such that the mobile terminals are uniformly distributed in the network surface S with $S = 4000[m] \times 4000[m]$.⁵ In a zone (see Figure 2), all the mobile terminals consume the same energy. Thus, the overall energy consumption within a zone is the sum of the energy consumed by every mobile terminal in this zone (see Table 2). Then, the energy consumption of the overall network is the sum of the energy consumed per zone that is traversed. Figure 3 gives the energy consumption associated with the delivery of a message of 1Kb to the destination node according to the number of hops, which is the sum of the energy consumed on all the

⁵This network is taken as an example, as it is used in [8] for the evaluation of ZRP.

Number of mobile terminals in the network surface	500	600	700
Energy consumption of non-destination terminals ($\mu W.sec$)	1389	18959	2136
ratio: $\frac{\varepsilon_{send} + \varepsilon_{dest} + \varepsilon_{forwarding\ mobile\ terminals}}{\varepsilon_{overall\ non-destination\ mobile\ terminals}}$	9	7	6.11

Table 3. Energy consumption according to the network density.

nodes involved in the communication. The figure clearly demonstrates that the energy consumption increases with the number of hops. This directly follows from the resulting increase of both mobile terminals forwarding the message and non-destination nodes receiving control messages. Table 3 further evaluates the impact of the network density on energy consumption. For a constant number of hops ($=3$), we see that increased density of mobile terminals results in additional energy consumption for non-destination nodes. But, the ratio highlights the weak impact of message reception on non-destination nodes, on the overall energy consumption. Indeed, for 600 mobile terminals and a destination node at 4 hops of the sender, the energy consumed by non-destination nodes is about 6 times less than the energy consumed by the sender, the receiver and the 3 forwarding mobile terminals.

Figure 2. Mobile terminals in the range of the sender and of the destination node.

Figure 3. Energy consumption in the network for the retrieval of Web data.

Having examined the energy consumption associated with data delivery, we now focus on the energy consumption induced by our ad-hoc cooperative caching protocol. A request message for a Web page includes broadcasting (with mobile terminals in the requester's zone) and peer-to-peer communication (with mobile terminals sharing the same interests). Before broadcasting a message, the sender listens to the channel; if no traffic is detected, the message is broadcasted. The energy consumed by the sender X (resp. destination A) is given by the equation $\varepsilon_{brsend} = 1.9 \times size + 266$ (resp. $\varepsilon_{brdest} = 0.5 \times size + 56$). The higher energy consumption associated with peer-to-peer communication (see Table

Mobile terminal in range of	Surface	number of mobile terminals	Total Energy for a zone
sender X & destination A	$S_{AX} = \theta(r^2 - a_1^2) + \beta(R^2 - a_2^2)$	$n_{AX} = \frac{N S_{AX}}{S}$	$n_{AX} \varepsilon_{AX}$
the sender X	$S_X = \pi r^2 - S_{AX}$	$n_X = \frac{N(\pi r^2 - S_{AX})}{S} = \pi r^2 - \theta r^2 \sin^2 \theta + \beta R^2 \sin^2 \beta$	$n_X \varepsilon_X$
the receiver A	$S_A = \pi R^2 - S_{AX}$	$n_A = \frac{N(\pi R^2 - S_{AX})}{S} = \pi R^2 - \theta r^2 \sin^2 \theta + \beta R^2 \sin^2 \beta$	$n_A \varepsilon_A$

Table 2. Energy consumption in a zone.

Figure 4. Energy consumption for broadcast and peer-to-peer communication.

1) compared to broadcasting is due to the emission of control messages. Figure 4 gives the energy consumed by cooperative caching according to the number of hops that defines the zone. Precisely, Figure 4 gives the cost associated with broadcasting and peer-to-peer communication at a distance that adds one hop to the number of hops that defines the zone. We find that the cost added by broadcasting is weak compared to the number of mobile terminals contacted, and thus supports such an approach.

6. Conclusion

Mobile technology has reached a stage that enables foreseeing easy access to information technology anywhere, anytime. However, enabling mobility comes with limitations that mainly relate to unstable connectivity and limited energy. Thus, it is necessary to devise adequate solutions at the level of both software and hardware to mask as much as possible the limitations of mobile environments. This paper has addressed one such solution, which focuses on enabling Web accesses from mobile terminals. Currently, Web access is easy to realize if the mobile terminal is in the communication range of a base station of either a WLAN or a global wireless network. However, this cannot always be assumed due to the financial cost associated with the deployment of the underlying infrastructure. In addition, the systematic use of a global wireless network is quite costly for users. Instead of relying on a base station for accessing the Web, an alternative solution is to exploit ad-hoc networking, which allows for remote communication at no financial cost and also reaching a base station of a WLAN in a number of hops. In that context, we have proposed a cooperative Web caching system for ad-hoc networks, which enables mobile terminals to share Web pages. Our system lies in implementing on each mobile terminal: (i) an ad-hoc cooperative caching protocol that selects the mobile termi-

nals from which a requested page should be retrieved, in a way that minimizes both energy consumption and network load, (ii) a caching strategy that maintains the local cache in a way that minimizes resource consumption and masks disconnection. Preliminary assessment of our proposal has been addressed, by providing the energy cost that our system incurs for mobile terminals. We are currently working on further assessment of our proposal, which lies in the implementation of a simulator to thoroughly examine the behavior of our system. Experiment using the simulator will in particular serves tuning the various weighing functions that we use for cooperative caching.

Acknowledgments: The authors would like to thank Laurent Viennot for helpful discussion on ad-hoc networking.

References

- [1] C. Bennett. *Practical Wap, Developing applications for the Wireless Web*. Cambridge University Press, 2001.
- [2] C. Bowman, P. Danzig, and D. Hardy. The harvest information discovery and access system. *Computer Networks and ISDN Systems*, 28(1-2), 1995.
- [3] T. Clause, P. Jacquet, and A. Laouti. Optimized link state routing protocol. IEEE INMIC Pakistan, 2001.
- [4] L. Fan, P. Cao, J. Almeida, and A. Broder. Summary cache: a scalable wide-area web cache sharing protocol. *IEEE/ACM Transactions on Networking*, 8(3), 2000.
- [5] L. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. *In proc. of IEEE Infocom*, 5(8), 2001.
- [6] P. Gauthier, D. Harada, and M. Stemm. Reducing power consumption for the next generation of pdas: It's in the network interface. <http://http.cs.berkeley.edu/stemm/>, 1996.
- [7] Z. Haas. A new routing protocol for the reconfigurable wireless networks. In *Proc. of the IEEE ICUP'97*, 1997.
- [8] Z. J. Haas and M. Pearlman. The performance of query control schemes for the zone routing protocol. *ACM/IEEE Transactions on Networking* 9(4), 2001.
- [9] Z. Jinag and L. Kleinrock. Web prefetching in a mobile environment. *IEEE Personal Communications*, 5(8), 1998.
- [10] D. Johnson and D. Maltz. Dynamic source routing in ad hoc wireless networks. In Imielinski and Korth, editors, *Mobile Computing*. Kluwer Academic Publishers, 1996.
- [11] E. Markatos and C. Chronaki. A top-10 approach to prefetching on the web. *In proc. of the INET'98*, 1998.
- [12] J. Menaud, V. Issarny, and M. Banatre. A scalable and efficient cooperative system for web caches. In *IEEE Concurrency*. 8(3), July-September 2000.
- [13] C. Perkins. Ad-hoc on-demand distance vector routing. In *Proc. of IEEE WMCSA'99*, 1999.