article because it was not found through search. However, *FindAll* indexes all web pages, and therefore can retrieve the article for Bob. The search interface in *FindAll* can also potentially leverage the re-finding characteristics of other apps, including social networking apps, maps, and movie databases.

Similarly, *FindAll* is useful not only to re-find previously viewed web pages, but also to search within previously viewed web pages. For example, a user searching for hotels in Alaska may want to look for "Valet Parking" within the hotels that she previously viewed. By building an index, *FindAll* allows users to issue new queries to their local cache.

9. CONCLUSIONS AND FUTURE WORK

In this work, we developed *FindAll*, a local search engine that supports re-finding on mobile phones. To aid with its design and to better understand mobile re-finding, we conducted a user study with 23 users over 30 days. The study showed that users have diverse browsing and re-finding habits, and 45% of the URLs are re-found within 50 minutes. Therefore, the key design goal in *FindAll* is to design a search engine that indexes web pages locally and soon after the URL is first visited, to improve availability. The challenge is in designing a search engine for a resource-constrained mobile environment, while adapting to the user's re-finding behavior. To this end, *FindAll* indexes web pages only when the expected energy benefit of indexing outweigh the indexing cost. *FindAll* estimates the benefits of indexing by learning the re-finding patterns of each user and predicting the re-finding probability. We implemented *FindAll* over Android by adapting a publicly available search engine called Galago. Our evaluations show that *FindAll* reduces search latency by two-folds for users who re-find often. *FindAll* also reduces 3G data usage by up to 100 MB a month by serving over 40% of the web pages locally.

10. ACKNOWLEDGEMENTS

We thank our shepherd, Y. Charlie Hu, and all anonymous reviewers. Their reviews and comments greatly helped improve the presentation of this paper. We thank all the participants of our user study who helped us collect data. This work was supported in part by an NSF Computing Innovation Fellowship and NSF grant CNS-1217644.

11. REFERENCES

[1] Alexa top 500 websites.: `http://www.alexa.com/topsites`.

[2] Galago search engine: `http://www.galagosearch.org/`.

[3] Monsoon power monitor.: `http://www.msoon.com/`.

[4] D. Abrams, R. Baecker, and M. Chignell. Information archiving with bookmarks: personal web space construction and organization. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 41–48, 1998.

[5] E. Adar, J. Teevan, and S. T. Dumais. Resonance on the web: web dynamics and revisitation patterns. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 1381–1390, 2009.

[6] A. Aula, N. Jhaveri, and M. Käki. Information search and re-access strategies of experienced web users. In *Proceedings of the 14th international conference on World Wide Web*, WWW '05, pages 583–592, New York, NY, USA, 2005. ACM.

[7] A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting mobile 3g using wifi. In *MobiSys '10: Proceedings of the 8th international conference on Mobile systems, applications, and services*, pages 209–222, New York, NY, USA, 2010. ACM.

[8] G. Barish and K. Obraczka. World wide web caching: Trends and techniques. *IEEE Communications Magazine*, 38:178–184, 2000.

[9] E. Benson, A. Marcus, D. Karger, and S. Madden. Sync kit: a persistent client-side database caching toolkit for data intensive websites. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 121–130, New York, NY, USA, 2010. ACM.

[10] A. Cockburn, S. Greenberg, S. Jones, B. Mckenzie, and M. Moyle. Improving web page revisitation: analysis, design and evaluation. *IT & Society*, 1:159–183, 2003.

[11] E. Cuervo, A. Balasubramanian, D.-k. Cho, A. Wolman, S. Saroiu, R. Chandra, and P. Bahl. Maui: making smartphones last longer with code offload. In *MobiSys '10: Proceedings of the 8th international conference on Mobile systems, applications, and services*, pages 49–62, New York, NY, USA, 2010. ACM.

[12] T. Fagni, R. Perego, F. Silvestri, and S. Orlando. Boosting the performance of web search engines: Caching and prefetching query results by exploiting historical usage data. *ACM Trans. Inf. Syst.*, 24:51–78, January 2006.

[13] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. Witten. The weka data mining software: an update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.

[14] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl. Anatomizing application performance differences on smartphones. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, MobiSys '10, pages 165–178, New York, NY, USA, 2010. ACM.

[15] E. Koukoumidis, D. Lymberopoulos, K. Strauss, J. Liu, and D. Burger. Pocket cloudlets. In *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems*, ASPLOS '11, pages 171–184, New York, NY, USA, 2011. ACM.

[16] E. Lagerspetz, T. Lindholm, and S. Tarkoma. Dessy: Towards flexible mobile desktop search. In *Proceedings of the DIALM-POMC International Workshop on Foundations of Mobile Computing, Portland, Oregon, August 16, 2007, CD-ROM*. ACM, 2007.

[17] N. Lester, A. Moffat, and J. Zobel. Efficient online index construction for text databases. *ACM Trans. Database Syst.*, 33:19:1–19:33, September 2008.

[18] D. Lymberopoulos, O. Riva, K. Strauss, A. Mittal, and A. Ntoulas. Pocketweb: instant web browsing for mobile devices. *SIGARCH Comput. Archit. News*, 40(1):1–12, Mar. 2012.

[19] Microsoft. See your search history: `http://onlinehelp.microsoft.com/en-us/bing/ff808483.aspx`.

[20] P. Shodjai. Your slice of the web: `http://tinyurl.com/24vhmw`, 2007.

[21] Sprint Capping Unlimited 3G Data Service at 5GB. `http://gizmodo.com/391887/oh-no-sprint-capping-unlimited-3g-data-service-at-5gb`, 2008.

[22] J. Teevan, E. Adar, R. Jones, and M. A. S. Potts. Information re-retrieval: repeat queries in yahoo's logs. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 151–158, 2007.

[23] S. K. Tyler and J. Teevan. Large scale query log analysis of re-finding. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 191–200, 2010.

[24] S. K. Tyler, J. Wang, and Y. Zhang. Utilizing re-finding for personalized information retrieval. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 1469–1472, 2010.

[25] Z. Wang, F. X. Lin, L. Zhong, and M. Chishtie. How effective is mobile browser cache? In *Proceedings of the 3rd ACM workshop on Wireless of the students, by the students, for the students*, S3 '11, pages 17–20, New York, NY, USA, 2011. ACM.

[26] I. Witten, A. Moffat, and T. Bell. *Managing gigabytes: compressing and indexing documents and images*. Morgan Kaufmann, 1999.

[27] J. Wortham. Customers Angered as iPhones Overload 3G. `http://www.nytimes.com/2009/09/03/technology/companies/03att.html?_r=2&partner=MYWAY&ei=5065/`, 2009.

[28] Y. Xie and D. R. O'Hallaron. Locality in Search Engine Queries and Its Implications for Caching. In *Proc. IEEE Infocom*, pages 1238–1247, June 2002.