Calendar based Contextual Information as an Internet Content Pre-Caching Tool

Andreas Komninos
University of Strathclyde
26 Richmond Street
Glasgow G11XH, UK
andreas@cis.strath.ac.uk

Mark D. Dunlop
University of Strathclyde
26 Richmond Street
Glasgow G11XH, UK
mdd@cis.strath.ac.uk

ABSTRACT

Motivated by the need to access internet content on mobile devices with expensive or non-existent network access, this paper discusses the possibility for contextual information extracted from electronic calendars to be used as sources for Internet content predictive retrieval (pre-caching). Our results show that calendar based contextual information is useful for this purpose and that calendar based information can produce web queries that are relevant to the users' task supportive information needs.

Categories and Subject Descriptors

H.3.3 [**Information Retrieval**]: Information Search and Retrieval – *Query Formulation, Relevance Feedback, Retrieval models*

General Terms

Algorithms, Performance, Design, Human Factors

Keywords

Pre-caching, Pre-fetching, mobile devices, electronic calendars

1. INTRODUCTION

One of the basic issues in the problem of effectively pre-caching internet content needs on a personal level, is the determination of exactly which documents should be pre-cached. This involves the observation of a user's behaviours, in order to create suitable models that would encompass these behaviours and allow accurate predictions to be made. Such algorithms pre-cache documents in a dynamic manner and presume that the user has a currently active connection, in order to improve navigation (surfing) speed.

In mobile devices such as palmtops, PDAs and smart-phones, traditional pre-caching is a technique that is largely inapplicable (or even unacceptable), due to the significant cost and the charging mechanism for accessing the Internet over a wireless connection. In fact, given the charging per Kilobyte that most service providers adopt for GPRS, it would seem obvious why a user might want to initiate data transfer from the Internet as only

they deem necessary. Furthermore, several PDAs are not equipped with wireless connection and many users have problems configuring data-services when travelling outside their home country. These problems lead to the success of applications such as AvantGo, which provide a means of pre-caching selected or interesting pages ("channels") based on declared user preferences and storing them for off-line browsing.

In our research, we attempt to determine whether it would be possible to employ a user's electronic calendar as a source of information, in order to make informed predictions and provide useful content to the user based on their context. The notion of context in our work encompasses the type of activities a user repeatedly performs, for which of these internet content is required, details on the internet content generally required for each activity type, the user's calendar input style (monolectic or detailed), and finally, the interpretation of the calendar input.

While it may be unrealistic to expect such an approach to cover all of the user's preferences and content needs or desires in a perfect manner, we expect that content that is relevant to the user can be picked up.

2. THE PRE-CACHING SYSTEM

2.1 Overview

In previous work[1], we described the principles for our pre-caching system, whose software comprises of two modules. The first resides on a desktop computer and works in order to extract entries from the user's electronic calendar. It then uses the desktop's connection to pre-cache documents.

The second module is responsible for presenting the pre-cached content on the user's mobile device and therefore resides therein. The second module is also responsible for evaluating the user's interactions with the pre-cached content and passes back this information to the desktop module, which then uses it to make more informed guesses at what it needs to pre-cache, by inferring the user's context.
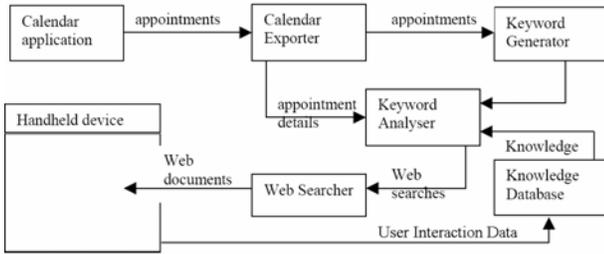
Figure 1. Overview model of the system components

2.2 Component description

Keywords are identified from the entry details and these are combined with the top *n* keywords (in terms of user-perceived importance) from the knowledge base, which are known to be related with the context of the entry keywords.

The knowledge base contains clusters of keywords, which we call category descriptors. Based on the entry keywords, the system will try to identify which category an entry belongs to, and it will associate related terms to it. For example, the entry keyword "Toronto", signifies that the entry probably describes a "travelling" type of event, supposing it is known that the user lives in "Glasgow". Therefore it might be associated with "map" or "hotel" or other relevant terms, picked from the category descriptors. The entry keywords might also be combined with other entry keywords, in order to form web queries. These are then posted to the Google search engine, which in turn provides an index of useful links to be pre-cached. More details of this technique can be found in [2]. Google was chosen as it is widely regarded as the best available Internet search tool at the moment, therefore using it allows the system to display its maximum potential. We did not use other search engines as the comparison of their performance was out of scope with our research.

The web pages with the query results, as generated by Google, are parsed and re-formatted into an XML structure, which holds the document title, the document description and the file system location of the pre-fetched document. The calendar entries, the keywords, the queries for each keyword and the filename for the accompanying Google-based XML are also stored in an XML structure. The structures described above and the pre-fetched documents are then passed on to the mobile device.

In turn, on the device, a log is kept of the user interactions with the content. Based on this information, we attempt to judge the relevance of a given keyword from the knowledge base to the user's context and determine which of these are likely to be wanted as part of a query in the future. We log the viewing of the document index for a given search, the viewing of the summary of a given document, the viewing of the document, the amount of time spent on a document that has been opened and finally any explicit feedback rating that a user might provide for the document. The reading time does not include the time taken to rate the document explicitly. These measures are in order of importance from least to most, and are thus weighed differently with a much greater emphasis on the reading time and explicit rating than the other metrics. We were especially interested in the use of reading time as per Claypool [5], Morita [6] and Kelly

[4],[7], and how it might differ on mobile devices, compared with the aforementioned studies that were solely performed on desktop PCs. Other methods of interaction with documents, such as scrolling and clicking, have been disproved [5], while a debate over reading time is still active.

3. Experiments

3.1.1 Experiment scenario

We gave test subjects the same scenario with details of their persona's living location, job and a list of names of people and how they would be related to them. The subjects were also given five tasks per week for a three week period that forms their hypothetical schedule of activities. These activities were given in the form of a calendar entry that contained a title, location and notes for each one. Some activities did not contain items in the location or notes fields, as they were based on real-world entries that we had collected in previous studies. The subjects were also given clear instructions on the exact meaning of each entry. Finally, they were allowed plenty of time to interact and familiarise themselves with the system fully before the experiment, using sample data. This was done to prevent any possible learning effects from affecting our observations.

The subjects were then allowed to freely navigate through the pre-cached content that was fetched for these hypothetical schedules, and try to locate content that they thought might be helpful to them. The subjects were told explicitly that they were under no obligation to open documents they did not find interesting. We also asked them to give an indication of whether they found the provided content for each activity useful.

Subjects were split into two groups: one group had fixed predictions over the three week period, while the other groups predictions for latter weeks were based on their earlier activity. Logs were kept of their activity and analyzed to provide content that was personalized on the basis of their individual behaviour. The subjects were not aware of this experimental condition, but they were only told after the end of the experiment. Regardless, the distinction is not relevant to the scope of this paper and the observations will be presented jointly for both groups.

3.1.2 Experiment setting

The experiment was performed in a quiet room. This setting might not appear to be realistic in the sense that there were no external distractions for the subjects, although they were given food and drink and were allowed to communicate and interact with each other. Mobile devices are used in both mobile and stationary environments and since the experiment shows that reading times are not long (around 25 seconds), it is expected that a mobile user could easily dedicate such small times to interact undisturbed. We perceive the notion of "mobile" to mean "out of office" rather than "walking" or "driving", therefore the setting seems adequate for the purposes of the experiment. In any case, a maximum time limit of 2 minutes, based on the observations from an initial test group (see 3.1.3), is imposed on the measuring to eliminate gross inaccuracies due to user distraction.

3.1.3 Initial experiment setup

An initial group of ten subjects volunteered to test the system before we proceeded with the actual experiment. All of them were from a similar background and considered themselves computer literate, although most did not have previous experience with a

PDA. The initial group was given similar data to those that would participate in the main experiment, with data only covering one week. The main goal of this initial experiment was to observe the average reading times for the web documents and their relation to explicit feedback, as we planned to use this metric for implicit relevance feedback. In addition the study acted as a pre-trial to ensure the system ran smoothly with users that were unfamiliar with it.

3.1.4 Cache hit-rate performance comparison
For our main experiment, we used two groups of ten people each, which did not include any of the members of the initial subject group. Unfortunately, due to data corruption on the logs of two members of one group, we were forced to exclude them from the analysis, removing also two random members from the other group, to make the figures directly comparable.

The analysis of the main experiment's results shows that from the documents cached through the queries that were automatically formed by the system, a significant percentage was deemed worth opening by the users (table 3). These hit-rate percentages show that the information found in calendars can be exploited to generate queries which, combined with a good web search engine (such as Google), can retrieve a significant number of documents that the user might find interesting.

The following table shows the numerical and percentile quantities of the opened documents (cache hits) vs. the total documents offered. It is apparent that some growth can be observed in the cache hit-rates, especially for week 3. This can be attributed to two factors: Firstly, because one of the groups was actively monitored and provided with documents from queries that were adjusted to suit their preferences. The number of queries this group was provided with was gradually smaller, as an attempt was made to provide to them less, but more relevant queries. Secondly, the variance in the contents of entries for each of the three weeks is also the deciding factor with regard to the amount of queries that could be generated, and consequently for the amount of documents that were fetched (compare for example a query for hotels in Barcelona, a large well-known city, and Paisley, a relatively small town in Scotland). Statistical tests performed on the data acquired from the two groups indicate the group whose interaction was monitored benefited consistently from improvements in their cache hit rate over the other group (P<0.02).

Table 1. Total documents vs. Opened documents

	Total	Opened	Percentage
Week 1	2928	484	16.5%
Week 2	3759	667	17.7%
Week 3	2313	645	27.8%

3.1.5 Viewed web document ratings as a query relevance indicator
We also examined the average document scores for each session (table 2). It would appear from this table that the scores remain at a constant level and in fact, close to the middle of the scoring table (M=2.46). While one might have expected a larger variance in the scoring results, due to the changing amount of information in the calendar entries, this result is consistent with the results in [3]. In this research, it is mentioned that users should expect approximately 1 of 2 web documents they view to be highly relevant, therefore our observations seem to confirm that this theory holds true.

Table 2. Average Document Ratings on a 1-5 scale (1-irrelevant, 5-most relevant)

	Group 1	Group 2	Joint
Week 1	2.26	2.37	2.44
Week 2	2.60	2.38	2.48
Week 3	2.49	2.36	2.42

We chose not to measure the individual scores attained by each group in order to establish a trend. These would be actually just measuring the ability of Google to return relevant results, where as we are interested in measuring the relevance of the web query in the context of the calendar entry and the user's needs.

4. Discussion and future work
In this paper, we described a pre-caching system which is dependent on the inference of user context based on electronic calendars, in order to provide useful content for a user with a small mobile computing device. While such a system in its own right would not be able to completely satisfy all of a user's internet content needs or desires, we show that it can indeed provide useful content for the appropriate entry categories, through the inference of user context and the observation of their interaction. Even in the case of entries where the calendar entry comprises of a single word, the automatic generation of web queries based on common knowledge and the users' observed preferences proves to be able to provide meaningful and useful content.

It could be argued that the cache hit rates observed in our experiment are existent simply because the content was already available to the subjects. We believe however that even if this is true, the consideration of the actual percentages shows that, at the very least, the use of calendar extracted contextual information can provide a useful supportive service to users, especially those of mobile devices. The significant cache hit-rates reported by our research provide motivation to investigate the combinatory performance that the calendar-based algorithm we described might have, together with other proven pre-caching techniques.

5. References
[1] Komninos, A., Dunlop, M.D. (2003): Towards a model for an Internet content pre-caching agent for small computing devices, Proceedings of the 10th International conference on Human Computer Interaction (HCII2003), Crete, Greece, 2003

[2] Komninos, A., Dunlop, M.D. (2004): Towards a model for an Internet content pre-caching agent for small computing devices, Proceedings of the Mobile and Ubiquitous Information Access Workshop in association with Mobile HCI04, Glasgow 2004

[3] Jansen, B. J., Spink, A., (2003) An analysis of web documents retrieved and viewed, Internet Computing Conference, Las Vegas

[4] Kelly, D., Belkin, N. J. (2001). Reading time, scrolling and interaction: Exploring implicit sources of user preference for relevance feedback. In Proceedings of the 24th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR '01), New Orleans, LA., 408-409.

[5] Claypool, M., Le, P., Waseda, M., Brown, D., (2001), Implicit Interest Indicators, In Proceedings of ACM Intelligent User Interfaces Conference (IUI), Santa Fe, New Mexico, January 14-17.

[6] Morita, M., Shinoda, Y., Information Filtering Based on User Behaviour Analysis and Best Match Text Retrieval, Proceedings of SIGIR on Research and Development, pp.272-281, 1994

[7] Kelly, J. D., Understanding Implicit Feedback and Document Preference: A naturalistic user study, PhD Thesis, State University of New Jersey, USA, 2004