

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
1 April 2010 (01.04.2010)

(10) International Publication Number
WO 2010/034056 A1

(51) International Patent Classification:
G06Q 30/00 (2006.01) **G09F 23/00** (2006.01)
G06Q 90/00 (2006.01)

(21) International Application Number:
PCT/AU2009/001261

(22) International Filing Date:
23 September 2009 (23.09.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
2008905019 26 September 2008 (26.09.2008) AU

(71) Applicant (for all designated States except US): **NATIONAL ICT AUSTRALIA LIMITED** [AU/AU]; ACN 102 206 173, Level 5, 13 Garden Street, Eveleigh, NSW 2015 (AU).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **PENEV, Alexander** [AU/AU]; 1/22 Villiers Street, Kensington, NSW 2033 (AU). **WONG, Kwok, Kay** [AU/AU]; Level 5, 13 Garden Street, Eveleigh, NSW 2015 (AU).

(74) Agent: **F B RICE & CO.**; Level 23, 44 Market Street, Sydney, NSW 2000 (AU).

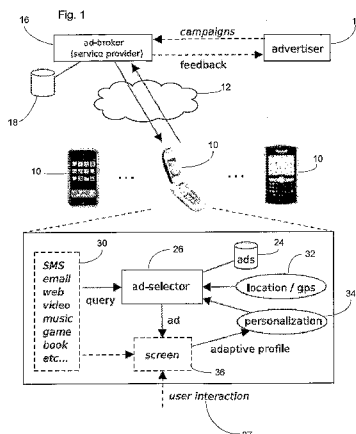
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: ADVERTISEMENT SELECTION FOR A PORTABLE COMPUTING DEVICE



(57) Abstract: The invention concerns the selection of relevant advertisements for presentation on a portable computing device. The device has stored in memory multiple advertisements where each advertisement is associated with one or more advertisement descriptors. Based on the dynamic and heterogeneous content presented (98) on the device (10), content descriptors are identified (100). Next, content descriptors are compared (102) to advertisement descriptors stored in memory to select and advertisement and present (104) it on the device (10). This invention identifies that many portable devices are continually used without being on-line (i.e. being offline). It is an advantage of the invention that relevant advertisements can be selected and displayed on a portable device while the device is offline and no reliance is placed on an external network or external device as the device can decide what advertisements to display and when. Aspects of the invention include a method performed by the device, the device and software for installation on the device.

WO 2010/034056 A1

ADVERTISEMENT SELECTION FOR A PORTABLE COMPUTING DEVICE

Technical Field

The invention concerns the selection of relevant advertisements for presentation on a portable computing device. Aspects of the invention include a method performed by the device, the device and software for installation on the device.

Background Art

Advertising on portable computing devices, such as PDAs, mobile phones and iPods, is achieved one of two ways:

1. Advertisements are sent directly to the portable device, usually based on some prior knowledge of the general demographic of the user. The timing and relevance of the advertisements is poor. Additional cost is incurred in the bandwidth for the transmission of the advertisement. These are commonly viewed by portable device users as spam.

2. The advertiser or content provider sends messages directly to the portable device. The message includes an advertisement that is inserted into the content that is downloaded by the portable device when online. These include Short Message Service (SMS) injections, advertisements on websites, and video prepended advertisements. SMS injections are where an SMS is intercepted by an advertiser, who then places an advertisement inside it and forwards it to the recipient. Web advertisements are advertisements placed on web sites that are displayed on the portable device. They cannot be tailored to be relevant to a portable device rather than a desktop computer. Again, the portable device must be online and the additional data costs are incurred for the transmission of the advertisement. Prepend video advertisements are short video clips that inserted into a downloaded video. For these types of advertising the portable device must be both online and currently downloading content. Additional cost is incurred in the bandwidth for the transmission of the advertisement.

Summary of the Invention

In a first aspect the invention provides a method performed by a portable computing device for selecting an advertisement for presentation on the device, the device having stored in memory multiple advertisements where each advertisement is associated with one or more advertisement descriptors, the method comprising:

- (a) presenting content on a presentation means of the device;
- (b) identifying content descriptors that describe the content presented on the device;
- (c) comparing content descriptors to advertisement descriptors stored in memory; and
- (d) based on the comparison, selecting the advertisement and causing it to be presented on the device.

This invention identifies that many portable devices are continually used without being online (i.e. being offline). It is an advantage of the invention that relevant advertisements can be selected and displayed on a portable device while the device is offline as the advertisements are stored locally on the device. The selection process is performed by the device itself meaning that no reliance is placed on an external network or external device as the device can decide what advertisements to display and when. No active open channel to a content provider through which content or advertisements can be received is required. It is a further advantage that the advertisements are relevant to what is actually being displayed on the device, giving the presented advertisement the benefits of being targeted.

The method may comprise repeating steps (a) to (d) when user input is received to change the content presented, and the content presented is changed.

The method may comprise repeating steps (a) to (d) at intervals.

The content may include text and the advertisement descriptors are based on at least the text.

The content may be stored in memory.

The device may be offline when the method is performed. That means the device is not online, having no active open channel with a content provider.

The presentation may be by way of display on a screen of the device and/or sound produced by a speaker.

Identifying content descriptors may comprise generating the content descriptors.

Alternatively, identifying content descriptors may comprise accessing content descriptors stored associated with the content as stored in memory.

The content descriptors may be keywords, classifications and/or descriptors to guide or bias the selection.

Comparing in step (c) may comprise using the content descriptors as a search query over the advertisement descriptors.

One or more advertisements may be associated with a geographic descriptor, and selecting an advertisement in step (d) may further be based on a comparison of the current location of the user and the geographic descriptor.

In a second aspect the invention provides a portable computing device having:

memory to store a set of advertisements, each advertisement associated with one or more advertisement descriptors;

presentation means to present content and a selected advertisement;

a processor to cause content to be presented on the presentation means, to identify content descriptors that describe the content presented, and to select the selected advertisement from memory based on a comparison of advertisement descriptors and content descriptors to cause the selected advertisement to be presented on the presentation means.

In yet another aspect the invention provides software, that is computer readable instructions stored on a computer readable medium, that when executed by a computer causes it to perform the method described above.

Brief Description of the Drawings

An example of the invention will now be described with reference to the accompanying drawings, in which:

Fig. 1 is an abstraction of the portable device and its connection to a mobile network;

Fig. 2 is a flow chart showing an example of the invention;

Fig. 3 is a schematic representation of the structure of a record in the advertisement database; and

Fig. 4 is pseudo code showing the method of selecting an advertisement for display.

Best Modes of the Invention

Portable computing devices by their nature are small in size to make them convenient to carry by the user. Common examples are iPods, portable game consoles, Wi-Fi Netbooks, MIDs, tablet PCs and e-readers. Portable computing devices have the following hardware features:

Input means, such as a keypad or touch screen using which a user is able to provide input so as to control the device, including requesting changes in what software application to run and what content to present on the device.

Presentation means, such as a screen that can display a user interface of a software application that in turn displays content. Often, a speaker is also used to present sound to the user.

Input port/output port, such as a connection to a mobile phone communications network, internet, GPRS, and/or GPS. Data can be sent and received by the device using this port.

Memory, such as internal memory and/or a removable memory disk that is able to store a database of advertisements. Further memory is used to store content that can be presented to the user. The content can be personal content, such as emails, SMS messages, diary entries and personal notes. The content also includes impersonal content such as music, movie clips, games, e-books, and pocket encyclopaedias. Executable software is also stored on the device, including operating software and application software, for example webpage browser and email, SMS and instant messaging applications.

A processor to execute the software. This includes processing the user's input and to execute the method of selecting an advertisement as described below. Usually these devices are not as computationally powerful as laptops or desktops.

The device can be online, that is, having an open communication channel on the input/out port through which data can be transmitted especially content. Each device

can be operated while not online (i.e. offline). In fact, it is common that these devices will be used offline mostly. That means that the content being displayed on the device is not displayed in real time as it is received from a third party at the input port, instead it is stored in memory in advance.

It should be understood that all the processes/functions/steps performed by the device is controlled by the processor executing stored software.

In Fig. 1, multiple portable devices 10 are shown, that are each connected to a mobile phone network 12 using which calls can be made and received by the device 10. For simplicity, the following description will refer to only one mobile phone 10.

An advertiser 14 is also provided who owns the advertisements and wishes to distribute them to mobile phone 10 users. Advertisements can be either graphic files, video files or sound files. These are usually bundled into campaigns and provided to an ad-broker 16. The ad-broker 16 receives campaigns from multiple advertisers 14 (not shown) and stores them in a data store, such as a database 18. The ad-broker 16 operates a server (not shown) that is connected to the mobile phone network 12 and is able to transmit a pre-selection of advertisements to the mobile phone 10. Each advertisement is associated with an advertisement descriptor (also referred to here as advertisement tag). In this example, a total of 10,000 separate advertisements are transmitted to the mobile phone 10. The advertisements are for products, such as e-books and movies. An advertisement for the movie 'Harry Potter and the Goblet of Fire' may have the advertisement descriptors <harry potter>, <movie>, <fantasy> and <entertain>.

In alternative embodiments the advertisements could be placed on the phone 10, or updated, using different methods. For example, it could be pre-stored on the phone 10 when the phone 10 is purchased by the user or it can be connected to a personal computer that is connected to the internet which downloads the advertisements from the ad-brokers internet server (not shown). In turn the mobile phone 10 downloads the advertisements from the personal computer using a direct connection.

In this example a set of advertisements are received by the device at the input/output port and are saved to an advertisement database 24 in the internal memory of the mobile device 10. The database 24 stores each advertisement and its descriptors in an associated manner, similarly to a relational database.

Currently the user is listening to a song and a suitable interface of a media player is displayed on the screen 36 which includes the name of the song. While this interface is displayed an advertisement may have been previously selected and displayed to the user using the method described below where the displayed name of the song was taken as the content.

Referring now to Fig. 2, the then user provides input 37 to the device 10 that causes the display 36 to change. For example, the user can provide click, character or scrolling input that is typically used when operating a mobile phone 10. In this example, the user provides input to the device to cause an emailing interface to be currently displayed, replacing at least in the most part the media player. In this example, the user then starts to enter text of the email being written by the user to suggest to a friend that they go to watch a movie. The text also appears in the current display of the email interface 98.

The content currently displayed, being the text of the email itself, is analysed as described in further detail below to identify 100 content descriptors. In this case the content descriptors <movie> and <social> are generated.

The processor then performs a comparison 102 of the identified content descriptors <movie> and <social> with the advertisement descriptors stored in memory 24, and identifies a match with advertisement descriptor <movie> for an advertisement for the movie Lord of the Rings. The location 32 of the user and the personal preferences 34 of the user can also be taken into account in this analysis. These are external factors that can influence the advertisement selector's 26 selection. Assuming (for simplicity here) that this is the only match of descriptors, then the processor selects the Lord of the Rings advertisement and causes it to be displayed 104 on the screen. This is performed in real time so that the advertisement is displayed while the email is still being composed by the user.

In this example, the advertisement may be displayed as a film banner for a predetermined period of time, such as 5 seconds. If the user clicks to view the advertisement, then the advertisement expands to fill substantially the most of the screen and includes additional information. If the user doesn't click on the banner, the banner is removed from the display.

It should be appreciated that the steps 100, 102 and 104 are repeated multiple times while the device is offline. In this example an advertisement is selected and displayed each time the content display is substantially updated, but no more often than at a predetermined regularity. For example, the user may be reading a long memo on their device. At the top of the memo there is a discussion about technology but as the user scrolls down the discussion changes to politics. Since the descriptors in this example are based on the content currently displayed, the content descriptors determined for the top and bottom of the memo may be different and accordingly the advertisements selected and presented may be different.

Alternatively, the method of selecting a new advertisement may be repeated on the identification of certain events, such as each time the user switches between different software applications that each present a different user interface on the screen.

Referring back to Fig. 1, usage information can be periodically sent from the mobile phone 10 to the advertisement broker 16 using the cell network 12. The feedback may be primarily usage statistics. Based on the received usage statistics, the advertisement broker 16 is able to bill the advertiser 14 based on the amount of times that their ads were displayed the portable devices 10 and/or the user interacted with their displayed advertisements.

The method of selecting the advertisements for presentation will now be described in more detail.

The advertisement database 24 stores a set of advertisements on the device 10. In the database each advertisement is associated with advertisement descriptors in a manner that the processor is able to dynamically filter through them. The advertisements are also stored in a way that the advertisements on the database can be updated by replacing some or all advertisements. In this example there are over 10,000 ads in the database that are indexed in memory using mContext [1].

Fig. 3 shows the record structure for each advertisement stored in memory of this example. There is an identifier 50, a title 52, a description 54, the advertisement itself 56 such as the graphic, some display configuration parameters 62, a location component (optional) 58 and set of advertisement descriptors 60.

Advertisement descriptors are similar to keywords. These tags are used for indexing, retrieval, scoring and personalization tasks in information retrieval. Being able to use tags to perform all these function makes the use of tags suitable on small devices with limited processing power. Unlike keywords, the tags do not need to be regular words, nor do they need to carry a legible meaning. The only requirement is the ad-selector 26 can match one string of characters of a tag associated with an advertisement to another string of characters that is determined as the content descriptors of the displayed content. In this example three types of tags are used as follows:

1. Regular keywords that describe the advertisement. For example, an advertisement for an iPhone may have tags such as <iphone>, <apple>, <mobile>, <phone> and <smart phone>.
2. Syntax strings that represent a classification. For example the same iPhone advertisement has additional classification tags of <category:electronics> or <:576>, where 576 symbolizes a category for electronic products. As can be seen from this example classification tags need not be real words but have a valid meaning within the system.
3. Syntax words designed to guide or bias the selection towards certain advertisement. For example a record company 'Energized' offering music downloads tags their advertisements and their music with a broker-approved tag <:78623> such that both the advertisement and the content share this tag and are likely to be matched together. The goal of this would be to show Energized's advertisements whenever Energized's music is played.

As shown at 60, each tag also has one or more numerical weights attached to it and stored in memory. These weights will represent the strength of the tag. For example, the tag <iphone> for the iPhone advertisement may have a strength of 1.0 because it's very relevant, whereas the tag <iphone> for an advertisement of a shop that sells electronics may only have a strength of 0.1. In this example, the weights of the tags are used for determining relevance and also for billing. The billing weight will represent how much the advertiser is willing to pay for the advertisement to be shown. If more than one weight is associate with a tag, in this example they are expressed as a single abstract weight such as by using their harmonic mean.

The ad-selector 26 is a software agent installed on the device 10. The processor executes the software to select and display advertisements. After display, the ad-selector 26 will also record in memory if the user clicked on the advertisement or not, and can update its internal book-keeping of these details so that it can personalize 34 future selections to better suit the user's taste.

Due to the randomness 102g of the selection process, the ad-selector 26 chooses a suitable, but not necessarily the most suitable, advertisement based on the currently viewed content.

The method initially comprises presenting content currently on the device 98. In this example the content identified is the text features of the currently displayed content, including any displayed file names. In some applications only the text actually displayed is used in step 100, and in other alternatives content analysed extends to content not actually displayed but associated with the displayed content, such as a the name of the file whose content is displayed but not the name of the file itself or the rest of a paragraph that is in part displayed.

Then, content descriptors for that content is identified 102. First, the content currently displayed is analysed to identify a succinct tag-like representation. If the content displayed is already stored and indexed in the device's memory using mContext then the content descriptors are previously determined when the content was stored and indexed. The ad-selector 26 can simply access 100aa these content descriptors from memory rather than re-analyse the content.

Otherwise, they are generated as needed by transforming 100a the text identified as currently displayed on the device. The transformation step uses techniques that are already known to those skilled in the art: case-folding, stop-word removal, stemming, extracting salient words, and finally TFIDF weighting. The goal of the transformation step is to convert arbitrary content text into a weight tag-form representation. This representation is superior to the original text because it contains only the most important words and therefore more efficient for matching purposes. It also gives weights to each content descriptor making each content descriptor more expressive. It also adds extra tags such as classification tags and match-guiding tags, so some new information is added.

After transforming 100a the input displayed content, the top-most tags are query expanded 100b to obtain a richer content-query. The new content descriptors are now considered a search query. This content-query may be longer than the input if the input is short (e.g. a few words). For example, if the content is the SMS message “want to see a movie tonight?”, the identified tags could be <movie> and <tonight>, with a suitable categorical tag to suggest that the input is related to television and cinema. The query expansion step 100b would add some related words, including <film>, <television>, <actor>, <night> and <evening>.

Combined the tags, each having an associated weight, are used as a query 102 by the advertisement selector 26 on the advertisement database 24. To perform the this analysis on a small device several files are stored on the device, such as a stoplist, an IDF table, a word relationship table, a stemmer and classifier as known in the art

Next, these contents descriptors are compared 102 with the advertisement descriptors 60 in the advertisement database 24 in order to select an advertisement for display. This includes the retrieval 102a of candidate advertisements from the database. Candidates advertisements are those that match at least one content descriptor. The content-query formulated above can match many advertisements, in this example hundreds.

Then, pre-filtering 102b to discard weak candidates based on approximating their relevance score is performed. Many of these matches will be quite weak. For performance reasons the weak advertisements are now removed so that only the most promising candidate advertisements progress to the scoring step 102c. Differentiating between the promising and the weak advertisements is done using an on-the-fly approximation of their relevance score. More specifically, weak candidates are discarded based on their approximated relevance score. To approximate their score, in this example the advertisement selector 26 sums the pair-wise products of their IDF values and their tag strengths in the content-query and the advertisement. The higher the score, the more relevant the advertisement is anticipated to be. Alternative methods can be used provided they can discard weak candidates using a very quick approximate of their relevance without doing a full relevance calculation.

After this step is completed only the top- k ads are kept, where the variable k may be a function of the database size. These remaining k advertisement candidates are scored 102c in a more in-depth manner. The previous step allows a known quota of the most-promising candidate advertisements to reach the scoring step 102c. Each advertisement A is now scored more thoroughly for the content-query Q in a manner that uses A 's data, A 's location and personalization factors.

$$S_{A,Q} = base(A) + rel_{A,Q}(1 + \alpha g(A))(1 + \beta p(A))$$

where $base(\cdot)$ is a base rating reflecting A 's all-suitability, rel is a relevance score against the currently viewed content 102d, $g(\cdot)$ is a geographic function to measure the appropriateness of A 's location 102e and $p(\cdot)$ is a personalization factor 102f.

The candidate advertisements will receive low scores if their rel is low or if they specify inappropriate locations. Their scores will be high if their rel is high, and even higher if it they specify an appropriate location or are aligned to the user's personal interests. The parameters α and β are tuneable.

The base score is useful for ads that always have a positive relevance and qualify for any query. This may include a service provider bulletins and the temporary boosting of ads that the user has historically shown strong interest in but has not yet seen since their most recent update (see personalisation 102f below).

The relevance between A 's tag and the content query is a greedy sum:

$$rel_{A,Q} = \sum_{t \in (T_A \cap T_Q)} w(t, T) \cdot w(t, T_Q) \cdot IDF_t \cdot titlematch(A, T_Q)$$

Where $w(t, \cdot)$ is a weight in $[0,1]$, IDF measures tag rarity and T_{object} is the set weighted tags for *object*. Generally, the more tags the advertisement and the content-query have in common, the higher the rel score will be.

Referring back to Fig. 1, location 102e sensitivity 32 refers to knowing the location of the device 10 to bias in favour of advertisements that are appropriate/close and against advertisements that are inappropriate/far. The device 10 can determine its location via GPS and triangulation. In alternate embodiments other small devices, the ad-selector

26 can prompt the user to place a marker on a map or to ask them for their zip/area code. This feature is optional since some advertisements are not location specific, such as an advertisement for Coca-Cola. On the other hand, businesses such as restaurants and shops have fixed addresses and can benefit from specifying a location. The geographic function $g(\cdot)$ will act as either a boost or penalty. It will boost advertisements with an appropriate location and penalize ads with an inappropriate location.

The location component 58 in this example provide a latitude, longitude and radius triplet that describes the area in which the advertisement is considered relevant. The radius may be as small as a few blocks or as large as a country.

Personalization 102f, 34 will monitor the user's prior interactions with previously-shown ads to adaptively learn what topics and what sort of advertisements they are interested in (e.g. sport topics, video clip advertisements). People naturally have both persistent and transient interests, which we refer to as long-term and short-term. The ad-selector 26 will model these using tags and build a user profile 34. We model long-term interest as repeated confirmation of short-term interests and we model short-term interests as viewing ads. For example, after viewing the Harry Potter advertisement, the user's short-term interest may be <movie>, <film>, <harry potter> and <Radcliffe>, while their long-term interest may be different, such as <technology>, <internet>, <movie>, <software>. These personalised tags 34 can be used to influence future selections. For example, if the user clicks an advertisement with tags <iphone>, <apple>, and <electronics> then the ad-selector 26 will make note of this and in future queries can slightly boost advertisements that share the same tags—such as other Apple products. This is a passive observation and has several advantages over the direct questionnaire to the user approach. It requires no user effort, it is self-updating, it can glean fine-grained interest instead of only coast-grain interest, and it can track short-term interest. It also has advantages over the advertiser tracking approach in that it can preserve user privacy because the profile stays on the device.

The ad-selector 26 will assign the candidate advertisements a probability of selection that is proportionate to their final relevance score 102d.

Post filtering 102g is used again to discard weak candidates. These advertisements may be weak because they scored poorly for their location or simply failed to live up to their expectations regarding their initial approximated score in the pre-filter step 102b.

This time the candidates will be filtered using a score-based filter, where candidates scoring below a certain value are dropped. This value is variable and depend on the final set of candidate advertisements. After removing the weak advertisements, the final small number of advertisements will be the very highest scoring candidates.

Finally, a lottery 102g is performed, where the final set of candidates enter a draw to randomize the choice. Aside from location and personalization factors, the selection process so far would produce the same result for the same content. If the device were to always choose the top candidate then there would be little variety, especially if the same content is seen again moments later. Thus the ad-selector 26 will randomize the final choice to provide some variety. Due to the lottery 102g, even with the usage data the advertisement broker 16 can't reverse engineer the content meaning privacy is maintained.

The algorithm performed by the ad-selector is shown in Fig. 4.

It should be understood that the techniques of the present invention might be implemented using a variety of technologies. For example, the methods described herein may be implemented by a series of computer executable instructions residing on a suitable computer readable medium. Suitable computer readable media may include volatile (e.g. RAM) and/or non-volatile (e.g. ROM, disk) memory, carrier waves and transmission media (e.g. copper wire, coaxial cable, fibre optic media). Exemplary carrier waves may take the form of electrical, electromagnetic or optical signals conveying digital data steams along a local network or a publically accessible network such as the internet.

It should also be understood that, unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that processes and transforms data represented as physical (electronic) quantities within the computer system's registers and memories

into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

Advantages of this example of the invention includes:

- Advertisements can be selected and displayed to the user regardless of the software program that is currently being used by the user to display content. For example, it is not tied to just webpage content and it can also be used when the user is sending an SMS or listening to music. The content is therefore dynamic and heterogeneous.
- This example addresses the problem that it is not suitable or practical for advertisers to send advertisements targeted to offline content displayed on a portable device. The advertiser does not know (and for privacy reasons shouldn't know) what is stored and displayed offline on the portable device. All the advertisement processing is entirely client-side.
- All advertisements can be saved locally so the need for external intelligence and transmission to select advertisements is removed while also keeping the user's content and data private.

It will be appreciated by persons skilled in the art that numerous variations and/or modifications may be made to the invention as shown in the specific embodiments without departing from the scope of the invention as broadly described.

For example, the ad-broker may in fact be the cell network service provider to the mobile phone 10.

The method of selecting and displaying advertisements can also be performed when the device is displaying content that is being received in real time when the device is online.

Updating the set of advertisements stored in memory can be done using any connection available to the device, with many devices having multiple input ports, such as to the mobile network, mini USB and blue tooth.

The present embodiments are, therefore, to be considered in all respects as illustrative and not restrictive.

References

- [1] www.mcontext.com

CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method performed by a portable computing device for selecting an advertisement for presentation on the device, the device having stored in memory multiple advertisements where each advertisement is associated with one or more advertisement descriptors, the method comprising:
 - (a) presenting content on a presentation means of the device;
 - (b) identifying content descriptors that describe the content presented on the device;
 - (c) comparing content descriptors to advertisement descriptors stored in memory; and
 - (d) based on the comparison, selecting the advertisement and causing it to be presented on the device.
2. The method according to claim 1, wherein method comprises repeating steps (a) to (d) when user input is received to change the content presented, and the content presented is changed.
3. The method according to claim 1 or 2, wherein the method comprises repeating steps (a) to (d) at intervals.
4. The method according to claim 1, 2 or 3, wherein the content includes text and the advertisement descriptors are based on at least the text.
5. The method according to any one of the preceding claims, wherein the content is stored in memory.
6. The method according to any one of the preceding claims, wherein the device is offline when the method is performed.
7. The method according to any one of the preceding claims, wherein the presentation is by way of display on a screen of the device and/or sound produced by a speaker.
8. The method according to any one of the preceding claims, wherein identifying content descriptors comprises generating the content descriptors.

9. The method according to any one of the preceding claims, wherein identifying content descriptors comprises accessing content descriptors stored associated with the content as stored in memory.
10. The method according to any one of the preceding claims, wherein the content descriptors are keywords, classifiers and/or descriptors to bias the selection.
11. The method according to any one of the preceding claims, wherein comparing in step (c) comprises using the content descriptors as a search query over the advertisement descriptors.
12. The method according to any one of the preceding claims, wherein one or more advertisements are associated with a geographic descriptor, and selecting an advertisement in step (d) further based on a comparison of the current location of the user and the geographic descriptor.
13. A portable computing device having:
 - memory to store a set of advertisements, each advertisement associated with one or more advertisement descriptors;
 - presentation means to present content and a selected advertisement;
 - a processor to cause content to be presented on the presentation means, to identify content descriptors that describe the content presented, and to select the selected advertisement from memory based on a comparison of advertisement descriptors and content descriptors to cause the selected advertisement to be presented on the presentation means.
14. Software, that is computer readable instructions stored on a computer readable medium, that when executed by a computer causes it to perform the method according to any one of claims 1 to 12.

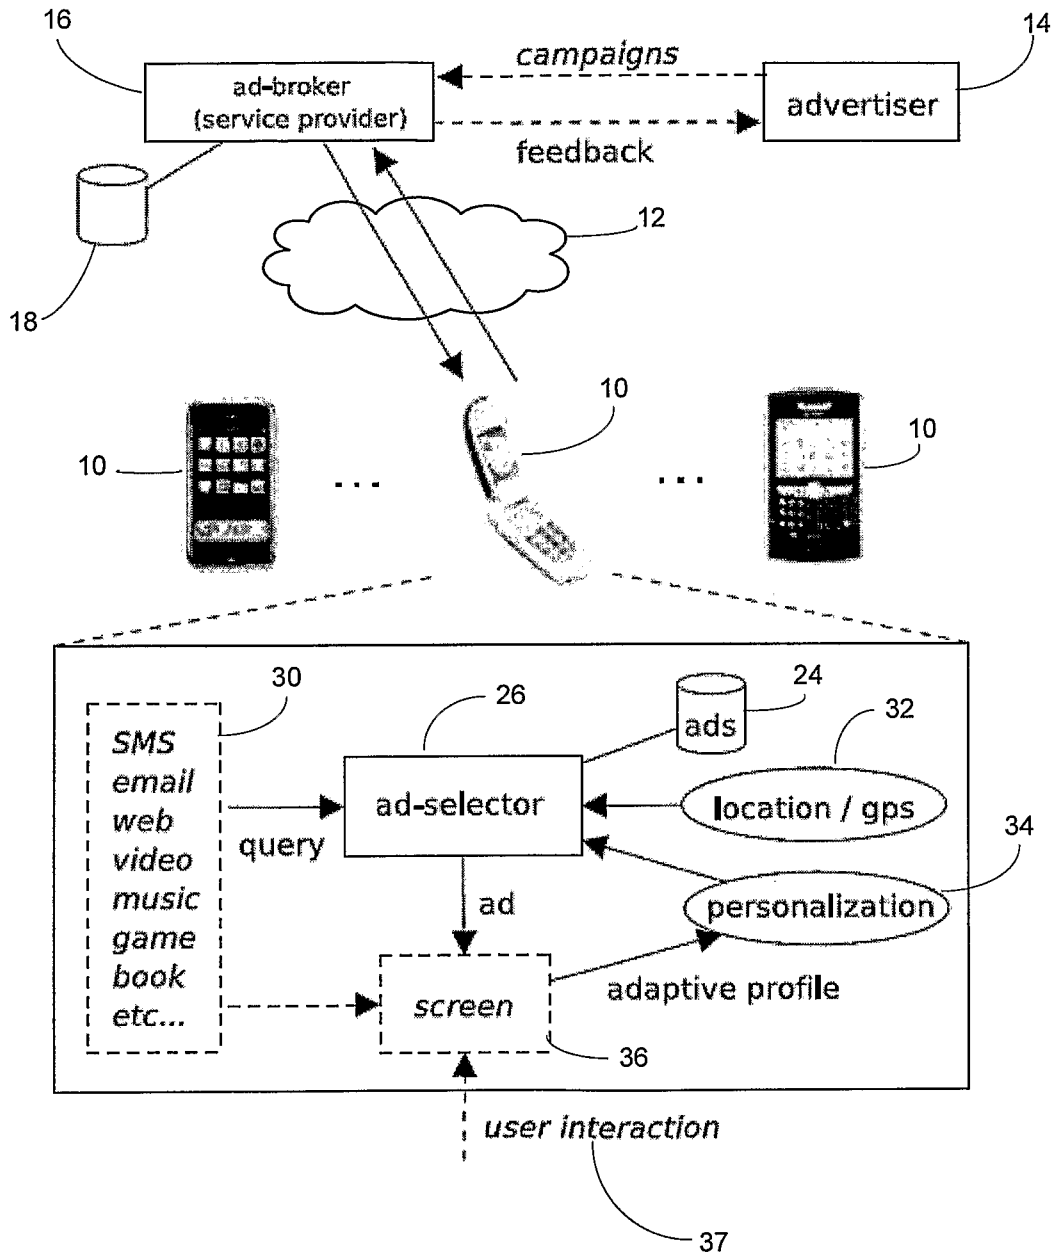


Fig. 1

2 / 4




Fig. 2

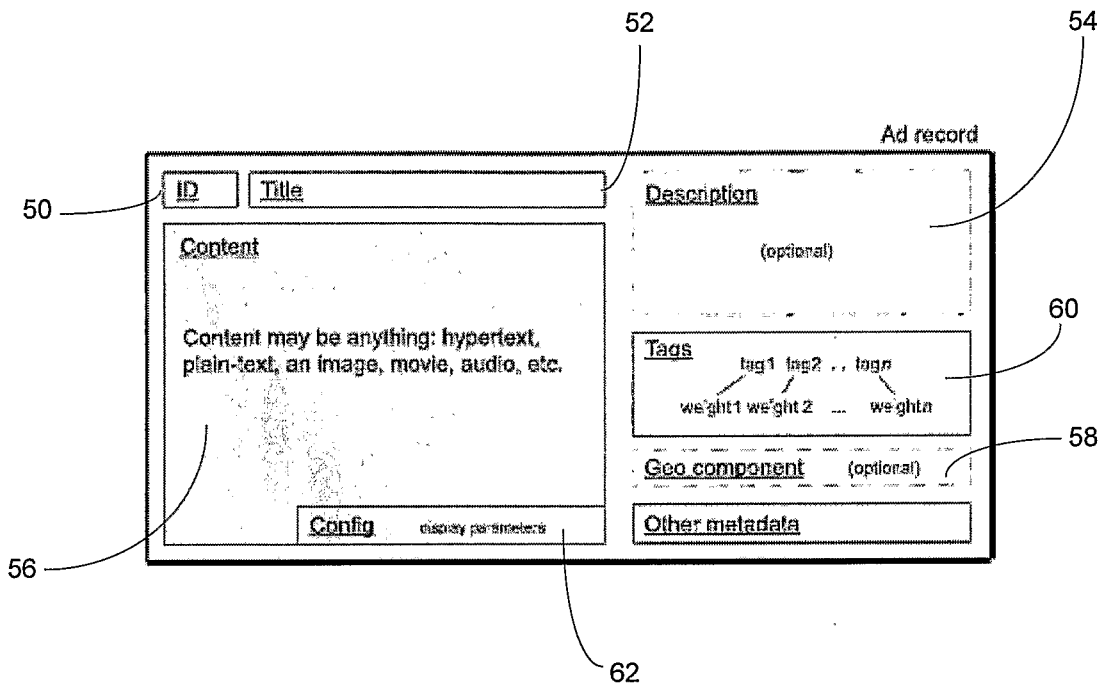


Fig. 3

4 / 4

```

ad-selector : choose best ad for query tags  $Q$ .
    //query expand Q to get a content-query
1.  $E := T_Q \cup \bigcup_{t \in T_Q} \text{queryexpand}(t)$ 
    //retrieve ads matching at least one tag
2.  $C := \{cand \in \text{ad-db} \mid T_{cand} \cap E \neq \emptyset\}$ 
3. apply prefilter to discard weak candidates
4. add ads with positive base rating
    //score candidates using more features
5. for each  $c \in C$ :
6.    $rel_{c,Q} := \sum_{t \in T_c \cap E} w(t, T_c) \cdot w(t, E) \cdot IDF_t \cdot \text{titlematch}(A, T_Q)$ 
7.    $s_{c,Q} := \text{base}(c) + rel_{c,Q}(1 + \alpha g(c))(1 + \beta p(c))$ 
8. end for
9. apply postfilter to discard weak candidates
10. set probabilities using  $L_1$ -norm
11. randomize the ad choice

```

Fig. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/AU2009/001261

A. CLASSIFICATION OF SUBJECT MATTER		
Int. Cl.		
G06Q 30/00 (2006.01), G06Q 90/00 (2006.01), G09F 23/00 (2006.01)		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols)		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
WPI and EPODOC: IPC G06Q, G09F, and keywords (portable device, mobile, advertisement) and like terms		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2008/0215429 A1 (RAMER ET AL.) 4 September 2008 Abstract, paras [0005, 0009, 0039, 0062, 0069, 0070-0071, 0074-0078, 0085, 0090, 0113, 0137, 0174, 0185, 0941] and figs 1, 3-5, 9-11, 19-21.	1-14
X	US 2008/0005695 A1 (OZZIE ET AL.) 3 January 2008 Abstract, paras [0012, 0033-0034, 0038, 0043-0045, 0062-0063, 0095, 0103-0105] and figs 1-4, 7-8, 15-17.	1-14
X	US 2007/0088852 A1 (LEVKOVITZ) 19 April 2007 Abstract, paras [0028, 0063, 0095, 0098, 0151, 0053, 0062, 0142, 0148, 0312] and figs 1-2, 7, 11-14.	1-14
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C <input checked="" type="checkbox"/> See patent family annex		
* "A"	Special categories of cited documents: document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed	
Date of the actual completion of the international search 19 October 2009		Date of mailing of the international search report 10 NOV 2009
Name and mailing address of the ISA/AU AUSTRALIAN PATENT OFFICE PO BOX 200, WODEN ACT 2606, AUSTRALIA E-mail address: pct@ipaaustralia.gov.au Facsimile No. +61 2 6283 7999		Authorized officer Rajni Verma AUSTRALIAN PATENT OFFICE (ISO 9001 Quality Certified Service) Telephone No : +61 2 6283 2816

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2009/001261

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P, X	WO 2009/099919 (Horizon capital securities limited) publication date 13 august 2009 having priority date of 31 January 2008. Entire document	1-14
A	US 6892217 B1 (Hanmann et al.) 10 May 2005 Entire document	
A	US 2006/0036495 A1 (Aufrecht et al.) 16 February 2006 Entire document	
A	US 2005/01215238 A1 (Macaluso) 29 september 2005 Entire document	

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/AU2009/001261

This Annex lists the known "A" publication level patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document Cited in Search Report		Patent Family Member					
US	2008215429	AU	2007219274	CA	2605315	CN	101152667
		EP	1905461	EP	1924936	EP	2076856
		JP	2008080403	KR	20080066676	US	7302993
		US	7369056	US	7548915	US	7577665
		US	7603360	US	2007060099	US	2007060109
		US	2007060114	US	2007060129	US	2007060136
		US	2007060173	US	2007061197	US	2007061198
		US	2007061211	US	2007061229	US	2007061242
		US	2007061243	US	2007061244	US	2007061245
		US	2007061246	US	2007061247	US	2007061300
		US	2007061301	US	2007061302	US	2007061303
		US	2007061317	US	2007061328	US	2007061331
		US	2007061332	US	2007061333	US	2007061334
		US	2007061335	US	2007061336	US	2007061363
		US	2007073717	US	2007073718	US	2007073719
		US	2007073722	US	2007073723	US	2007094042
		US	2007100650	US	2007100651	US	2007100652
		US	2007100653	US	2007100805	US	2007100806
		US	2007109142	US	2007118533	US	2007168354
		US	2007192294	US	2007192318	US	2007198485
		US	2007239724	US	2007260635	US	2007288427
		US	2008009268	US	2008118771	US	2008121365
		US	2008191897	US	2008214148	US	2008214149
		US	2008214150	US	2008214151	US	2008214152
		US	2008214153	US	2008214154	US	2008214155
		US	2008214156	US	2008214157	US	2008214162
		US	2008214166	US	2008214204	US	2008215428
		US	2008215475	US	2008215557	US	2008215623
		US	2008242279	US	2008270220	US	2009029687
		US	2009222329	US	2009234711	US	2009234745
		US	2009234861	US	2009240568	US	2009240569
		US	2009240586	WO	2007033358	WO	2008052205

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/AU2009/001261

		WO 2009002999					
US	2008005695	NONE					
US	2007088852	EP	1940486	EP	1974280	EP	1977326
		EP	2013744	US	2007088801	US	2007088838
		US	2007088851	WO	2007046086	WO	2007083301
		WO	2007086051	WO	2007129298		
US	6892217	US	6826614				
US	2006036495	AU	46744/00	GB	2371902	US	6341316
		US	6421717	US	6553412	US	6779042
		US	6839744	US	7000032	US	7263547
		US	7392308	US	2001047394	US	2001051979
		US	2002046296	US	2002052781	US	2002052916
		US	2006036493	US	2006036494	US	2006136571
		US	2008120184	US	2008134031	US	2008215672
		WO	0118688				
US	20051215238	NONE					
WO	2009198580	US	2009198580				
Due to data integration issues this family listing may not include 10 digit Australian applications filed since May 2001.							
END OF ANNEX							