



US 20050283468A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2005/0283468 A1**

Kamvar et al. (43) **Pub. Date: Dec. 22, 2005**

(54) **ANTICIPATED QUERY GENERATION AND PROCESSING IN A SEARCH ENGINE**

(52) **U.S. Cl. 707/3**

(76) **Inventors: Sepandar D. Kamvar, Palo Alto, CA (US); Taher H. Haveliwala, Mountain View, CA (US); Glen M. Jeh, San Francisco, CA (US)**

(57) **ABSTRACT**

Correspondence Address:
MORGAN, LEWIS & BOCKIUS, LLP.
2 PALO ALTO SQUARE
3000 EL CAMINO REAL
PALO ALTO, CA 94306 (US)

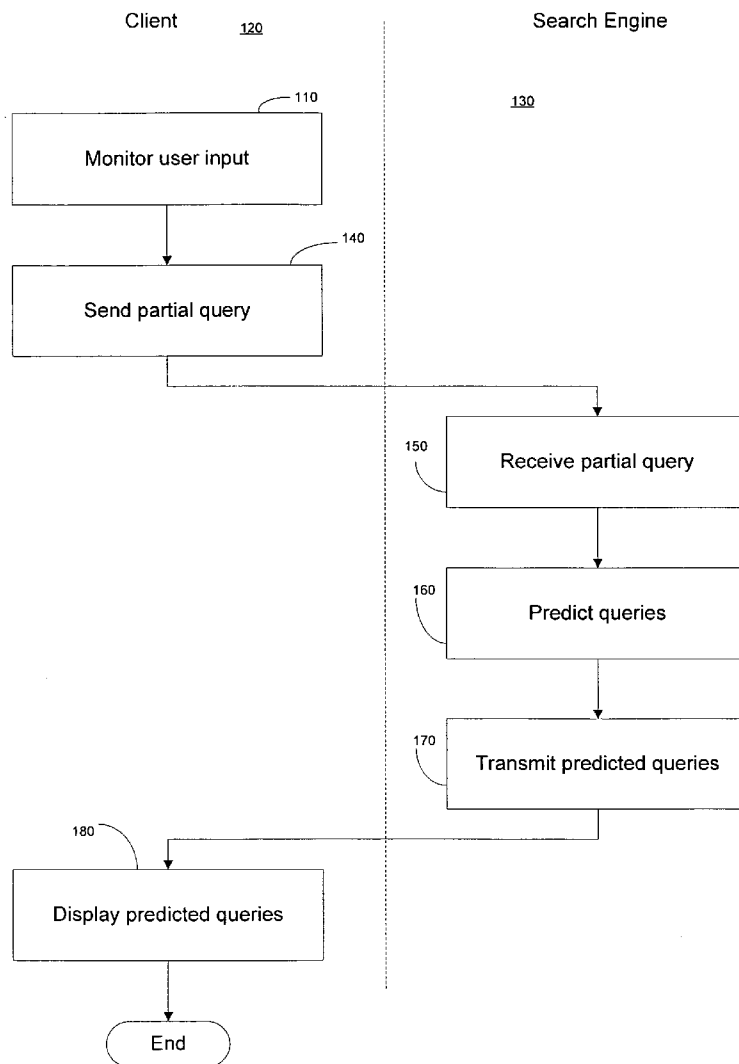
A search system monitors the input of a search query by a user. Before the user finishes entering the search query, the search system identifies and sends a portion of the query as a partial query to the search engine. Based on the partial query, the search engine creates a set of predicted queries. This process may take into account prior queries submitted by a community of users, and may take into account a user profile. The predicted queries are sent back to the user for possible selection. The search system may also cache search results corresponding to one or more of the predicted queries in anticipation of the user selecting one of the predicted queries. The search engine may also return at least a portion of the search results corresponding to one or more of the predicted queries.

(21) **Appl. No.: 10/875,143**

(22) **Filed: Jun. 22, 2004**

Publication Classification

(51) **Int. Cl.⁷ G06F 17/30; G06F 7/00**



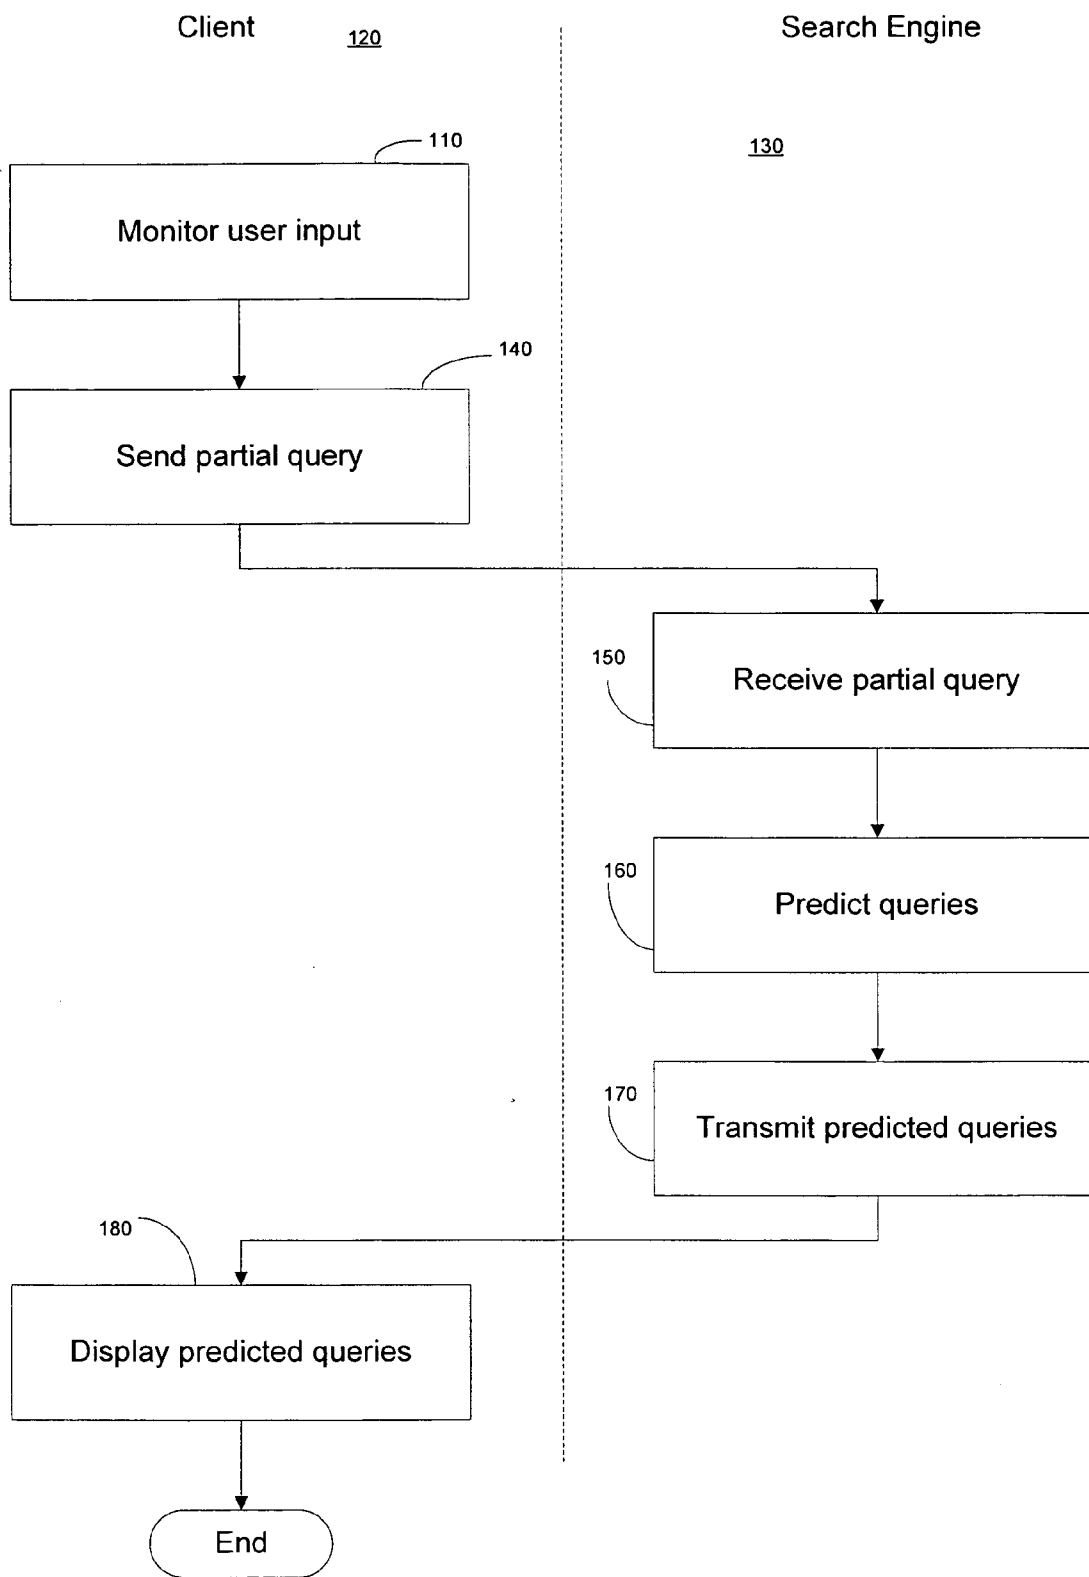


Fig. 1

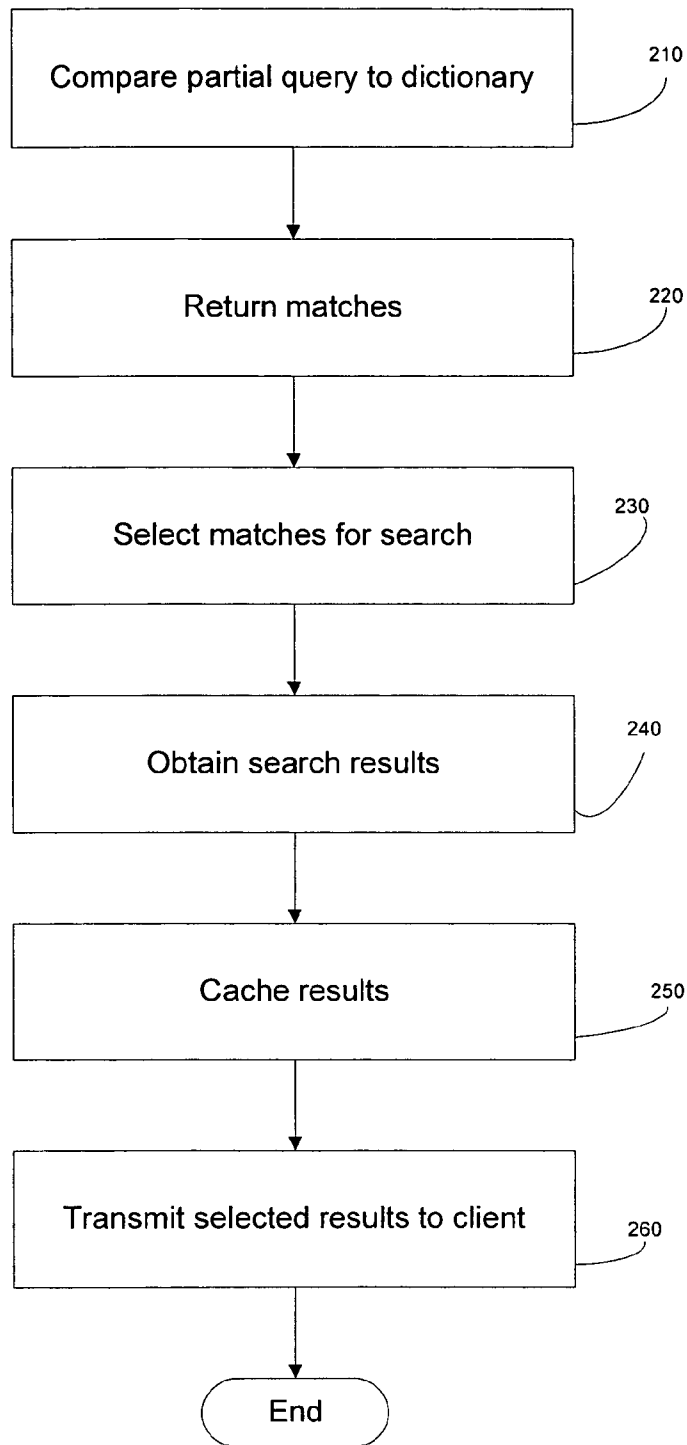


Fig. 2

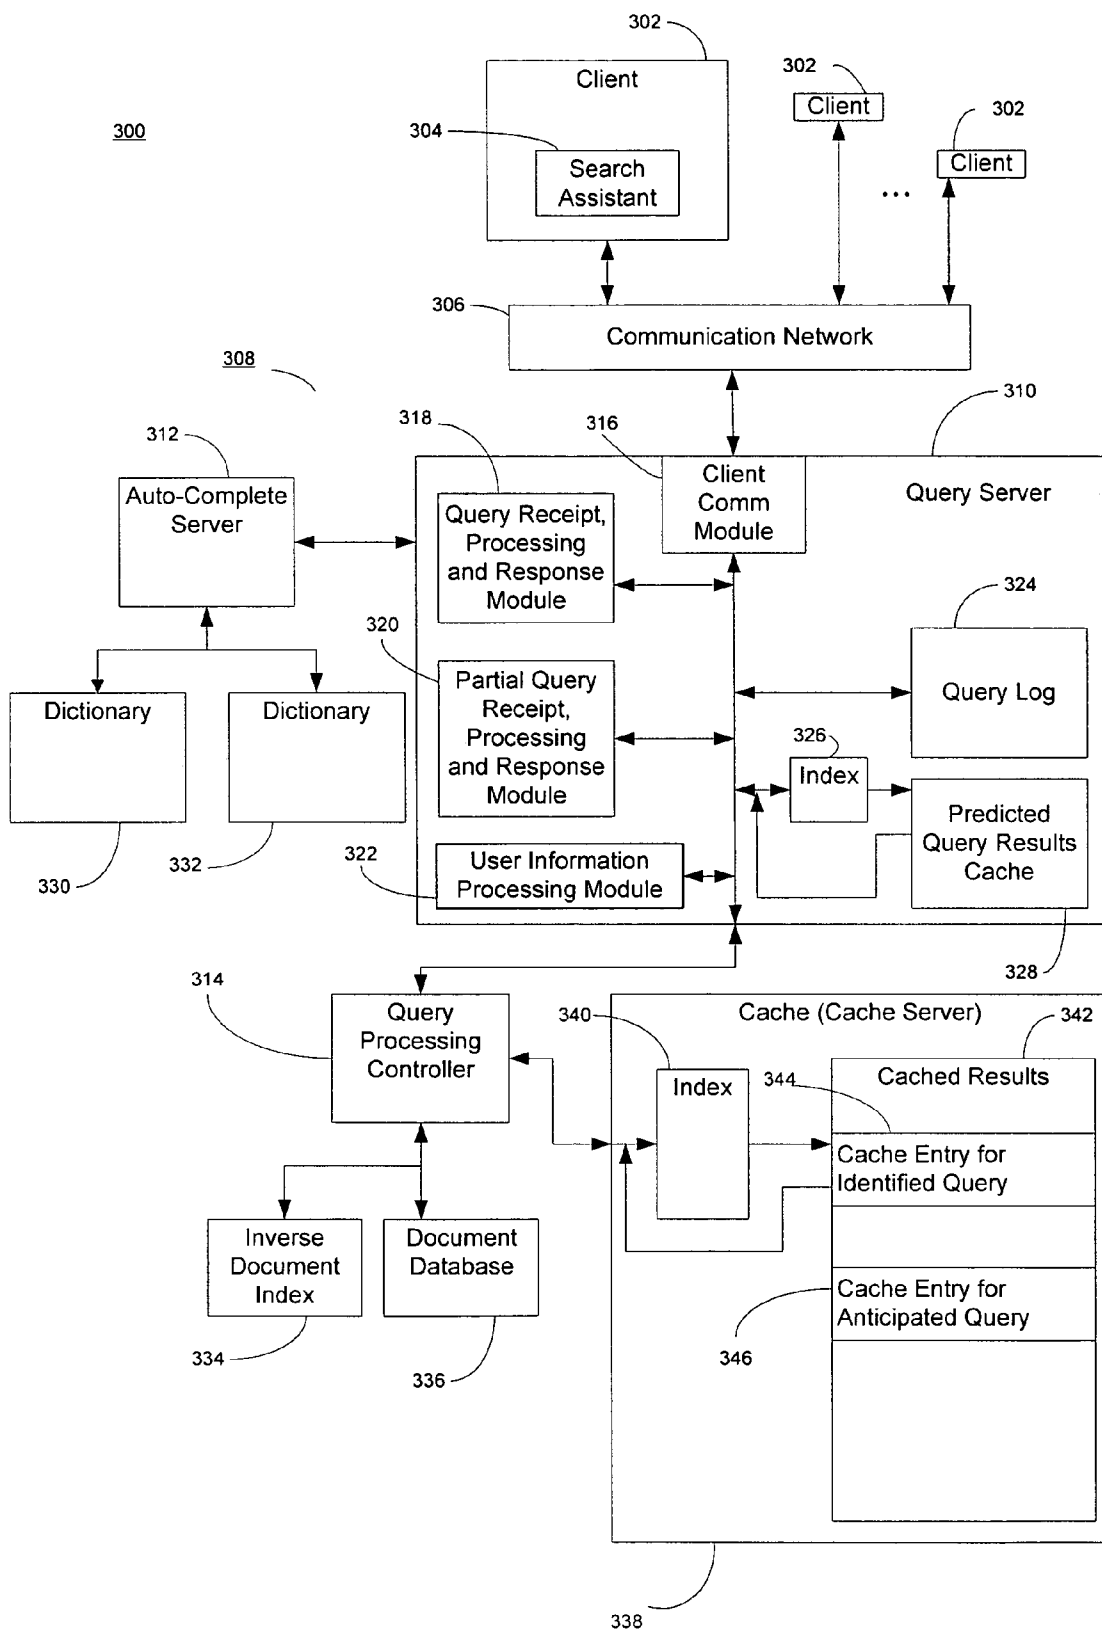


Fig. 3

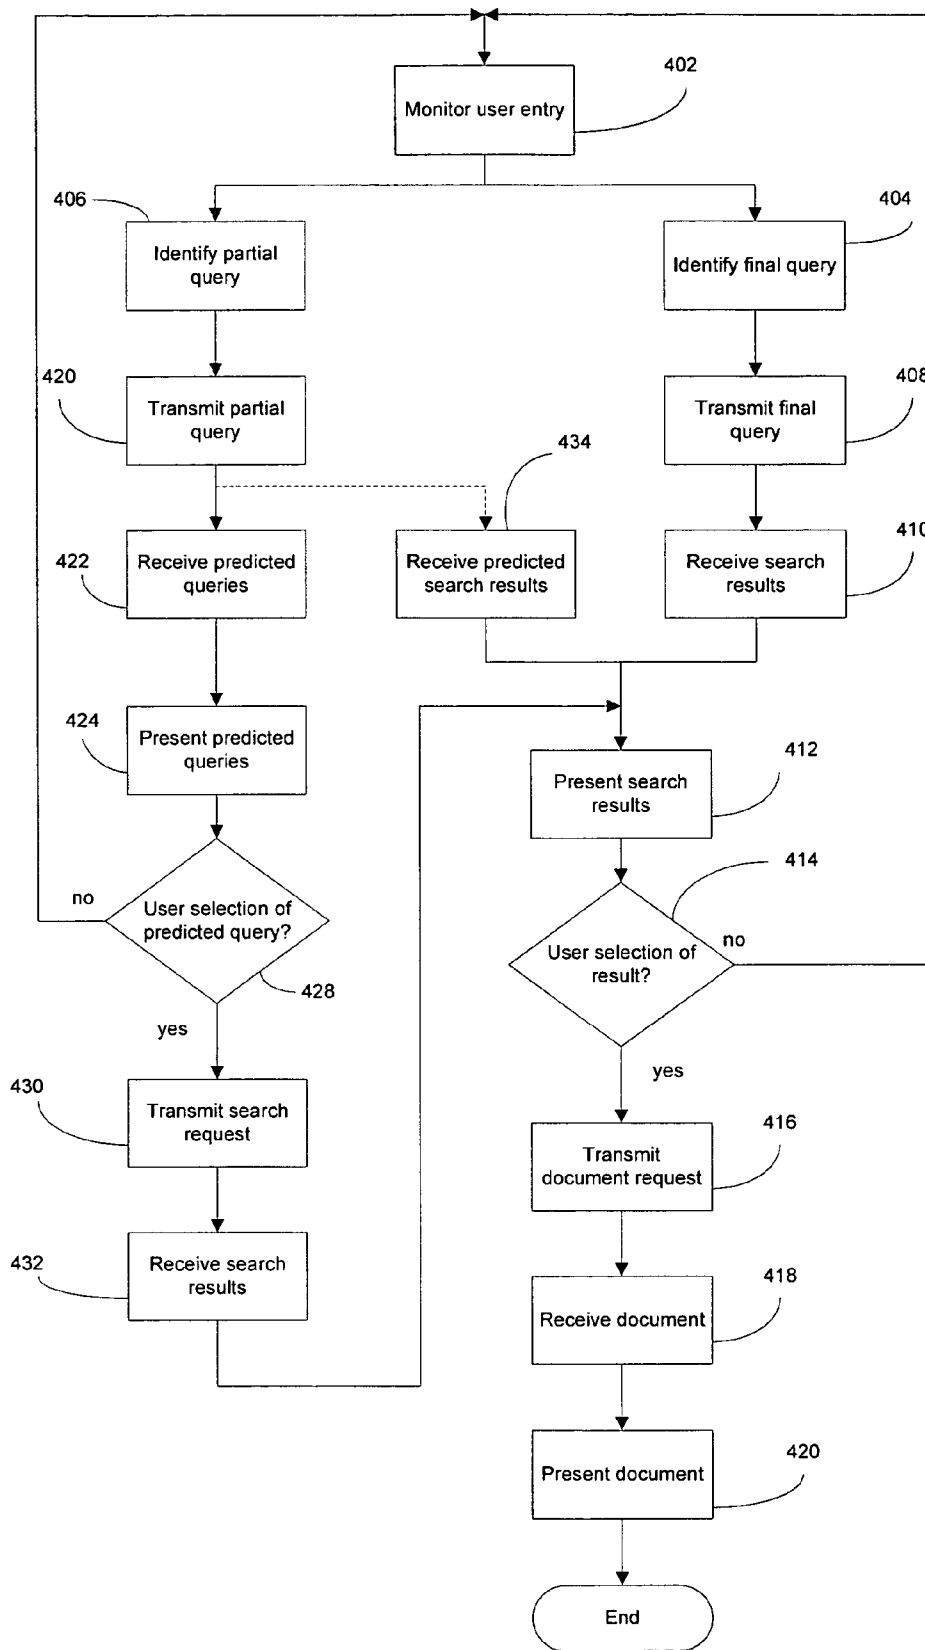


Fig. 4

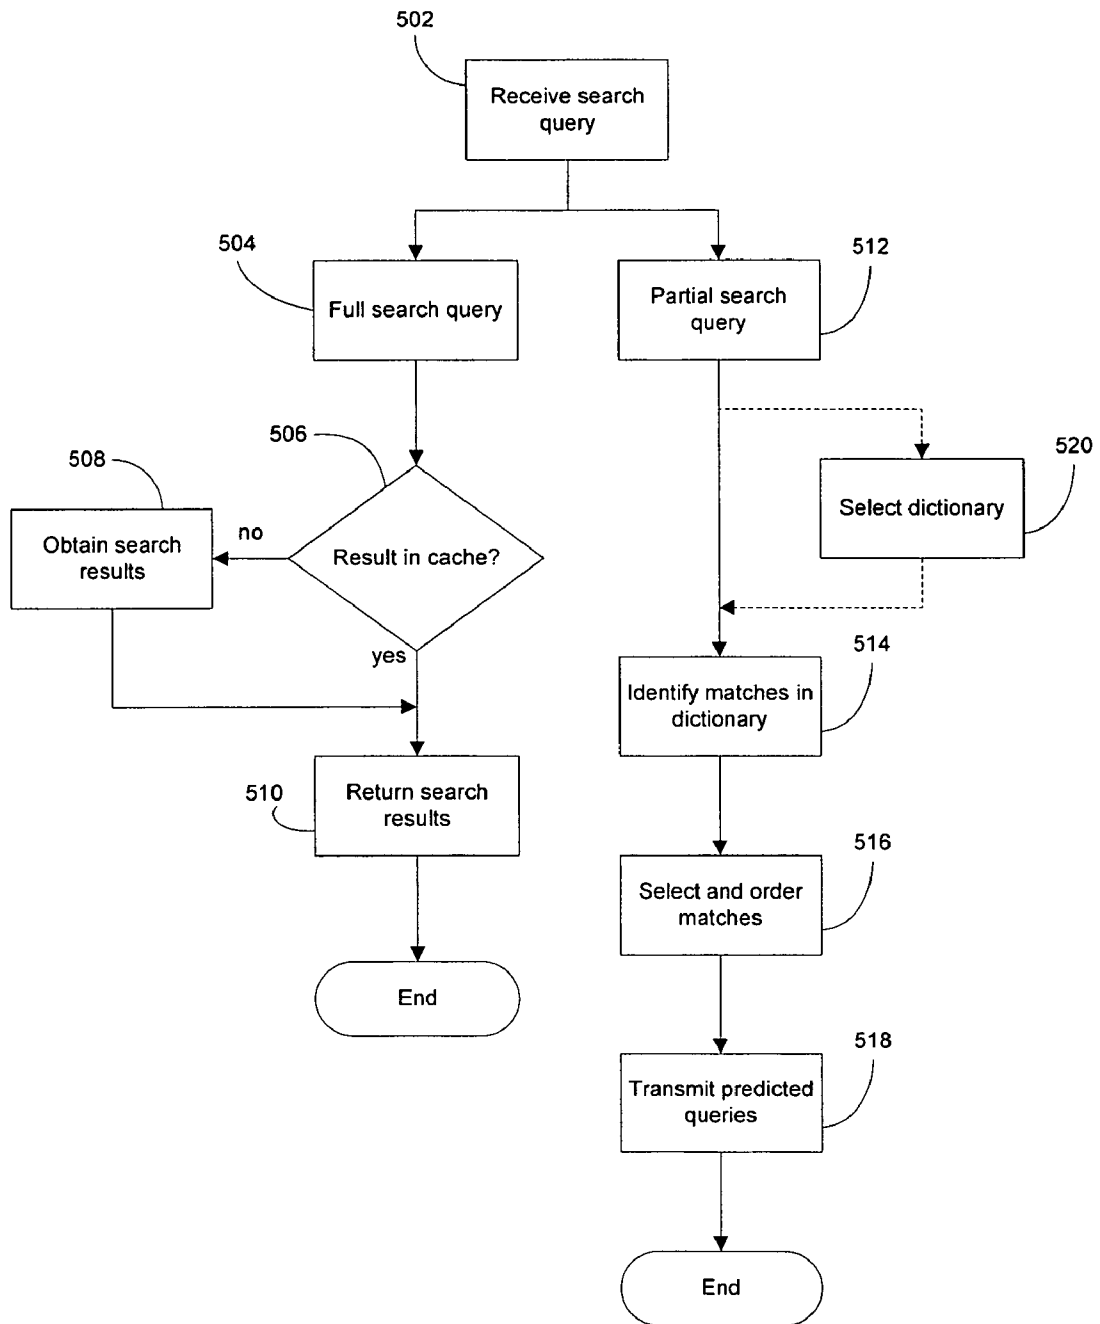


Fig. 5

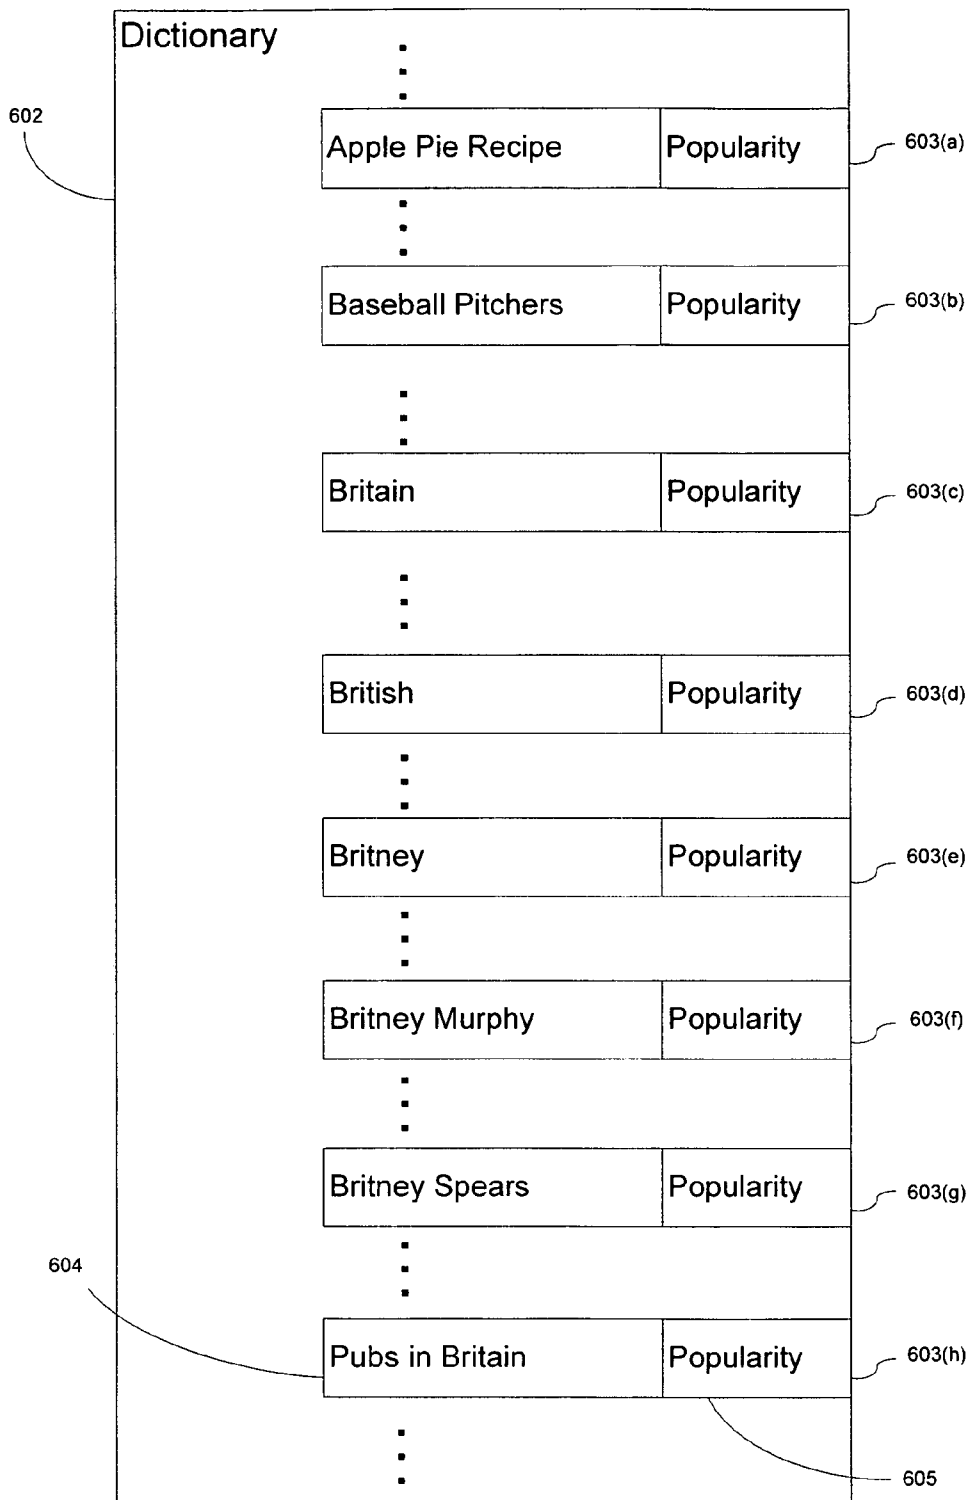


Fig. 6

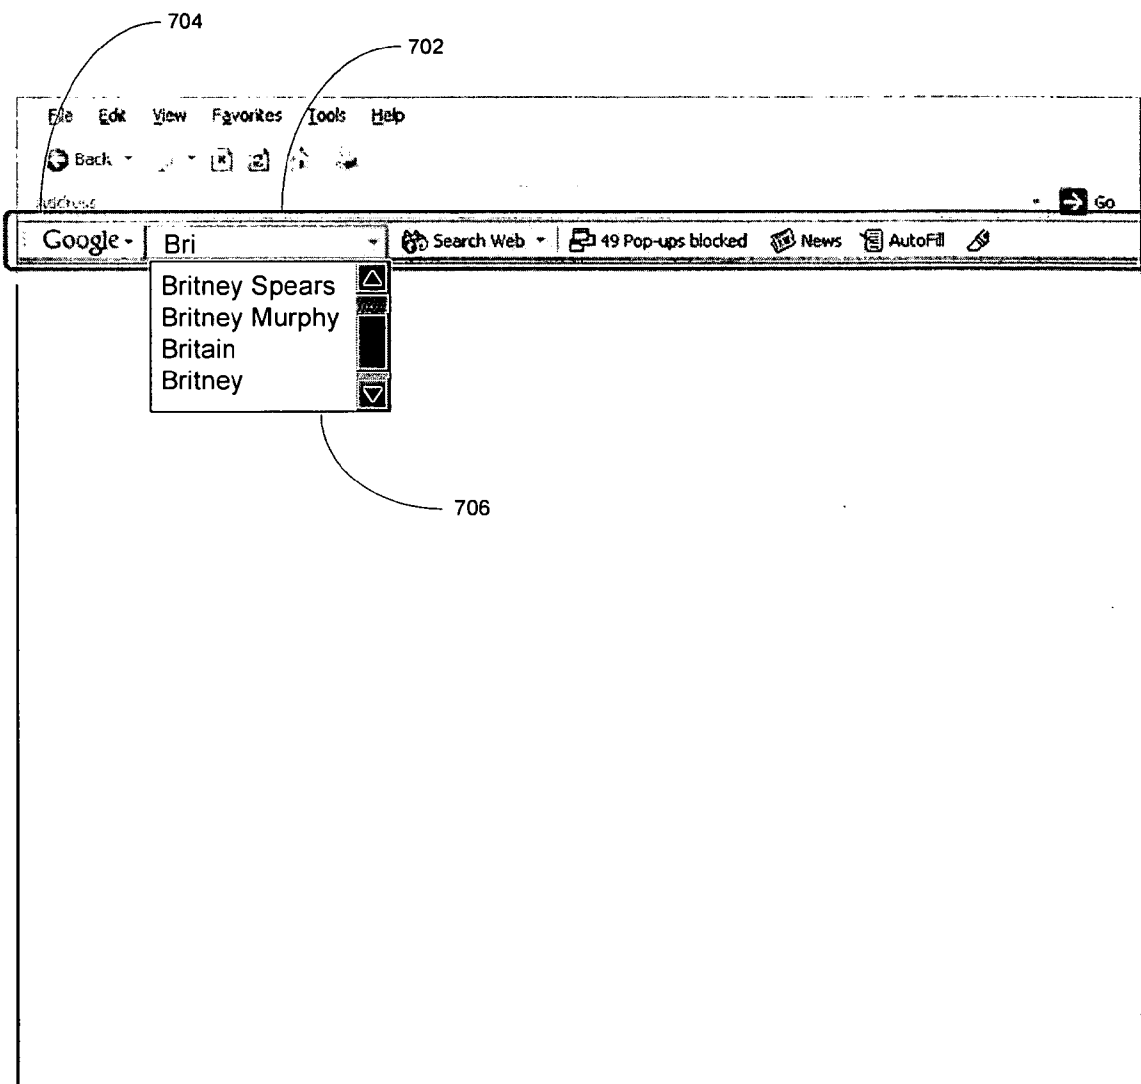


Fig. 7




Fig. 8

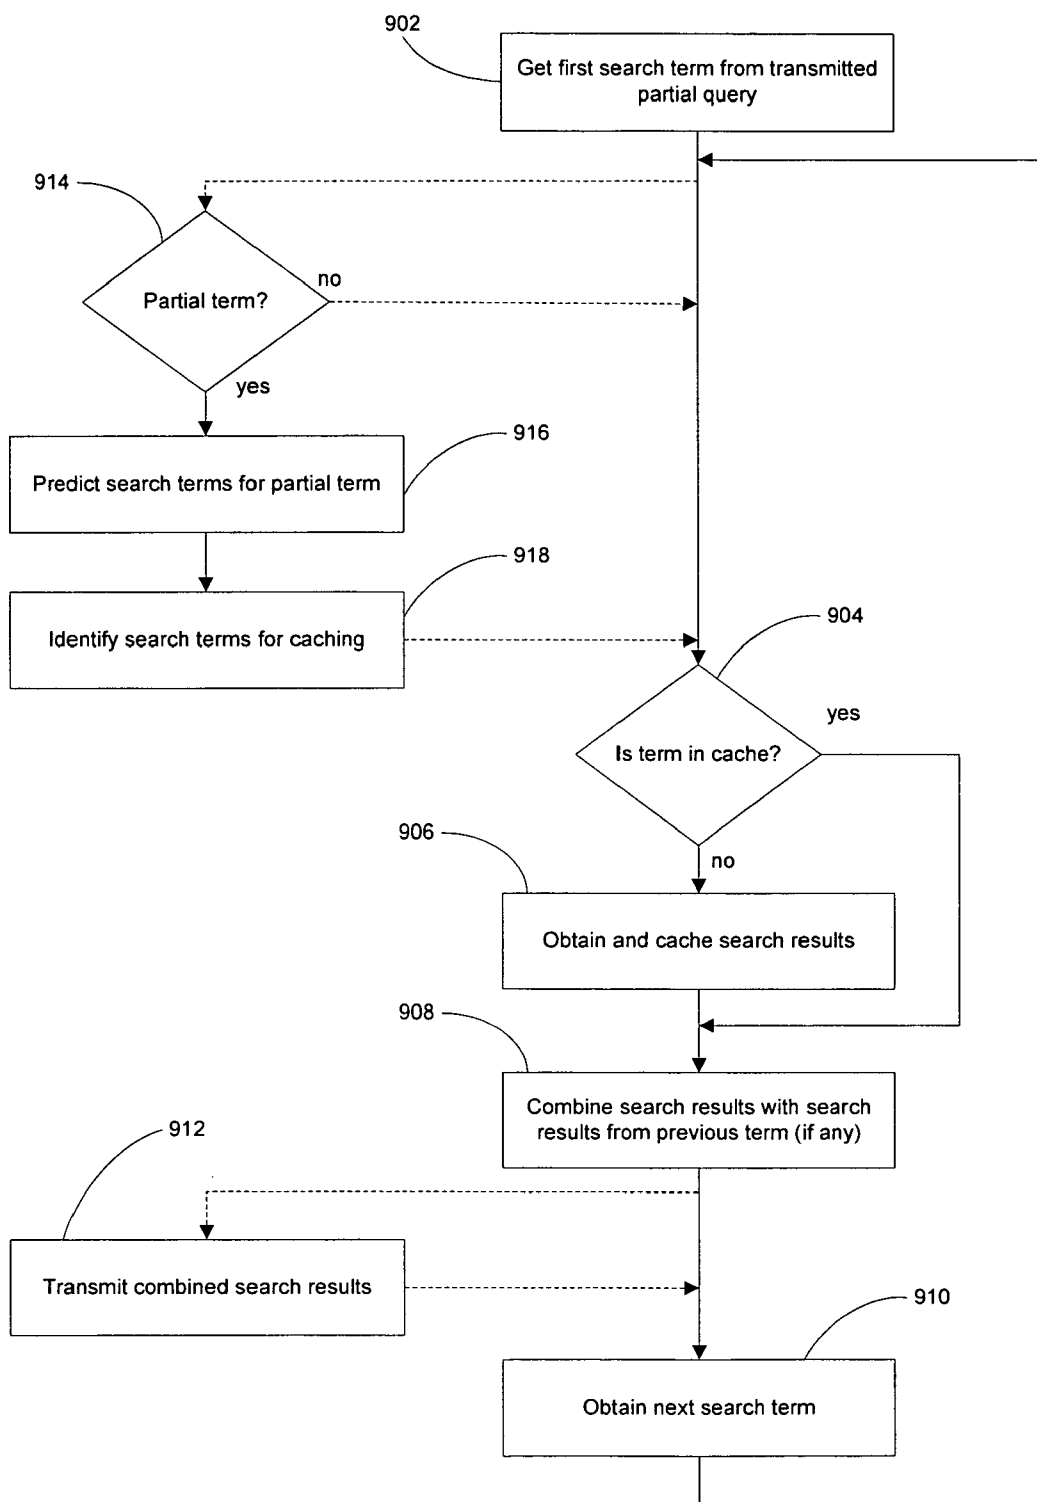


Fig. 9

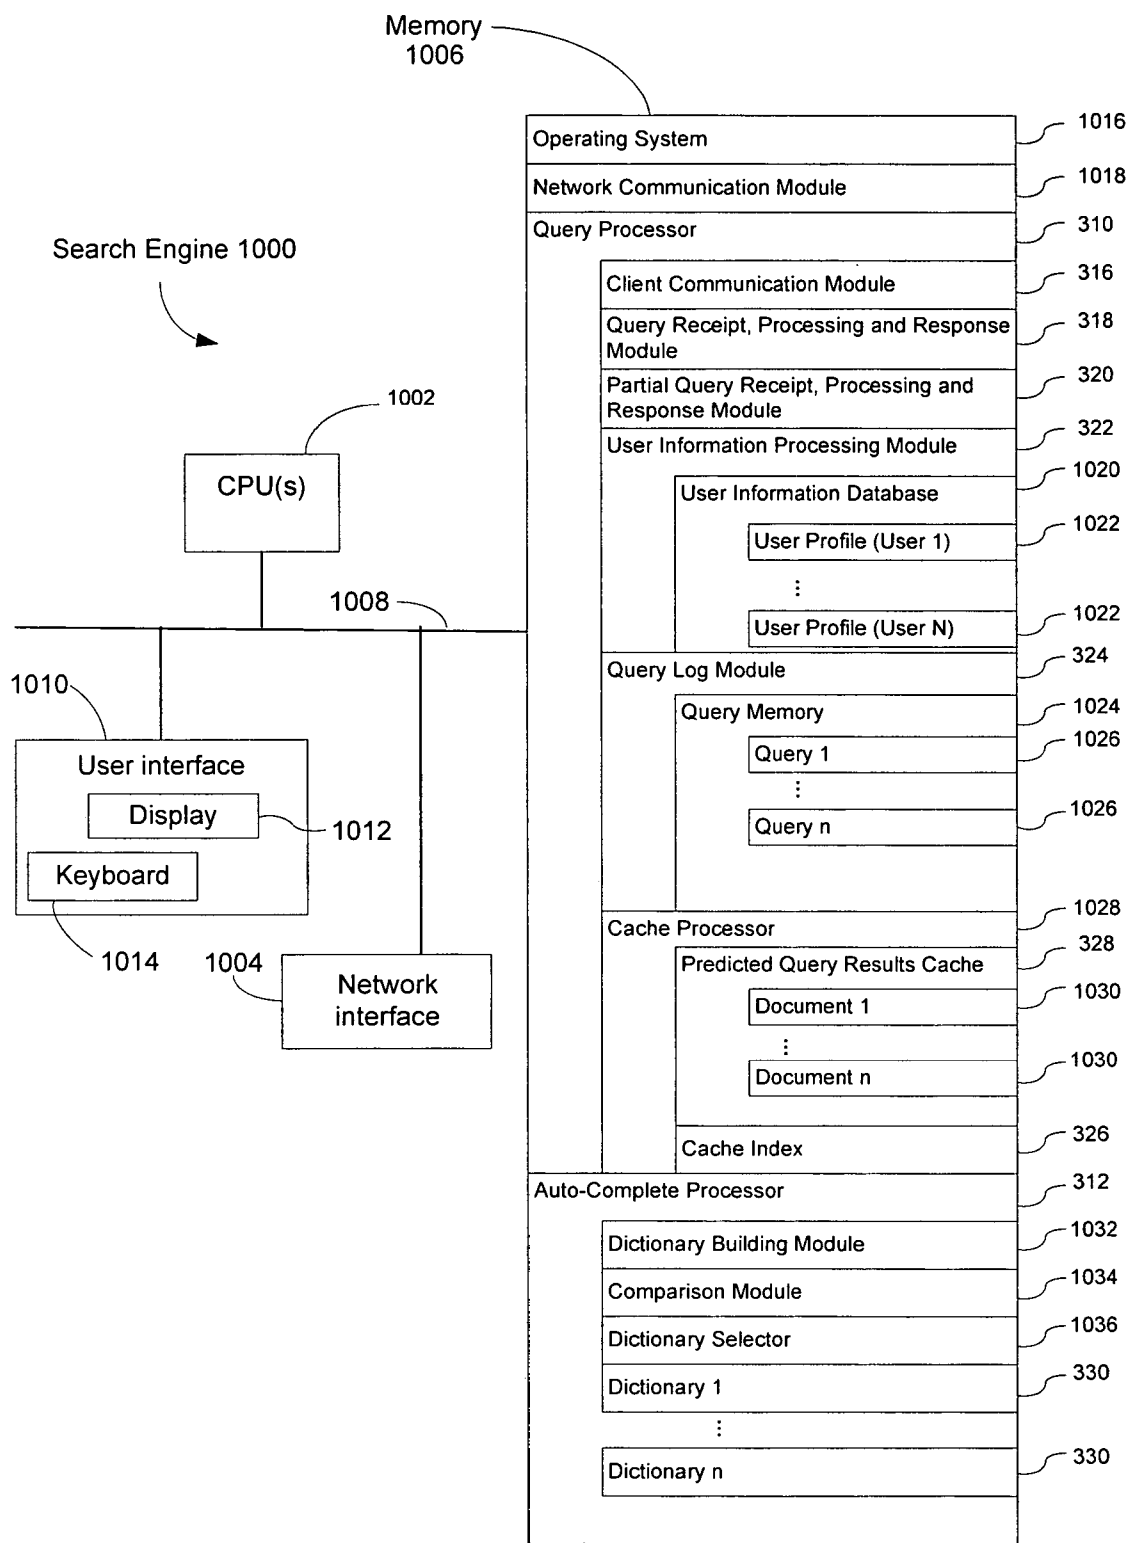


Fig. 10

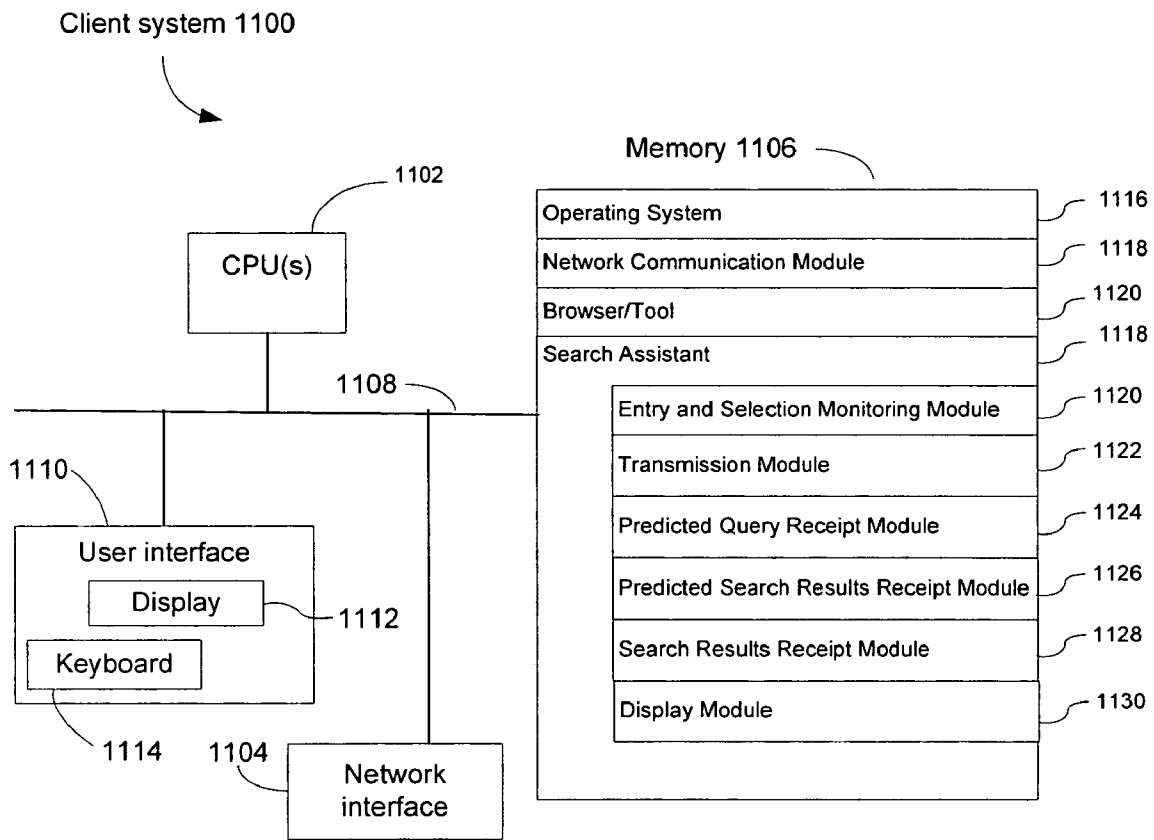


Fig. 11

ANTICIPATED QUERY GENERATION AND PROCESSING IN A SEARCH ENGINE

FIELD OF THE INVENTION

[0001] The present invention relates generally to the field of search engines for locating documents in a computer network system, and in particular, to a system and method for speeding up a desired search by anticipating a user's request.

BACKGROUND OF THE INVENTION

[0002] Search engines provide a powerful tool for locating documents in a large database of documents, such as the documents on the World Wide Web (WWW) or the documents stored on the computers of an Intranet. The documents are located in response to a search query submitted by a user. A search query may consist of one or more search terms.

[0003] In one approach to entering queries, the user enters the query by adding successive search terms until all search terms are entered. Once the user signals that all of the search terms of the query are entered, the query is sent to the search engine. The user may have alternative ways of signaling completion of the query by, for example, entering a return character by pressing the enter key on a keyboard or by clicking on a "search" button on the screen. Once the query is received, the search engine processes the search query, searches for documents responsive to the search query, and returns a list of documents to the user.

[0004] Because the query is not sent to the query processor until the query is completed, time passes while the user is building the full search query. It would be desirable to have a system and method of speeding up this process.

SUMMARY

[0005] A search system monitors the input of a search query by a user. Before the user finishes entering the search query, the search system identifies and sends a portion of the query as a partial query to the search engine. Based on the partial query, the search engine creates a set of predicted queries. This process may take into account prior queries submitted by a community of users, and may take into account a user profile. The predicted queries are sent back to the user for possible selection. The search system may also cache search results corresponding to one or more of the predicted queries in anticipation of the user selecting one of the predicted queries. The search engine may also return at least a portion of the search results corresponding to one or more of the predicted queries.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The aforementioned features and advantages of the invention as well as additional features and advantages thereof will be more clearly understood hereinafter as a result of a detailed description of preferred embodiments of the invention when taken in conjunction with the drawings.

[0007] FIG. 1 is a flow chart of a query entry and handling process.

[0008] FIG. 2 is a flow chart of a process for handling a partial query.

[0009] FIG. 3 is a block diagram of a client system 120 and a search engine 130.

[0010] FIG. 4 is a flow chart of tasks handled by a search assistant in a client.

[0011] FIG. 5 is a flow chart of tasks handled by a search engine.

[0012] FIG. 6 is a block diagram of a dictionary of search queries.

[0013] FIG. 7 depicts a listing of predicted search queries.

[0014] FIG. 8 is a flow chart of a process for caching search results corresponding to predicted search queries.

[0015] FIG. 9 is a flow chart of a process for processing a search query having a plurality of search terms.

[0016] FIG. 10 is a block diagram of a search engine.

[0017] FIG. 11 is a block diagram of a client system.

[0018] Like reference numerals refer to corresponding parts throughout the several views of the drawings.

DESCRIPTION OF EMBODIMENTS

[0019] In an embodiment of the invention, portions of a user's query may be transmitted to a search engine before the user has finished entering the complete query. The search engine may use the transmitted portions of the query to predict the user's final query. These predictions may be sent back to the user. If one of the predictions is what the user had intended the query to be, then the user can select that predicted query without having to complete entry of the query.

[0020] FIG. 1 illustrates an exemplary embodiment of the invention including a client system 120 and a search engine 130. As a user enters the search query, the input is monitored by the client system (stage 110). Prior to the user signaling completion of the search query, a portion of the user's query is sent from the client system 120 to the search engine 130 (stage 140). The portion of the query may be a few characters, a search term, or more than one search term. The search engine 130 receives the partial query for processing (stage 150) and makes predictions of what the user's contemplated query might be (stage 160). The search engine 130 may use a dictionary to assist in making the predictions. The dictionary may be created using previously entered search queries received by the search engine 130. The previous queries may include search queries from a community of users. The predicted queries are sent back to the client system 120 (stage 170) and then presented to the user (stage 180). If one of the predicted queries is what the user intended as the desired query, the user may select this predicted query and proceed without having to finish entering the desired query. If the predicted queries do not reflect what the user had in mind, then the user may continue to enter more of the desired search query.

[0021] In another embodiment of the invention, predicted search results are obtained and may be presented to the user prior to the user signaling that the complete query has been entered. In this embodiment, the portion of the search query received at stage 150 may be compared against the entries of the dictionary as illustrated in FIG. 2 (stage 210). As a result of comparing the partial query against the dictionary, entries

in the dictionary matching the partial search query are then returned (stage 220). Various criteria may be used for identifying matching entries when a partial query is compared against the dictionary, as described below. It may be desired that not all of the matches be further processed. In that case, one or more matches from the set of returned matches may be chosen to have search results obtained for them (stage 230). The search results obtained (stage 240) may then be cached in memory for later access (stage 250). In some embodiments, search results are obtained only for the top N matches, where N is a predefined positive integer. In some embodiments, one or more caches of stored search results are checked prior to generating such search results, so as to avoid regenerating search results that have already been cached.

[0022] If it is desired that the search results be returned to the user, then results are transmitted to the client system 120 (stage 260), and may be presented to the user while the user is still entering the complete query. It may be that one of the search results for the predicted query satisfies the user's intended query. If so, the search engine 130 has, in effect, reduced the latency of a search from a user's perspective to zero.

[0023] FIG. 3 illustrates a searching system 300 according to an embodiment of the invention showing various functional components which will be referred to in the detailed discussions which follow. The searching system 300 may include one or more client systems 302. Each client system 302 may have a search assistant 304. The client systems 302 are connected to a communications network 306. The communications network 306 is connected between the client systems 302 and a search engine 308. Search engine 308 includes a query server 310 connected to the communications network 306, an auto-complete server 312, and a query processing controller 314.

[0024] A query server 310 includes a client communications module 316, a query receipt, processing and response module 318, a partial query receipt, processing and response module 320, a user information processing module 322, a query log 324, an index 326, and a predicted query results cache 328, all interconnected. In some embodiments, a results cache 328 is not included. In some embodiments, fewer and/or additional modules or functions are included in the query server 310. The modules shown in FIG. 3 as being part of query server 310 represent functions performed in an exemplary embodiment.

[0025] An auto-complete server 312 is connected to dictionary 330 and to dictionary 332. Although illustrated as connected to only two dictionaries, one of ordinary skill in the art will recognize any number of dictionaries could be present and connected as will be described below.

[0026] The query processing controller 314 is connected to an inverse document index 334, and document database 336, and a query cache 338. The cache 338 may include an index 340 the function of which is to locate entries in the cached results 342. The cached results may include a cache entry for an identified query 344 and a cache entry for an anticipated query 346. The index 334 and document database 336 are sometimes collectively called the document database. In some embodiments, "searching the document database" means searching the inverse document index 334 to identify documents matching a specified search query or term.

[0027] Although illustrated as discrete blocks in the figure, FIG. 3 is intended more as a functional description of an embodiment of the invention rather than a structural mapping of the functional elements. One of ordinary skill in the art would recognize that an actual implementation might have the functional elements grouped or split among various components.

[0028] FIG. 4 illustrates an embodiment of the invention that may be implemented in the search assistant 304. The search assistant 304 monitors the user's entry of a search query on the client system 120 (stage 402). The user may enter a search query in a number of ways including a browser window, a search tool, or any other input mechanism. The search assistant 304 may identify either a final search query (stage 402) or a partial query (stage 406).

[0029] A final search query may be identified by the search assistant 304 in a number of ways such as when the user enters a carriage return, or equivalent character, selects a search button presented to the user during entry of the search query, or by possibly selecting one of a set of possible queries presented to the user during entry of the search query. One of ordinary skill in the art would recognize a number of ways to signal the final entry of the search query. Once the final search query has been identified (stage 404), the search query is transmitted to the search engine 308 (stage 408) for processing and return of search results. After the search results generated from the query are received (stage 410), they are presented to the user such that the user may select one of the documents for further examination (stage 412). For example, the search results could be visually or audibly presented to the user. One of ordinary skill in the art would recognize a number of ways to present the search results to the user for potential selection.

[0030] If the user does not select any of the presented search results (stage 414), then the user may begin to enter another search request and monitoring will begin again (stage 402). If, on the other hand, the user selects one of the presented results (stage 414), the document is requested from its host (e.g., a host web server or host document server) (or from the search engine 308) for further examination by the user (stage 416). The document is received from the host or the search engine 130 (stage 418) and then presented to the user (stage 419). Presentation of the document may be similar to, or different from, presentation of the search results (stage 412) (e.g., visually or audibly).

[0031] The stages described above follow when the user signals that the query is complete. Prior to the end of the entry of the search query, however, a partial query may be identified (stage 406). A partial query may be identified in a number of ways. A partial query might include a single search term of the search query, multiple search terms, or it might include a number of characters of a search term.

[0032] One embodiment of the invention may include identifying a search term by detecting entry of a space character or other search term delimiting character (e.g., without limitation, a quote character, or a parenthesis character). Entry of a delimiting character may indicate that a user has finished entering a desired search term and is moving onto the next search term.

[0033] Another embodiment may include identifying a partial query by entry of a pre-determined number of char-

acters. In this embodiment, it may be that a search term contains a number of characters and it is desirable to identify a partial query before the user has entered all of the characters. This technique may be desirable, for example, if the search term contains a large number of characters or if the pre-determined number of characters is large enough to result in useful predictions.

[0034] Yet another embodiment might include identifying a partial query by the absence of a character being entered within a period of time, representing a pause by the user. It may be that the pause in entry signifies that the user has entered one search term but has not entered the space key (or other delimiting character) to start entering another term or signifies that the search query is complete but the user has not yet signaled that the search query is complete.

[0035] Regardless of the way the partial query is identified, it is transmitted to the search engine 308 (stage 420) for processing. In response to the partial search query, the search engine 308 returns a set of predicted search queries (stage 422) which may be presented to the user (stage 424). One skilled in the art will recognize a number of ways to present the predicted search queries to the user. For example, the predicted search queries might be presented in a drop down menu. Regardless of the manner in which the predicted queries are presented to the user, the user may select one of the queries if the user determines that one of the predicted queries matches a desired query. If the user selects one of the predicted search queries (stage 426), the predicted query is transmitted to the search engine 308 as a search request (stage 430). After the search results generated from the search request are received (stage 432), they may be presented and selected for further study as described above (stages 414 through 419).

[0036] In an embodiment of the invention, the search engine 308 may optionally return predicted search results (stage 434). This activity may overlap with receiving predicted queries and is so indicated by the dashed line in FIG. 4 from stage 420 to stage 434. The received predicted search results may be presented and selected as described earlier (stages 414 through 419).

[0037] Accordingly, the user may have predicted search results presented that match a desired query before the user finishes entering the query. In such situations, the query processing latency as viewed by the user is effectively reduced to less than zero because the user did not have to complete the query to obtain the desired results.

[0038] An embodiment of the invention for receiving queries or partial queries and returning search results and predicted queries is illustrated in FIG. 5. After a query is received (stage 502) at a search engine 308, for example, it is identified as either a full search query (stage 504) or a partial search query (stage 512). If the search query is a final search query, it is determined whether search results corresponding to the search query are present in a cache of the search engine (stage 506). For example, the search results may have been cached from an earlier search on the same query. The earlier search may be as a result of a final query or a predicted query (as described below). Referring back to FIG. 3, the cached search results could be in the query server 310 or the cache 338. Typically, results stored in a cache will be available more quickly (i.e., results are presented with lower latency as viewed from the perspective of

the user) than producing those search results from a search of the search engine's document database. Producing search results from a cache also reduces the workload of the search engine, which improves overall performance of the search engine. If the results of the search query are not present in any cache, then the search results are obtained by searching the inverse document index 334 (stage 508). Once the results are obtained from the cache or document index, they are sent to the client system 120 (stage 510).

[0039] If a partial query is identified (stage 512), then an attempt is made to anticipate queries that the user might be in the process of entering based on that partial query. In one embodiment, the partial query is compared against entries in a dictionary to generate possible queries (stage 514). For example, the partial query receipt, processing and response module 320 (of FIG. 3) could send the partial query to the auto-complete server 312. The auto-complete server 312 could return a set of entries from the dictionary that match the partial query. From the returned set, some matches could be selected and placed in an order based on a metric or score representative of how likely each entry is to match the user's search query (stage 516). This selecting and ordering is described in more detail below. The selected and ordered matches may be sent to the client system 120 as predicted queries (stage 518). In some embodiments, matching entries are selected based on a metric or score, but are ordered based on other criteria, such as alphabetical order.

[0040] In an embodiment of the invention, there may be more than one dictionary against which to match partial queries. The dotted lines in FIG. 5 represent an alternate path between stage 512 and stage 514, where a specific dictionary (or a plurality of specific dictionaries) may be selected (stage 520). Multiple dictionaries may allow for personalization of the matched entries in a dictionary. For example, a user profile may identify one or more categories of information that are of interest to a respective user. Each dictionary could correspond to one or more categories of users or categories of interest (e.g., sports, music, news, finance, food, popular culture, etc.).

[0041] FIG. 6 schematically depicts an exemplary dictionary 602. A dictionary 602 according to an embodiment of the invention may be created from previously submitted search queries. These previous search queries could consist not only of queries submitted by the user, but also of queries submitted by a community of users. Thus, the entries in the dictionary 602 would tend to reflect commonly submitted searches as well as those recent search requests that might still be stored in a cache somewhere in the searching system 300. As noted above, if a search request is stored in a cache, the search results may be more quickly presented to the user than if the search required searching the document database 336 from scratch.

[0042] The dictionary 602 includes a number of exemplary dictionary entries 603(a) through 603(h). Each of the entries 603 includes a term portion 604 and a popularity value 605. The term portion 604 stores the single or multi-word terms which could be identified as a possible query search term. The popularity value 605 is a value that indicates how popular the associated term portion 604 may be at any given moment in time. The popularity value 605 can be used to rank the predicted search queries according to current popularity and therefore increase the likelihood that the user will be presented with an intended query term.

[0043] For the purposes of illustration, only a few exemplary entries are shown. And, while the entries 603 in dictionary 602 are shown in alphabetical order, the entries may be stored in any order. Furthermore, while the exemplary dictionary 602 includes both single search term entries and multiple search term entries, dictionaries consisting of entries of either or both are also contemplated by embodiments of the invention.

[0044] Consider matching “Bri” as a partial search query against the entries 603 in dictionary 602. In one embodiment, matches in the dictionary 602 are identified if a respective dictionary entry begins with the same letters as the partial search query. Using this criteria, the partial query “Bri” would match entry 603(c) “Britain”, entry 603(d) “British”, entry 603(e) “Britney”, entry 603(f) “Britney Murphy”, and entry 603(g) “Britney Spears”, but would not match “Apple Pie Recipe”. In this example, capitalization is taken into account, but in another embodiment would not be.

[0045] Consider another example where “Britney” is the partial search query. In this example, the matches could include entry 603(e) “Britney”, entry 603(f) “Britney Murphy”, and entry 603(g) “Britney Spears”. As more characters are included in the partial search query, the number of dictionary entries matching the partial search query may diminish, increasing the likelihood that the predicted search query will match the user’s desired query. However, there is a tradeoff between waiting for more information to be included in the partial search query (and thus increasing the likelihood that the predicted search query will match the desired search query) against more quickly returning a larger number of search queries, many of which may not be the desired search query.

[0046] In another embodiment, entries 603 matched to a partial search query in the dictionary 602 are identified if any search term in the entry begins with the same letters as the partial search query. For example, if the partial search query was “Bri” this partial query would match entry 603(c) “Britain”, entry 603(d) “British”, entry 603(e) “Britney”, entry 603(f) “Britney Murphy”, and entry 603(g) “Britney Spears” as in the previous example, but would also match 603(h) “Pubs in Britain”. In entry 603(h), the third search term “Britain” was matched and thus the entire query is considered a match. Although only a few exemplary matching embodiments have been discussed, one of ordinary skill in the art will recognize any number of possible ways to match dictionary entries.

[0047] Referring back to FIG. 5, once the set of matching entries has been identified (stage 514), the matching entries are selected and ordered (stage 516) prior to transmission to the client system 120. Any number of selection and ordering strategies are contemplated as embodiments of the invention. For example, all matches could be selected and then ordered in alphabetical order. In another embodiment, the N most popular searches from all of the matched entries could be selected and presented in order of popularity. In another embodiment, the N most likely results could be selected which most closely match the user’s partially entered query, based on a predefined matching metric (which might also take into account a user profile).

[0048] [We will replace FIG. 7 with a schematic screen shot.] FIG. 7 illustrates one possible embodiment of the invention where “Bri” was entered as a partial search query

in a browser toolbar. In FIG. 7, the user enters text into a search entry box 702 in a browser tool bar 704. Referring to both FIG. 3 and FIG. 7, the partial query “Bri” is monitored and may be sent by the search assistant 304 through communications network 306 to query server 310. The client module 316 may receive the partial query which is then routed to the partial query, receipt processing and response module 320. The partial query receipt, processing and response module 320 sends the partial query to auto-complete server 312 which may return matches against a dictionary 330 which satisfy predefined matching criteria such as the first letters of the partial query matching the first letters of the dictionary entries. The auto-complete server may return a set of matches similar to ones as described above in an earlier example: 603(c) “Britain”; entry 603(d) “British”; entry 603(e) “Britney”; entry 603(f) “Britney Murphy”; and entry 603(g) “Britney Spears”. The partial query receipt, processing and response module 320 could receive those matches and use information in the query log 324, in the cache 338 itself, or in the dictionary 330 to determine the popularity of each of those matches.

[0049] For the purposes of this example, consider the popularity ranking to be, in decreasing order of popularity: “Britney Spears”; “Britney Murphy”; “Britain”; “Britney”; and “British”. Also, for this example, the partial query receipt, processing and response module 320 is configured to return only the top four queries in popularity. By accounting for popularity, the search engine 308 may be able to take into account popularity spikes due to such reasons as breaking news issues or popularity fads. The partial query receipt, processing and response module 320 would return the selected and ordered set of: “Britney Spears”; “Britney Murphy”; “Britain”; and “Britney”. This set could then be presented to the user as shown in select box 706 of FIG. 7.

[0050] In an embodiment, the query log 324 could contain a popularity ranking for each of the queries stored in it. In another embodiment, the query log 324 could contain information indicating when in time the query was submitted. In this way, an embodiment could select and order the matches based on when the query was last submitted. One could order the matches such that the most recently requested queries are presented to the user before queries less recently submitted. Other embodiments could be envisioned including selecting the most popular N matches as determined from the query log 324 and ordering them based on which queries were most recently submitted. One of ordinary skill in the art would recognize many different ways to select and order the matching entries. In some alternate embodiments, information about the popularity of queries whose search results have been cached is stored in the cache itself. For instance, the cache may maintain a “reuse count” indicating the number of times each entry in the cache has been re-used. The reuse count for each specified entry may be returned by the cache in response to a re-use count inquiry. In such embodiments, a query log 324 might not be used for selecting and/or ordering predicted search queries.

[0051] In yet another embodiment, user profile information about users is stored in user information processing module 322. The user information processing module 322 might contain such information as the user having a preferred interest in entertainers. As such, the entries “Britney Spears” and “Britney Murphy” would be given more weight in the matching, selecting, and ordering than other terms. By

use of information which might be contained in (or accessible by) the user information processing module 322, the search engine 308 might be able to further increase the likelihood that the predicted search queries match the user's desired query.

[0052] If the predicted search queries present in select box include the user's desired query, the user can select that query using an appropriate keyboard, mouse or other input. Imagine that the user's desired search query was "Britney Spears". The user could select the "Britney Spears" entry in select box 706 and this would be transmitted to the search engine as a final query.

[0053] In yet another embodiment of the invention, some or all of the search results from the predicted queries are cached in the search engine such that when the user selects one of the predicted queries in for example, select box 706, the search engine is not required to search the document database. This embodiment would serve to anticipate possible final queries and store the search results of those queries in a cache so that they are available more quickly than if the search engine was required to complete the search from scratch. An embodiment of this type is described next.

[0054] Referring to FIG. 8, the predicted queries that are generated as a result of stage 516 of FIG. 5 are received (stage 802). In another embodiment, the received set could be the matches identified at stage 514 prior to selecting and reordering. Out of the received set, certain of the queries are selected to have their search results cached in anticipation that those predicted queries might be selected by the user as the final query (stage 804). Any number of factors may be used to select queries for having their results cached. Identifying which search results to cache could be done similarly to identifying which queries to transmit to the user as identified in stages 516 and 518 of FIG. 5. That is, one could use the popularity of a query or how recently in time a query was accessed to determine which queries' search results to cache. For example, one could select the top N most popular queries from the set and cache the search results of the most recently accessed queries in that subset, or simply cache the search results of the N most popular queries. In another embodiment, results could be cached for those queries which most closely match a predefined matching metric (which might also take into account a user profile).

[0055] Once the queries are selected, they are checked against current cache entries (stage 806 and stage 808). In an exemplary embodiment, a cache exists in the query server 310 and in the cache server 338. One of ordinary skill would recognize that a single cache or multiple caches may be used in other embodiments. Referring to both FIGS. 3 and 7, checks are initially made against the predicted query results cache 328 in the query server to determine if the results for the query are present (stage 806). If the query results are not in the predicted query results cache 328, then the cached results 342 in the cache 338 are checked (stage 808) to determine whether the query search results are present. If the query search results are not present, then the search results for that query are obtained by using, for example, the query processing controller 314 and searching the document database (stage 810). The search results may be stored in one or more caches for later access (stage 812). If there are more queries to check against the caches (stage 814), then cache checking continues (stage 806). If the search results for a

query are found in a cache at stage 806 or stage 808, then it would need to be determined whether any additional queries needed to be checked (stage 814).

[0056] In some embodiments, it may be that certain of the predicted search results are returned to the client and presented to user during processing of the partial query in the hopes that included within the predicted search results is the desired search result. In such a situation, the user would be receiving a desired search result even more quickly than described earlier because the user would not need to spend the time in choosing one of the predicted queries. In any event, it makes it more likely that the desired search results are presented to the user prior to completion of the final query.

[0057] Such embodiments may be represented, for example, by stages 816 through 820 of FIG. 8. At stage 816, it is determined which search results to return to the client system 120. This may be the same or a subset of the search results for the queries selected for caching and may be based on any number of factors previously described, for example, most recently accessed, most frequently requested, most closely matching a predefined matching metric, or any combination thereof. Those of ordinary skill in the art would recognize many different factors. Similarly, stage 818 determines an order for the selected searches. The ordering of the searched may also be based on any number of factors including those previously discussed. Finally, at stage 820 the search results are sent to the client, and subsequently presented to the user.

[0058] Although the stages are shown in a particular order, one of ordinary skill would recognize that they need not necessarily be performed in that order. For example, checking and obtaining search results could be overlapped with identifying, ordering, and sending search results to the client.

[0059] In another embodiment, the anticipated search results that are cached are based on individual query terms, including anticipated search terms, and are not based on executions of anticipated multi-term queries. As more query terms are entered, predicted search results are generated based on the search results generated for the individual query terms. For example, this might be done by determining which documents in the (previously generated) respective search result sets fully satisfy the query entered by the user, and for each of those qualifying documents, combining query scoring values (e.g., information retrieval scores) associated with the individual query terms to produce a combined score for each qualifying document that is then used for ordering search results. If the number of qualifying documents is less than a threshold value, the full query is executed by the search engine. FIG. 9 provides an example of this embodiment.

[0060] Initially the first search term is obtained from the partial query (stage 902). This search term may be identified in any number of ways, including those described earlier regarding identification of a partial query, such as monitoring when a space or other delimiting character is entered by the user. It is determined whether search results corresponding to the search term are resident in a cache (stage 904). For example, it may be determined whether the search results for the search term are currently in the predicted query results cache 328. Although described as using the predicted query

results cache **328**, one of ordinary skill in the art would recognize that another embodiment could use the cached results **342** alone or in combination with the predicted query results cache **328**.

[**0061**] If the search results for that term are not in cache, then the search results are obtained and cached (stage **906**). If the search results for that term are in the cache, then flow bypasses stage **906** and proceeds to the combing stage **908**.

[**0062**] If only one search term has been identified and the cache has been checked for the corresponding search results, then the next search term in the partial search query, if any, should be obtained (stage **910**). If however, a preceding search term has been addressed, the results of the two search terms are combined to yield a set of search results that is more closely aligned with the multiple terms (stage **908**). One way to do this might be to identify which of the documents in the set of results for the first search term are also in the set of results for the second search term to create a combined search results set. This set could be ordered by, for example, combining query scoring values associated with the individual search terms to create a combined score for each document in the set. This process could continue until all of the search terms are included in the combined set. Using this approach might let a search result set be built incrementally while the user is entering the full query. By the time the user has indicated that the query is complete, the combined results might be more quickly available to the user than if the full search was obtained using the document database directly with no prior processing.

[**0063**] In another embodiment, the combined search results are transmitted to the client as each additional search term is encountered (stage **912**). This enhancement is shown by the dotted lines to and from stage **912**. Accordingly, as the user adds additional search terms, the user may be able to identify a desired result without having to enter all the intended search terms.

[**0064**] An additional embodiment would account for partial terms and is shown via dotted lines in **FIG. 9**. In this embodiment, flow may proceed from stage **902** to stage **914** for identification of partial terms. If a partial term was received (stage **914**) then the techniques described above for predicting search terms from partial search terms would be used to generate a set of predicted search terms (stage **916**). From this set of predicted search terms, certain ones could be identified for caching the associated search results (stage **918**). Techniques for selecting which predicted search terms to cache may be the same or similar to those described earlier regarding which queries to cache.

[**0065**] Referring to **FIG. 10**, an embodiment of a search engine **1000** that implements the methods described above includes one or more processing units (CPU's) **1002**, one or more network or other communications interfaces **1004**, memory **1006**, and one or more communication buses **1008** for interconnecting these components. The search engine **1000** may optionally include a user interface **1010** comprising a display device **1012** (e.g., for displaying system status information) and/or a keyboard **1014** (e.g., for entering commands). Memory **1006** may include high speed random access memory and may also include non-volatile memory, such as one or more magnetic or optical storage disks. Memory **1006** may include mass storage that is remotely located from CPU's **1002**. The memory **1006** may store:

[**0066**] an operating system **1016** that include procedures for handling various basic system services and for performing hardware dependent tasks;

[**0067**] a network communication module (or instructions) **1018** that is used for connecting the search engine **1000** to other computers via the one or more communications network interfaces **1004** (wired or wireless), such as the Internet, other wide area networks, local area networks, metropolitan area networks, and so on;

[**0068**] a query processor **310** for receiving full or partial queries and returning search results and predicted queries and predicted search results;

[**0069**] a query log module **324** for storing information about previously submitted queries in a query memory **1024**, which includes information for a plurality of previous queries **1026** such as the last time the query was submitted or how frequently the query is submitted over time;

[**0070**] a cache processor **1028** for storing the search results of previously submitted queries in a predicted query results cache **328**, the cache for storing a plurality of documents **1030**, and a cache index **326** used to assist in determining locations of documents in the cache; and

[**0071**] an auto-complete server **312** for receiving a partial query and returning predicted search terms or queries.

[**0072**] In some embodiments, the query processor **310** does not include a cache processor **1028**. In some embodiments, the query processor **310** does not include a query log module **324**.

[**0073**] In some embodiments, the query processor **310** includes: a client communications module **316** for receiving and transmitting information; a query receipt, processing and response module **318** for receiving and responding to full search queries; a partial query receipt, processing and response module **320** for receiving and responding to full search queries; a user information and processing module **322** for accessing user information from a user information database **1020**, which includes a respective user profile **1022** for a plurality of users. In some embodiments, the query processor **310** includes a subset of these modules. In some embodiments, the query processor **310** includes additional modules.

[**0074**] In some embodiments, the auto-complete server **312** includes a dictionary building module **1032** for creating various dictionaries from previously submitted queries, a comparison module **1034** to compare partial search terms or queries to the entries in various dictionaries, a dictionary selector **1036** to select one or more dictionaries for the comparison module to use, and one or more dictionaries **330** including entries of previously submitted search terms and/or queries.

[**0075**] **FIG. 10** depicts the internal structure of a search engine **1000** in one embodiment. It should be understood that in some other embodiments the search engine **1000** may be implemented using multiple servers so as to improve its throughput and reliability, or cost, other factors. For instance the auto-complete server **312** could be implemented on a

distinct server that communicates with and works in conjunction with other ones of the servers in the search engine **1000**.

[**0076**] Referring to **FIG. 11**, an embodiment of a client system **1100** that implements the methods described above includes one or more processing units (CPU's) **1102**, one or more network or other communications interfaces **1104**, memory **1106**, and one or more communication buses **1108** for interconnecting these components. The search engine **1100** may optionally include a user interface **1110** comprising a display device **1112** and/or a keyboard **1114**. Memory **1106** may include high speed random access memory and may also include non-volatile memory, such as one or more magnetic or optical storage disks. Memory **1106** may include mass storage that is remotely located from CPU's **1102**. The memory **1006** may store:

[**0077**] an operating system **1116** that include procedures for handling various basic system services and for performing hardware dependent tasks;

[**0078**] a network communication module (or instructions) **1118** that is used for connecting the client system **1100** to other computers via the one or more communications network interfaces **1104** (wired or wireless), such as the Internet, other wide area networks, local area networks, metropolitan area networks, and so on; and

[**0079**] a browser/tool **1119** for interfacing with a user to input search queries, and for displaying search results; and

[**0080**] a search assistant **1120**.

[**0081**] In some embodiments, the search assistant **1120** is separate from the browser/tool **1119**, while in other embodiments the search assistant is incorporated in the browser/tool **1119**.

[**0082**] The search assistant **1120** includes: an entry and selection monitoring module **1121** for monitoring the entry of search queries and selecting partial queries and final queries for transmission to the search engine; a transmission module **1122** for transmitting the partial search queries to the search engine; a predicted query receipt module **1126** for receiving predicted queries; a predicted search results receipt module **1128** for receiving predicted search results; and optionally, a search results receipt module **1128** for receiving search results. The transmission of final (i.e., completed) queries, receiving search results for completed queries, and displaying such results may be handled by the browser/tool **1119**, the search assistant **1120**, or a combination thereof.

[**0083**] The foregoing description, for purpose of explanation, has been described with reference to specific embodiments. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the invention and its practical applications, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated.

What is claimed is:

1. A computer-implemented method for processing a search query, comprising:

receiving a portion of the search query from a search requester; predicting from the portion of the search query a set of predicted queries; and transmitting the set of predicted queries to the search requester.

2. The method of claim 1, further comprising:

predicting the set of predicted queries using a collection of queries from a community of users.

3. The method of claim 1, further comprising:

comparing at least one or more characters of the portion of the search query to entries in a dictionary; and creating the set of predicted queries from one or more of the entries in the dictionary whose entries include the at least one or more characters.

4. The method of claim 3, wherein the set of predicted queries is created from entries in the dictionary whose entries begin with the at least one or more characters.

5. The method of claim 3, wherein the portion of the search query includes a complete search term of the search query.

6. The method of claim 1, further comprising:

transmitting one or more predicted search results to the search requester.

7. The method of claim 6, further comprising:

obtaining search results corresponding to one or more of the set of predicted queries; and caching the search results.

8. The method of claim 7, further comprising:

transmitting one or more of the cached results to the search requester.

9. The method of claim 7, further comprising:

selecting for which of the one or more of the set of predicted queries to obtain the search results based on predefined query selection criteria.

10. The method of claim 9, further comprising:

selecting for which of the one or more of the set of predicted queries to obtain the search results based on popularity of the one or more of the set of predicted queries.

11. The method of claim 9, further comprising:

selecting for which of the one or more of the set of predicted queries to obtain the search results based on information about the search requester.

12. The method of claim 3, further comprising:

selecting the dictionary from a set of one or more dictionaries.

13. The method of claim 12, further comprising:

selecting the dictionary based on search requester information.

14. The method of claim 12, further comprising:

creating the dictionary based on a collection of queries from multiple search requestors.

15. A computer-implemented method for processing a search query, comprising:

creating a dictionary from a community of users; receiving a portion of the search query from a search

requester; identifying entries in the dictionary which match the portion of the search query; selecting one or more of the matching entries in the dictionary; ordering the one or more matching entries to create a set of predicted queries; transmitting the set of predicted queries to the search requestor; obtaining search results for at least one of the predicted queries; caching the search results; and transmitting at least a portion of the search results to the search requestor.

16. A search query processing system, comprising:

one or more central processing units for executing programs; and

a server executable by the one or more central processing units, the server comprising:

instructions for receiving a portion of the search query from a search requestor;

instructions for predicting from the portion of the search query a set of predicted queries; and

instructions for transmitting the set of predicted queries to the search requestor.

17. The system of claim 16, wherein the predicting instructions further comprise:

instructions for predicting the set of predicted queries using a collection of queries from a community of users.

18. The system of claim 16, wherein the instructions for predicting further comprise:

instructions for comparing at least one or more characters of the portion of the search query to entries in a dictionary; and

instructions for creating the set of predicted queries from one or more of the entries in the dictionary whose entries include the at least one or more characters.

19. The system of claim 18, wherein the instructions for predicting further comprise instructions for predicting the set of predicted queries is created from entries in the dictionary whose entries begin with the at least one or more characters.

20. The system of claim 18, wherein the one or more characters of the portion are search terms.

21. The system of claim 16, wherein the server further comprises:

instructions for transmitting one or more predicted search results to the search requestor.

22. The system of claim 21, wherein the server further comprises:

instructions for obtaining search results corresponding to one or more of the set of predicted queries; and

instructions for caching the search results.

23. The system of claim 22, wherein the server further comprises:

instructions for transmitting one or more of the cached results.

24. The system of claim 18, wherein the server further comprises:

instructions for selecting the dictionary from a set of one or more dictionaries.

25. The system of claim 24, wherein the instructions for selecting the dictionary further comprise:

instructions for selecting the dictionary based on search requestor information.

26. The system of claim 18, wherein the server further comprises:

instructions for creating the dictionary based on a collection of queries from multiple search requestors.

27. The system of claim 22, wherein the instructions for obtaining search results further comprise:

instructions for selecting for which of the one or more of the set of predicted queries to obtain the search results based on predefined query selection criteria.

28. The system of claim 27, wherein the instructions for selecting for which of the one or more of the set of predicted queries further comprise:

instructions for selecting for which of the one or more of the set of predicted queries to obtain the search results based on popularity of the one or more of the set of predicted queries.

29. The system of claim 27, wherein the instructions for selecting for which of the one or more of the set of predicted queries further comprise:

instructions for selecting for which of the one or more of the set of predicted queries to obtain the search results based on information about the search requestor.

30. A computer program product, for use in conjunction with a computer system, for processing a search query, the computer program product comprising:

instructions for receiving a portion of the search query from a search requestor;

instructions for predicting from the portion of the search query a set of predicted queries; and

instructions for transmitting the set of predicted queries to the search requestor.

31. The computer program product of claim 30, further comprising:

instructions for predicting the set of predicted queries using a collection of queries from a community of users.

32. The computer program product of claim 30, further comprising,

instructions for comparing at least one or more characters of the portion of the search query to entries in a dictionary; and

instructions for creating the set of predicted queries from one or more of the entries in the dictionary whose entries include the at least one or more characters.

33. The computer program product of claim 32, wherein the set of predicted queries is created from entries in the dictionary whose entries begin with the at least one or more characters.

34. The computer program product of claim 32, wherein the one or more characters of the portion are search terms.

35. The computer program product of claim 30, further comprising:

instructions for transmitting one or more predicted search results to the search requestor.

36. The computer program product of claim 35, further comprising:

instructions for obtaining search results corresponding to one or more of the set of predicted queries; and

instructions for caching the search results.

37. The computer program product of claim 36, further comprising:

instructions for transmitting one or more of the cached results.

38. The computer program product of claim 32, further comprising:

instructions for selecting the dictionary from a set of one or more dictionaries.

39. The computer program product of claim 38, further comprising:

instructions for selecting the dictionary based on search requestor information.

40. The computer program product of claim 32, further comprising:

instructions for creating the dictionary based on a collection of queries from multiple search requestors.

41. The computer program product of claim 36, further comprising:

instructions for selecting for which of the one or more of the set of predicted queries to obtain the search results based on predefined query selection criteria.

42. The computer program product of claim 41, further comprising:

instructions for selecting for which of the one or more of the set of predicted queries to obtain the search results based on popularity of the one or more of the set of predicted queries.

43. The computer program product of claim 41, further comprising:

instructions for selecting for which of the one or more of the set of predicted queries to obtain the search results based on information about the search requestor.

44. A method for processing a search query, comprising:

monitoring a user's entry of a search query;

transmitting a portion of the search query prior to the user's completion of the search query;

receiving a set of predicted queries based on the portion of the search query; and

presenting one or more of the set of predicted queries to the user.

45. The method of claim 44, further comprising:

monitoring a user's selection of one of the set of predicted queries; and

transmitting the user's selection to a query server.

46. The method of claim 44, further comprising:

receiving one or more predicted search results; and

presenting the one or more predicted search results to the user.

47. The method of claim 44, further comprising:

transmitting the portion of the search query after a predefined number of characters are entered.

48. The method of claim 44, further comprising:

transmitting the portion of the search query after a period of at least a predefined duration during which no characters are entered.

49. The method of claim 44, further comprising:

transmitting the portion of the search query when a delimiting character is entered.

50. The method of claim 46, further comprising:

ordering the one or more predicted search results; and

presenting the one or more predicted search results to the user based on the ordering.

51. The method of claim 50, further comprising:

ordering the one or more predicted search results based on user information.

52. A system for processing a search query, comprising:

one or more central processing units for executing programs; and a client executable by the one or more central processing units, the client comprising:

instructions for monitoring a user's entry of a search query;

instructions for transmitting a portion of the search query prior to the user's completion of the search query;

instructions for receiving a set of predicted queries based on the portion of the search query; and

instructions for presenting one or more of the set of predicted queries to the user.

53. The system of claim 52, wherein the client further comprises:

instructions for monitoring a user's selection of one of the set of predicted queries; and

instructions for transmitting the user's selection to a query server.

54. The method of claim 52, wherein the client further comprises:

instructions for receiving one or more predicted search results; and

instructions for presenting the one or more predicted search results to the user.

55. The method of claim 52, wherein the instructions for transmitting a portion of the search query further comprise:

instructions for transmitting the portion of the search query after a predefined number of characters are entered.

56. The method of claim 52, wherein the instructions for transmitting a portion of the search query further comprise:

instructions for transmitting the portion of the search query after a period of at least a predefined duration during which no characters are entered.

57. The method of claim 52, wherein the instructions for transmitting a portion of the search query further comprise:

instructions for transmitting the portion of the search query when a delimiting character is entered.

58. The method of claim 54, wherein the client further comprises:

instructions for ordering the one or more predicted search results; and

instructions for presenting the one or more predicted search results to the user based on the ordering.

59. The method of claim 58, wherein the instructions for ordering the one or more predicted results further comprise:

instructions for ordering the one or more predicted search results based on user information.

60. A computer program product, for use in conjunction with a computer system, for processing a search query, the computer program product comprising:

instructions for monitoring a user's entry of a search query;

instructions for transmitting a portion of the search query prior to the user's completion of the search query;

instructions for receiving a set of predicted queries based on the portion of the search query; and

instructions for presenting one or more of the set of predicted queries to the user.

61. The computer program product of claim 60, further comprising:

instructions for monitoring a user's selection of one of the set of predicted queries; and

instructions for transmitting the user's selection to a query server.

62. The computer program product of claim 60, further comprising:

instructions for receiving one or more predicted search results; and

instructions for presenting the one or more predicted search results to the user.

63. The computer program product of claim 60, further comprising:

instructions for transmitting the portion of the search query after a predefined number of characters are entered.

64. The computer program product of claim 60, further comprising:

instructions for transmitting the portion of the search query after a period of at least a predefined duration during which no characters are entered.

65. The computer program product of claim 60, further comprising:

instructions for transmitting the portion of the search query when a delimiting character is entered.

66. The computer program product of claim 62, further comprising:

instructions for ordering the one or more predicted search results; and

instructions for presenting the one or more predicted search results to the user based on the ordering.

67. The computer program product of claim 66, further comprising:

instructions for ordering the one or more predicted search results based on user information.

68. A method for processing a search query, comprising:

receiving a portion of the search query at a query server from a client;

predicting a set of queries based on the transmitted portion of the search query;

obtaining search results corresponding to one or more of the predicted set of queries; and

caching the search results.

69. The method of claim 68, further comprising:

transmitting one or more of the cached searched results from the query server to the client.

70. The method of claim 68, further comprising:

upon completion of user entry of the search query, and upon determining that the completed search query matches one of the cached search results, transmitting the matching search results from the query server to the client.

71. A method of processing a search request, comprising:

receiving a portion of a search query;

extracting one or more search terms from the portion of the search query;

obtaining individual search results for each of the one or more search terms;

caching the individual search results;

combining the individual search results to produce final search results;

transmitting the final search results to a search requestor.

72. A method of processing a search request, comprising:

receiving portions of a search query;

extracting one or more search term from each portion of the search query;

obtaining individual search results for each of the one or more search terms;

caching the individual search results;

combining the individual search results to produce final search results;

transmitting the final search results to a search requester.

73. A method for processing a search query, comprising:

receiving a first portion of the search query at a query server from a client;

extracting a first search term from the first portion of the search query;

predicting at least one first query based on the first search term;

obtaining first search results for the at least one first query;

caching one or more of the first search results;

receiving a second portion of the search query at the query server from the client;

extracting a second search term from the second portion of the search query;

predicting at least one second query based on the second search term;

obtaining second search results for the at least one second query;

caching one or more of the second search results;

combining the first search results and second search results to produce final search results; and

transmitting the final search results to the client.

* * * * *