



(22) Date de dépôt/Filing Date: 2008/04/03

(41) Mise à la disp. pub./Open to Public Insp.: 2008/11/16

(45) Date de délivrance/Issue Date: 2013/04/02

(30) Priorité/Priority: 2007/05/16 (US11/804,030)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01),
G06F 9/44 (2006.01)


(72) Inventeurs/Inventors:
GRECHANIK, MARK, US;
CONROY, KEVIN MICHAEL, US;
PROBST, KATHARINA ANDREA, US

(73) Propriétaire/Owner:
ACCENTURE GLOBAL SERVICES LIMITED, IE

(74) Agent: SMART & BIGGAR


(54) Titre : OUTIL DE RECHERCHE D'APPLICATIONS POUR LA MISE AU POINT RAPIDE DE PROTOTYPES ET DE NOUVELLES APPLICATIONS

(54) Title: APPLICATION SEARCH TOOL FOR RAPID PROTOTYPING AND DEVELOPMENT OF NEW APPLICATIONS



(57) Abrégé/Abstract:

A code search tool greatly reduces time, cost, and other resource expenditures associated with implementing a new application. The tool is a search, navigation and visualization tool that accepts high-level processing concepts as inputs to identify, rank, and



(57) **Abrégé(suite)/Abstract(continued):**

return the code of relevant existing applications. A software developer may use the relevant applications to rapidly build prototypes, identify requirements, and develop new applications. The tool provides an efficient way to improve the reuse of application logic to realize the high-level processing concepts, and more efficiently deliver proof of concept.

ABSTRACT

A code search tool greatly reduces time, cost, and other resource expenditures associated with implementing a new application. The tool is a search, navigation and visualization tool that accepts high-level processing concepts as inputs to identify, rank, and return the code of relevant existing applications. A software developer may use the relevant applications to rapidly build prototypes, identify requirements, and develop new applications. The tool provides an efficient way to improve the reuse of application logic to realize the high-level processing concepts, and more efficiently deliver proof of concept.

**APPLICATION SEARCH TOOL FOR RAPID
PROTOTYPING AND DEVELOPMENT OF NEW APPLICATIONS**

INVENTORS:

**Mark Grechanik
Kevin Michael Conroy
Katharina Andrea Probst**

BACKGROUND OF THE INVENTION

1. Technical Field.

[001] This disclosure concerns finding existing program logic and reusing it to rapidly build prototypes and develop new applications. In particular, this disclosure relates to a search, navigation and visualization tool that accepts high-level processing concepts as inputs that drive a multi-layered search to identify applications and application programming interface (API) calls for reuse.

2. Background Information.

[002] Software professionals widely recognize logic (e.g., source code) reuse as a technique that reduces the time, money, and other costs associated with creating a new application. Software professionals recognize API calls as forms of abstraction for high-level processing concepts, which drives the wide acceptance of API calls as reusable logic. For example, implementing an existing API call that produces a pull-down menu eliminates the need to write all the underlining logic necessary to deliver the functionality of a pull-down menu. However, current logic mining techniques and mining tools fail to retrieve highly relevant software components from application repositories that developers can use to prototype requirements in support of high-level processing concepts. Modern search engines do not ensure that applications identified by the search engines can serve as highly relevant application prototypes (HRAPs). Software professionals consider the mismatch between the high-level processing concepts (e.g., the intent reflected in the descriptions of applications) and low-level implementation details (e.g., API calls and actual run-time behaviour) found in application logic a fundamental technical challenge to identifying highly relevant

applications (HRAs). Software professionals intend to author meaningful descriptions of applications, in the course of depositing applications into software repositories. The mismatch between the description of an application and the actual behaviour of the application represents one example of the "vocabulary problem", which states that no single word or phrase best describes a programming concept.

[003] In the spiral model of software development, stakeholders describe high-level processing concepts to development teams, and together the stakeholders and development teams identify requirements in support of the high-level processing concepts. In addition, a development team builds a prototype based on the requirements, and the development team demonstrates the prototype to the stakeholders to receive feedback. Prototypes attempt to approximate the desired high-level processing concepts (e.g., features and capabilities) of the new application stakeholders desire development teams to build. The feedback from stakeholders often leads to changes to the prototype and the original requirements, as stakeholders iteratively refine their vision. In the event the stakeholders make a substantial number of changes to the requirements, the development team often discards the prototype and builds a new prototype, and another iteration of refinements repeats. Building prototypes repeatedly without reusing existing application logic costs organizations a great deal in the form of wasted project resources and time.

[004] Development teams find it cost-effective to identify existing applications that approximate the high-level processing concepts and requirements of new software projects as the basis for prototypes. In the context of prototyping, software development professionals consider such existing applications as HRAs. Many application repositories (e.g., open source repositories and source control management systems maintained by stakeholders internally) contain hundreds of thousands of different existing applications (e.g., potential HRAs). Unfortunately, developers find it difficult to identify applications (e.g., HRAs) ideal for prototyping because of the time and expense involved in searching (e.g., querying) application repositories and source control management systems.

[005] The amount of intellectual effort that a developer must expend to move a software system from one stage of development to another may be considered the

"cognitive distance." For example, using current search tools developers expend significant intellectual effort to identify potentially relevant applications and confirm HRAs from potentially relevant applications. Many developers employ search engines that identify exact matches between keywords and the words found in application repositories. The application repositories may include descriptions, application logic comments, program variables names, and variable types of existing applications. Such search engines actually increase the difficulty of identifying HRAs, because of the poor quality of information contained in application repositories, and the inability to reduce the cognitive distance required to identify HRAs, as well as other factors. Additionally, many application repositories include incomplete, misleading and inaccurate descriptions of applications identified in the application repositories. Consequently, even matching keywords with words in the application descriptions found in application repositories does not guarantee that the search engine will identify HRAs.

[006] Effective software reuse techniques (e.g., prototyping using existing applications) reduce the cognitive distance between the initial concept of a system (e.g., high-level processing concepts that expressly and implicitly describe the features and capabilities of a new application), establishing discrete requirements, and the production implementation of the new system. Unfortunately, current search engines lack the ability to reduce the cognitive distance related to identifying HRAs.

[007] For example, an application description may indicate that an application includes an encryption feature when in fact the application uses compression as a crude form of encryption. A developer entering "encryption" (e.g., as a high-level processing concept and specific requirement) as a keyword may waste precious time to review a search engine result containing the incorrectly described application, and ultimately discard the result, because the application fails to meet the encryption requirement. The developer must download the application identified in the search result, locate and examine fragments of the application logic that allegedly implements encryption before determining that the application fails to meet the requirement. The developer may spend scarce project development budget resources and significant amount of time to analyze the application before determining that an application is not relevant. The developer may even observe the

runtime behavior of the application to ensure that the behavior matches the high-level processing concepts desired by the stakeholders, and meets the requirements in support of the high-level processing concepts before establishing that the application qualifies as a HRA. Current search engines also lack the ability to assist developers to

5 rapidly identify requirements in support of high-level processing concepts described by stakeholders.

[008] Some search tools return code snippets (e.g., segments of application logic), however, code snippets do not give enough background or context to assist developers to create rapid prototypes, and such search tools require developers to

10 invest significant intellectual effort (e.g., cognitive distance) to understand how to use the code snippets in broader scopes. Other existing approaches and tools retrieve snippets of code based on the context of the application logic that developers work on, but while these approaches and tools improve the productivity of developers, they do not return relevant applications from high-level processing concepts as inputs.

15 **[009]** A need has long existed for a system and method that efficiently identifies HRAs usable to rapidly build prototypes and develop new applications.

SUMMARY

[010] In some embodiments, the EXEcutable exaMPLes ARchive system

20 (Exemplar) rapidly and efficiently identifies highly relevant applications (HRAs) from large application repositories. In these embodiments, using Exemplar, a developer enters high-level processing concepts (e.g., toolbar, download, smart card) as input (e.g., initial query keywords), and Exemplar uses information retrieval and program analysis techniques to retrieve HRAs that implement the high-level processing

25 concepts. Exemplar may also accept various types of inputs that describe high-level processing concepts (e.g. concept text identifiers, concept visual identifiers, concept audio identifiers, and any other sensory identifier usable to identify high-level processing concepts). Exemplar may use the help pages and help documentation of third-party libraries, software development kits, and other middleware to produce a list

30 of names of API calls that Exemplar in turn uses to expand an initial query ("query expansion") to identify the application logic of HRAs and the API calls included in the

HRAs. Exemplar may determine the behavior of the application logic and API call logic, and ranks the HRAs and API calls included in the HRAs.

[011] Exemplar may use help documentation or other trusted sources that describe API calls to expand queries. An application provider typically provides the help pages and help documentation for their applications, which developers consider reliable and a trusted source. In particular, developers may consider application providers trusted sources for help pages and help documentation of popular and widely used applications written by large development teams, produced under rigorous testing and development best practices, and used by other developers who provide feedback regarding documentation using different forums (e.g., user groups). Developers may trust help documentation over the descriptions of applications included in application repositories, because application providers generally produce more verbose and accurate help documentation than the descriptions of applications included in application repositories. Developers may also trust help documentation because many different people and review procedures are typically used to produce help documentation.

[012] Exemplar query expansion may increase the probability of identifying logic matches usable to build highly relevant application prototypes (HRAPs) and new applications, and may address the vocabulary problem mentioned above by expanding an initial query to include new keywords, metadata, and semantics information found in help pages and other help documentation determined to have similar meanings to the keywords originally used by a developer in the initial query. Exemplar may expand an initial query to include the names of API calls with semantics that reflect (in many cases unequivocally) specific behaviour of the matched applications. Exemplar may locate application logic containing the API calls that exhibit desired semantics by identifying API calls through help pages and help documentation. Exemplar may provide a user interface that developers can use to navigate directly to the various locations to determine how an HRA implements high-level processing concepts.

[013] Exemplar may rank HRAs according to the number of high-level processing concepts implemented by each API call found in the HRAs, or based on other ranking metrics. In other words, since API calls implement high-level processing concepts, the

more high-level processing concepts implemented by an HRA the more relevant the HRA and the higher the rank assigned to the HRA. Exemplar may consider keywords included in queries to represent logically connected concepts. Often a question structured as a sentence forms the basis for a query, from which a developer extracts
5 keywords to form the query. For example, consider the query "send receive secure XML." Where a query presents a relation between multiple concepts (e.g., send secure XML), then a relation should exist between API calls that implement the concepts in the corresponding application logic (e.g., API calls that encrypt, process or handle XML formatted content, and transmit content). Application logic often preserves the
10 relations between concepts (e.g., control flow and data flow links), an instance of the software reflection model concept and known as connectivity heuristics. Exemplar may calculate HRAs rankings based on analyzing the connectivity heuristics of API calls that implement the concepts included in the queries. Exemplar may use program analysis algorithms, and computes control flow graphs (CFG), and data flow graphs
15 (DFG) to analyze the connectivity heuristics of API calls.

In one aspect of the invention, there is provided a method comprising: accessing a first search result comprising at least one application programming interface (API) call identifier, wherein said first search result is associated with a first search performed using a first query, and wherein said first search is performed on a
20 first database; combining said first query with said at least one API call identifier to form a second query; and performing a second search using said second query to obtain a second search result comprising at least one application identifier of at least one application, wherein said second search is performed on a second database.

In another aspect of the invention, there is provided a non-transitory
25 computer-readable medium having computer-readable program code embodied therein for causing a computer system to perform a method comprising: accessing a first search result comprising at least one application programming interface (API) call identifier, wherein said first search result is associated with a first search performed using a first query, and wherein said first search is performed on a first database;
30 combining said first query with said at least one API call identifier to form a second

query; and performing a second search using said second query to obtain a second search result comprising at least one application identifier of at least one application, wherein said second search is performed on a second database.

5 In a further aspect of the invention, there is provided a system comprising a processor and a memory, wherein said memory comprises instructions that when executed by said processor implement a method comprising: accessing a first search result comprising at least one application programming interface (API) call identifier, wherein said first search result is associated with a first search performed using a first query, and wherein said first search is performed on a first database; combining said
10 first query with said at least one API call identifier to form a second query; and performing a second search using said second query to obtain a second search result comprising at least one application identifier of at least one application, wherein said second search is performed on a second database.

[014] Other systems, methods, and features of the invention will be, or will
15 become, apparent to one with skill in the art upon examination of the following figures and detailed description. It is intended that all such additional systems, methods, features and advantages be included within this description, be within the scope of the invention, and be protected by the following claims.

20 **BRIEF DESCRIPTION OF THE DRAWINGS**

[015] The disclosure can be better understood with reference to the following drawings and description. The components in the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention. Moreover, in the figures, like referenced numerals designate corresponding parts or
25 elements throughout the different views.

[016] Figure 1 illustrates the Exemplar system configuration.

[017] Figure 2 shows an Exemplar data flow diagram.

[018] Figure 3 illustrates an Exemplar query interface that may be used to input an original query.

[019] Figure 4 shows an Exemplar system prototyping and application development (EPAD) project area.

[020] Figure 5 shows a more detailed view of the Exemplar system.

[021] Figure 6 shows other features of the Exemplar system.

[022] Figure 7 shows API graphs for two different potentially relevant applications.

[023] Figure 8 shows the acts that the Exemplar system may take to obtain a relevant application list.

[024] Figure 9 shows the acts that the Exemplar system heuristic ranking engine may take to assign an application heuristic relevance ranking to a potentially relevant application.

[025] Figure 10 shows the processing that the selection logic and application creation logic may take to generate a new application.

DETAILED DESCRIPTION

[026] The EXEcutable exaMPLeS ARchive system (Exemplar) solves the technical problem of providing a tool that accepts high-level processing concepts as queries to identify, determine the behavior, rank and return the application logic of HRAs. Exemplar solves an instance of the difficult vocabulary problem that exists when users and developers describe processing concept with different words. Exemplar is not limited to basic keyword matching used in queries against application descriptions and comments included with application logic. Accordingly, when an application is highly relevant, and where a query contains keywords different from the words used by the developer to describe application logic and API call logic, Exemplar nevertheless returns the application as a highly relevant application.

[027] Exemplar matches high-level processing concepts (e.g., expressed using keywords) with the descriptions of various API calls found in help documents or other trusted descriptive sources. Because a typical application invokes API calls from several different libraries, several different people who use different vocabularies often author help documents associated with API calls. The richness of different

vocabularies increases the probability of finding matches and producing a long list of potentially relevant applications and API calls. Searching help documents or other trusted descriptive sources produces additional benefits. For example, help documents including an API call often indicate where the application logic implements the API call. Consequently, Exemplar may direct a developer to the location in application logic where an API call implements a high-level processing concept. The developer may then determine the relevance of the application logic and API call logic. In other words, the developer may determine whether the application logic and API call logic actually support the high-level processing concept.

[028] Although specific components of Exemplar will be described, methods, systems, and articles of manufacture consistent with Exemplar may include additional or different components. For example, a processor may be implemented as a microprocessor, microcontroller, application specific integrated circuit (ASIC), discrete logic, or a combination of other type of circuits or logic. Similarly, memories may be DRAM, SRAM, Flash or any other type of memory. Logic that implements the processing and programs described below may be stored (e.g., as computer executable instructions) on a computer readable medium such as an optical or magnetic disk or other memory. Alternatively or additionally, the logic may be realized in an electromagnetic or optical signal that may be transmitted between entities. An example of such a signal is a physical layer Ethernet signal bearing TCP/IP packets that include program source code or executable programs. Flags, data, databases, tables, and other data structures may be separately stored and managed, may be incorporated into a single memory or database, may be distributed, or may be logically and physically organized in many different ways. Programs may be parts of a single program, separate programs, or distributed across several memories and processors. Furthermore, the programs, or any portion of the programs, may instead be implemented in hardware.

[029] Figure 1 illustrates the Exemplar system environment 100 ("Exemplar environment 100). The Exemplar environment 100 may include an Exemplar prototyping and application development (EPAD) user interface 102, a help content processor 104, and help content 106. The Exemplar environment 100 also includes

an API calls dictionary 108, expanded search engine 110, logic repository 112, heuristic relevance ranking engine 114, and a logic analyzer 116. Any or all of the elements shown in Figure 1 may be co-located or distributed and in communication over one or more networks 118 (e.g., the Internet).

[030] In one implementation, the EPAD user interface 102, expanded search engine 110, heuristic relevance ranking engine 114 and logic analyzer 116 form an Exemplar system 124 within the Exemplar environment 100. The Exemplar system 124 may include additional or different components. The Exemplar system 124 may communicate with the help content processor 104, help content 106, API calls dictionary 108, and logic repository 112, as well as other systems, through the networks 118 (e.g., Internet) as external systems.

[031] The logic repository 112 may include application logic 120 and API call logic 122. The Exemplar system 124 accepts high-level processing concepts (e.g., "send secure XML") as input and produces output identifying which application logic 120 and API call logic 122 developers may use to prototype and develop new applications implementing the high-level processing concepts. In one implementation, the Exemplar environment 100 implements the help content 106 and the logic repository 112 with multiple storage devices (e.g., multiple databases on different disk drives), and interfaces to help content 106, application logic 120 and API call logic 122 from various available source (e.g., local or remote help databases, websites, knowledge exchanges, document repositories, or other sources).

[032] In one implementation, the help content processor 104 may be implemented as a web crawler that traverses available application repositories, and downloads help content 106 (e.g., application descriptions), and logic repository 112 content (e.g., application logic 120, and API logic 122). The help content processor 104 may also perform full text indexing on the help content 106 and the logic repository 112 content. The help content processor 104 may further produce an API calls dictionary 108 that includes sets of tuples (a form of ordered list) that link selected words from the descriptions of the API calls to the names of the API calls.

[033] The description above used the examples of application logic 120 and API call logic 122. These types of logic may be program source code (e.g., C or C++ code),

for example. However, the Exemplar environment 100 may search, analyze, and determine relevance for many other types of logic. As examples, the logic repository 112 may include programs or program components expressed in a visual programming language using graphical program elements and spatial arrangements of text and graphic symbols. The visual programming logic may include icon-based logic, form-based logic, diagram-based logic or other types of visual expression. The visual expression may be consistent with dataflow languages, flow-based programming, domain-specific modelling, or other programming paradigms.

[034] Figure 2 shows an Exemplar data flow diagram 200. Exemplar system 124 accepts a high-level processing concept as input to create an original query 202 that Exemplar system 124 may forward to the help content processor 104. The help content processor 104 may produce a basis API call list 204 from the API calls dictionary 108 by matching the words in the high-level processing concepts (e.g., "send secure XML") found in the original query 202 executed to search the help content 106.

[035] The expanded search engine 110 may combine the original query 202 and the basis API call list 204 to form an expanded query 206. The expanded search engine 110 may execute an expanded search using the expanded query 206 to search through the logic repository 112 to obtain an expanded search result 208. In one implementation, the logic repository 112 may return the expanded search results 208 to the heuristic relevance ranking engine 114. The expanded search result 208 may contain a list of potentially relevant applications 210 and potentially relevant API calls 212 that the heuristic relevance ranking engine 114 analyzes using the logic analyzer 116. In one implementation, the heuristic relevance ranking engine 114 may include the logic analyzer 116. The logic analyzer 116 may include a parser generator such as ANTLR ("ANother Tool for Language Recognition") available from www.antlr.org that provides support for generating data flow graphs and control flow graphs.

[036] The logic analyzer 116 may return connectivity rankings 214, discussed in detail below, to further determine an application heuristic relevance ranking 216 and an API call heuristic relevance ranking 218. The heuristic relevance ranking engine 114 may return the application heuristic relevance ranking 216 and an API call

heuristic relevance ranking 218 to the EPAD user interface 102. The expanded search engine 110 may also return a relevant applications list 220 and a relevant API calls list 222 to the EPAD user interface 102. The Exemplar system 124 may assign an application heuristic relevance ranking 216 to one or more relevant applications found in the relevant applications list 220 to indicate how closely each relevant application supports the high-level processing concept represented by the original query 202. Similarly, Exemplar system 124 may assign an API call heuristic relevance ranking 218 to one or more relevant API calls found in the relevant API call list 222 to indicate how closely each relevant API call supports the high-level processing concept represented by the original query 202.

[037] Figure 3 illustrates an Exemplar query interface 300 that may be used to input an original query 202. The original query 202 may represent a high-level processing concept such as "compress uncompress ZIP file," as shown in the text entry field 302. Several developers may have implemented the example high-level processing concept "compress uncompress ZIP file" 302 in different ways with various API calls described in the help content 106, API calls dictionary 108 and logic repository 112. A user may specify search refinement criteria 304 using interface elements such as a drop down box, menu or user input field. The search refinement criteria 304 may dictate the types of logic of interest (e.g., C, C++, JAVA, or other types of logic), may dictate the information sources searched (e.g., search only compiled Help files, or *.doc files), or may specify other search criteria. The Exemplar query interface 300 may include graphical user interface elements (e.g., the search button 306) used to execute the original query 202.

[038] Table 1 shows an example of an original query 202 in the form of a structured query language statement (SQL) that represents the high-level processing concept "compress uncompress ZIP file" 302. Table 1 shows that the original query 202 will search the help content 106 (e.g., Java Help Documents) to identify a basis API calls list 204.

Table 1 – original query 202


```

SELECT basis API Calls
FROM Java Help Documents
WHERE Words in these Documents =
        compress or uncompress or ZIP or file.

```

[039] Table 2 shows one example of the help content 106 represented by a fragment of Java Help Documentation released by Sun Microsystems, Inc. that describes the functionality of classes exported from the Java.util package. The Java.util package defines a number of classes, primarily collections classes that a developer may use when working with groups of objects. Referring to Table 2, the help content processor 104 may identify partial matches for the class ZipEntry to the original query 202. The help content processor 104 may search the help content 106 and identify a fragment of the help documentation for the ZipEntry class shown in Table 3.

Table 2 – help content 106 (e.g. a fragment of Java Help Document)	
Class	Summary
Adler32	A class that can be used to compute the Adler-32 checksum of a data stream.
CheckedInputStream	An input stream that also maintains a checksum of the data being read.
CheckedOutputStream	An output stream that also maintains a checksum of the data being written.
CRC32	A class that can be used to compute the CRC-32 of a data stream.
Deflater	This class provides support for general purpose compression using the popular ZLIB compression library.
DeflaterInputStream	Implements an input stream filter for compressing data in the "deflate" compression format.
DeflaterOutputStream	This class implements an output stream filter for

	compressing data in the "deflate" compression format.
GZIPInputStream	This class implements a stream filter for reading compressed data in the GZIP file format.
GZIPOutputStream	This class implements a stream filter for writing compressed data in the GZIP file format.
Inflater	This class provides support for general purpose decompression using the popular ZLIB compression library.
InflaterInputStream	This class implements a stream filter for uncompressing data in the "deflate" compression format.
InflaterOutputStream	Implements an output stream filter for uncompressing data stored in the "deflate" compression format.
ZipEntry	This class is used to represent a ZIP file entry.
ZipFile	This class is used to read entries from a zip file.
ZipInputStream	This class implements an input stream filter for reading files in the ZIP file format.
ZipOutputStream	This class implements an output stream filter for writing files in the ZIP file format.

[040] Table 3 shows the descriptions of two different methods (e.g., `getCompressedSize`, and `setMethod`) for the `ZipEntry` class that include the terms compress and uncompress found in the high-level processing concept "compress uncompress ZIP file" 302. The basis API call list 204 may include the `getCompressedSize` and `setMethod` methods.

Table 3 – help content 106 (e.g., fragment of help documentation for <code>ZipEntry</code> class)	
Method	Summary
Object <code>clone()</code>	Returns a copy of this entry.
String <code>getComment()</code>	

	Returns the comment string for the entry, or null if none.
long	getCompressedSize() Returns the size of the compressed entry data, or -1 if not known.
long	getCrc() Returns the CRC-32 checksum of the uncompressed entry data, or -1 if not known.
byte[]	getExtra() Returns the extra field data for the entry, or null if none.
int	getMethod() Returns the compression method of the entry, or -1 if not specified.
String	getName() Returns the name of the entry.
long	getSize() Returns the uncompressed size of the entry data, or -1 if not known.
long	getTime() Returns the modification time of the entry, or -1 if not specified.
int	hashCode() Returns the hash code value for this entry.
boolean	isDirectory() Returns true if this is a directory entry.
void	setComment(String comment) Sets the optional comment string for the entry.
void	setCompressedSize(long csize) Sets the size of the compressed entry data.
void	setCrc(long crc) Sets the CRC-32 checksum of the uncompressed entry data.
void	setExtra(byte[] extra) Sets the optional extra field data for the entry.
void	setMethod(int method) Sets the compression method for the entry.
void	setSize(long size)

	Sets the uncompressed size of the entry data.
void	setTime(long time) Sets the modification time of the entry.
String	toString() Returns a string representation of the ZIP entry.

[041] Table 4 shows an example of two equivalent forms of an expanded query 206 that expand the original search from the help content 106 (e.g., Java Help Documents) to the logic repository 112 using the basis API call list 204 from the original query 202. Table 4 statement A shows the `getCompressedSize` and `setMethod` that may be included in the basis API call list 204. Table 4 statement B shows the expanded query 206 as a nested query, where the original query 202 and the basis API call list 204 (e.g., `getCompressedSize` and `setMethod`) drive the outer query that searches the logic repository 112 for potentially relevant applications 210 to obtain the expanded query result 208 including potentially relevant applications 210 and potentially relevant API calls 212. The expanded query 206 may improve upon the original query 202 by targeting the search performed against the logic repository 112 to obtain application logic 120 with a high probability of including potentially relevant applications 210 and potentially relevant API calls 212.

Table 4 – expanded query 206	
A.	SELECT Potentially Relevant Applications FROM Logic Repository WHERE API Calls in Source Code Files of these Application = <code>getCompressedSize</code> or <code>setMethod</code> . /***** The SQL statement above also expressed below. *****/
B.	SELECT Potentially Relevant Applications

FROM Logic Repository WHERE API Calls in Source Code Files of these Application = { SELECT basis API Calls FROM Java Help Documents WHERE Words in these Documents = compress or uncompress or ZIP or file }.
--

[042] Table 5 shows another example of two equivalent forms of an expanded query 206 that expand the original search from the help content 106 (e.g., Java Help Documents) to the logic repository 112 by combining the original query 202 and the basis API call list 204 to form the expanded query 206. Table 5 statement A shows the getCompressedSize and setMethod (e.g., the basis API call list 204) combined with the original query 202. Table 5 statement B shows the expanded query 206 as a nested query, where the original query 202 and the basis API call list 204 (e.g., getCompressedSize and setMethod) drive the outer query that searches the logic repository 112 for potentially relevant applications 210 to obtain the expanded query result 208 including potentially relevant applications 210 and potentially relevant API calls 212. The expanded query 206 may improve upon the original query 202 by targeting the search performed against the logic repository 112 to obtain application logic 120 with a high probability of including potentially relevant applications 210 and potentially relevant API calls 212.

Table 5 – expanded query 206

A. **SELECT Potentially Relevant Applications**

FROM Logic Repository

WHERE API Calls in

Source Code Files of these Application

= getCompressedSize or setMethod or

compress or uncompress or ZIP or file.

/*** The SQL statement above also expressed below. *****/**

B. **SELECT Potentially Relevant Applications**

FROM Logic Repository

WHERE API Calls in

Source Code Files of these Application

= {

SELECT basis API Calls

FROM Java Help Documents

WHERE Words in these Documents =

compress or uncompress or ZIP or file

} or compress or uncompress or ZIP or file.

[043] Table 6 shows an example of a fragment of logic extracted from the logic repository 112 (e.g., potentially relevant application 210) that includes a potentially relevant API call 212 (e.g., getCompressedSize).

Table 6 – potentially relevant application 210 (e.g., extracted logic fragment)


```

public static void addFilesToExistingZip(File zipFile, File[ ] files)
    throws IOException {
    // get a temp file

    File tempFile = File.createTempFile(zipFile.getName(), null);

    // delete it, otherwise you cannot rename your existing zip to it.

    tempFile.delete();

    int sz = zipFile.getCompressedSize();
    boolean renameOk=zipFile.renameTo(tempFile);

    if (!renameOk && sz == -1)
    {
        throw new RuntimeException("could not rename the file
"+zipFile.getAbsolutePath()+" to "+tempFile.getAbsolutePath());
    }
    byte[ ] buf = new byte[1024];

```

[044] Figure 4 shows an Exemplar system 124 prototyping and application development (EPAD) project area 400. The EPAD project area 400 may include a relevant application list display area 402, relevant API list display area 404, heuristic relevance ranking results display area 406, logic display area 408, and application creation display area 410. The relevant application list display area 402 may include the relevant application list 220 produced by the Exemplar system 124 based on the expanded query 206 search results. The relevant applications list 220 may include a relevant application identifier-1 412 (e.g., a program name, repository identifier, file name, or other program specifier) of relevant application logic 414. A user may select any identifier, as indicated by the arrow 416, to display the relevant application logic 414 (e.g., source code for the program) in the logic display area 408. The EPAD user interface 102 may include a keyboard, mouse, a microphone (or other sensors), joystick, game pad, or the like for the user to interact with the EPAD project area 400.

[045] The relevant API list display area 404 may include the relevant API call list 222 returned by Exemplar system 124 based on the original query 202. The relevant API call list 222 may include a relevant API call identifier-1 418 (e.g., a function call

name) and a relevant API call identifier-2 420 of the relevant API call logic-1 422 (e.g., source code for the function call) and relevant API call logic-2 424, respectively. The EPAD project area 400 may present the relevant API call identifier-1 418 and the relevant API call identifier-2 420 as user selectable, indicated by the arrow 426, to display and highlight the relevant API call logic-1 422 and the relevant API call logic-2 424 in the logic display area 408. In one implementation, the logic display area 408 may highlight the relevant application logic 414, and relevant API call logic-1 422 and relevant API call logic-2 424 so that the user can further determine the relevance of the logic to the high-level processing concept represented in the original query 202.

[046] The heuristic relevance ranking results display area 406, shown in Figure 4, may include an application relevance threshold 428, an API call relevance threshold 430, data flow graph 432, control flow graph 433, and an API call graph 434. The heuristic relevance ranking results display area 406 may display heuristic relevance ranking engine 114 information to assist the user to determine the relevance of user selected logic. As will be described in more detail below, the heuristic relevance ranking engine 114 may determine the application heuristic relevance ranking 216 for the relevant application logic 414 based on the number of relevant API calls (e.g., relevant API call logic-1 422 and relevant API call logic-2 424) found in the relevant application logic 414 in comparison to other relevant application logic 436 identified by Exemplar system 124. For example, the high-level processing concept example "compress/uncompress ZIP file" 302 may be entirely implemented in relevant application logic 414, but only partially implemented in the other relevant application logic 436. As a result, the heuristic relevance ranking engine 114 may assign the relevant application logic 414 a higher application heuristic relevance ranking 216 than the other relevant application logic 436. In another implementation, the heuristic relevance ranking engine 114 may determine the API call heuristic relevance rankings 218 of the relevant API call logic-1 422 and the relevant API call logic-2 424, based on analyzing semantics derived from the expanded query 206 and the expanded search result 208, which establish the behaviour of the relevant API call logic-1 422, the relevant API call logic-2 424, and the relevant applications logic-1 414.

[047] The application relevance threshold 428 and API call relevance threshold 430, shown in Figure 4, may be user selectable and/or pre-configured with system default values. In another implementation, Exemplar system 124 may determine the application relevance threshold 428 and the relevance threshold 430 based a number of factors (e.g., the complexity of the high-level processing concept represented by the original query 202, and the number of potentially relevant applications 210 and potentially relevant API calls 212 identified by the expanded search result 208). Exemplar system 124 may use the application relevance threshold 428 and the relevance threshold 430 to further refine the relevant applications list 220 and the relevant API calls list 222, respectively. In one implementation, the application relevance threshold 428 and the relevance threshold 428 may determine an application heuristic relevance ranking 216 value that the potentially relevant applications 210 must meet to be included on the relevant applications list 220. The API call relevance threshold 430 may also determine the API call heuristic relevance ranking 218 value that the potentially relevant API calls 212 must meet to be included on the relevant API calls list 222. For example, an application relevance threshold 428 of 1 may indicate a low relevance requirement (e.g., requiring loosely relevant applications, and low application heuristic relevance rankings 216) and allow a large number of potentially relevant applications 210 to qualify as relevant applications (e.g., relevant application logic-1 414). In another example, an application relevance threshold 428 of 10 may indicate a high relevance requirement (e.g., requiring highly relevant applications, and high application heuristic relevance rankings 216) and allow only a fewer number of potentially relevant applications 210 to qualify as relevant applications. The heuristic relevance ranking engine 114 may also use the data flow graph 432 and control flow graph 433 to determine the application heuristic relevance ranking 216 and API call heuristic relevance ranking 218, and visually describe the relationships between the relevant application logic 414, the relevant API call logic-1 422, and the relevant API call logic-2 424, discussed in further detail below.

[048] The relevant API call logic-1 422 and the relevant API call logic-2 424 may be user selectable (indicated by the arrow 442), and provide the user the ability to generate a new application 440 with the selected logic. To that end, the EPAD

project area 400 may implement point-and-click, drag-and-drop functionality for a user to select relevant API call logic-1 422 and relevant API call logic-2 424 to generate the new application 440. The EPAD project area 400 may also build the new application 440 by combining user selectable other relevant application logic 436, relevant API call logic-1, and relevant API call logic-2. The application creation display area 410 may also identify requirements 444 for the high-level processing concept represented by the original query 202. For example, a developer may desire to identify and confirm the requirements 444 for implementing a high-level processing concept (e.g., "send secure XML"). In one implementation, Exemplar may generate requirements documentation and end user documentation based on the help content 106 related to the other relevant application logic 436, the relevant API call logic-1, and the relevant API call logic-2 used to build the new application 440, and identify the requirements 444 in support of the new application 440.

[049] Figure 5 shows a more detailed view of the Exemplar system 124. The Exemplar system 124 may include a communications interface 504 used to communicate with various resources internal and external to Exemplar system 124, memory 506, and a processor 508. The processor 508 may execute any of the logic described below. The memory 506 may include the EPAD user interface 102 that employs the interface logic 510 to generate the Exemplar query interface 300, and the EPAD project area 400. The interface logic 510 may include graphics libraries, window rendering calls, and other user interface logic operable to display interface elements, receive input, and pass the input to any particular program logic in the Exemplar system 124.

[050] The memory 506 may also include expanded search logic 514. Table 5, above, shows an expanded query 206 where the search logic 514 forms the expanded query by combining the original query 202 and the basis API call list 204 to form the expanded query 206. More generally, the expanded search logic 514 combines the original query 202 and the basis logic results 516 to form the expanded query 206, and executes an expanded search using the expanded query 206. The basis logic results 516 may include the basis API call list 204, including zero or more basis API call identifiers (e.g., the basis API call identifier-1 518), and a basis application list 520, including zero or more basis application identifiers (e.g.,

the basis application identifier-1 521). The expanded search logic 514 thereby obtains the expanded search results 208. The expanded search result 208 may include potentially relevant applications 210, and potentially relevant API calls 212 that include zero or more potentially relevant application identifiers-1 522 and zero or more potentially relevant API call identifiers (e.g., potentially relevant API call identifier-1 524 and potentially relevant API call identifier-2 526).

[051] Figure 6 shows other features of the Exemplar system 124. The memory 506 may also include the heuristic relevance ranking engine 114 with the heuristic relevance ranking logic 602 that generates the application heuristic relevance ranking 216 and API call heuristic relevance ranking 218. The threshold logic 604 may apply the application relevance threshold 428 and API call relevance threshold 430 to the application heuristic relevance rankings 216 and API call heuristic relevance rankings 218 to determine whether potentially relevant applications 210 and potentially relevant API calls 212 qualify for inclusion in the relevant applications list 220 and the relevant API calls list 222. In other words, the threshold logic 604 may implement comparison logic to determine when potentially relevant logic qualifies as relevant logic.

[052] The memory 506 may also include analyzer logic 606 that the processor 508 executes to identify application metadata 608 and API metadata 610 of the potentially relevant applications 210, and the potentially relevant API calls 212, respectively. Examples of application metadata 608 include application descriptions, application logic comments, application parameter names, and application parameter types of existing applications. Similarly, examples of API metadata 610 include API descriptions, API logic comments, API parameter names, and API parameter types.

[053] The analyzer logic 606 may generate the data flow graph 432 and control flow graph 433 to obtain the API call graph 434. The API call graph 434 may include nodes (e.g., node-1 612 and node-2 614) that represent potentially relevant API calls 212 and data flow edges (e.g., data flow edge 616) between the potentially relevant API calls 212 to indicate data flow. Figure 7 provides additional examples. The analyzer logic 606 may determine the data flow edge count 618 corresponding to the number of connections between potentially relevant API calls 212 within the potentially relevant application 210. A graph with 'n' nodes has as many as $n(n-1)$

edges between nodes. The data flow edge count 618 provides insight into the degree of connectedness for the data flow graph 432. The analyzer logic 606 may also assign link values 620 to the edges between nodes, discussed in detail below. In one implementation, the analyzer logic 606 may determine the connectivity rankings 214 (e.g., strong connectivity ranking 622 and weak connectivity ranking 624) for each connection between the potentially relevant API calls 212 based on common API parameters 626, discussed in detail below.

[054] Figure 6 further illustrates that memory 506 may include selection logic 628 and application creation logic 630. The processor 508 may execute the selection logic 628 to allow a user to select relevant application logic 414, and relevant API call logic (e.g., the relevant API call logic-1 422 and the relevant API call logic-2 424) to develop the new application 440. In another implementation, selection logic 628 may provide a user drag-and-drop point-and-click functionality to select other relevant application logic 436 to combine with the relevant API call logic-1 422, and the relevant API call logic-2 424 to build the new application 440. The processor 508 may execute the application creation logic 630 to identify requirements 444 for the high-level processing concept represented by the original query 202 by identifying the help content 106 used to obtain the basis logic results 516 corresponding to the user selected other relevant application logic 436, relevant application logic 414, relevant API call logic-1 422, and relevant API call logic-2 424. In one implementation, the application creation logic may form a query using the other relevant application logic 436, relevant application logic 414, relevant API call logic-1 422, and relevant API call logic-2 424 to obtain the help content 106 that describes the requirements 444. The application creation logic 630 may generate customized requirements documents from the help content 106 corresponding to the user selected other relevant application logic 436, relevant application logic 414, relevant API call logic-1 422, and relevant API call logic-2 424.

[055] Figure 7 shows API call graphs 434 for two different potentially relevant applications (e.g., a potentially relevant application A and potentially relevant application B). The heuristic relevance ranking engine 114 may assign a higher heuristic relevance ranking 216 to the potentially relevant application A than the potentially relevant application B based on the number of potentially relevant API

calls 212, the connectivity rankings 214 and link values 620 assigned to each connection between potentially relevant API calls 212 included in the potentially relevant application A and potentially relevant application B, respectively.

[056] In one implementation, the logic analyzer 116 produces the API call graphs 434. The logic analyzer 116 may identify the application metadata 608 and API metadata 610 of the potentially relevant applications 210, and the potentially relevant API calls 212, respectively, to analyze the data flow paths and connectivity between the potentially relevant API calls 212. The logic analyzer 116 may provide the application metadata 608 and API metadata 610 to the heuristic relevance ranking engine 114. In an alternative implementation, the heuristic relevance ranking engine 114 may identify application metadata 608 and API metadata 610, and produce the data flow graph 432 and control flow graph 433 using logic analysis formulas, rules and equations to obtain the API call graphs 434. The data flow graphs 432, control flow graphs 433 and API call graphs 434 may be represented as mathematical structures. The logic analyzer 116 may obtain the API call graphs 434 as a result of comparing data flow and control flow between potentially relevant API calls 212.

[057] In one implementation, the logic analyzer 116 may perform control flow analysis on the potentially relevant application 210 to obtain control flow graphs 433, and perform data flow analysis on the control flow graphs 433 to obtain data flow graphs. The data flow graphs 432, control flow graphs 433, and API call graphs may similarly include nodes and edges. The logic analyzer 116 may obtain a control flow graph 433 by logically partitioning a potentially relevant application 210 as a result of parsing the logic of the potentially relevant application 210 into nodes that represent logic that includes API calls. The logic analyzer 116 may assign parsed logic of the potentially relevant application 210 to an assigned node until the logic analyzer 116 identifies a potentially relevant API call or branching logic (e.g., if-then, switch-case, and do-while), and add the assigned node to the control flow graph 433. Where a program includes multiple potentially relevant applications 210, the logic analyzer 116 may merge the control flow graphs 433 produced for each potentially relevant application into a single control flow graph 433. The logic analyzer 116 may obtain the API call graph 434 by comparing the edges in the control flow graphs 433 with the edges in the data flow graph 432. For example, where a control flow graph 433

includes an edge that a data flow graph 432 does not include, the logic analyzer 116 may not include the edge in the corresponding API call graph 434. However, where a control flow graph 433 includes an edge that the data flow graph 432 also includes, the logic analyzer 116 may include the edge in the API call graph 434.

[058] In one implementation, the logic analyzer 116 may receive user input to determine particular dependencies between API calls. For example, where a potentially relevant application 210 uses a function pointer (e.g., a type of pointer used in C, and C++ languages) to reference a potentially relevant API call 212 and a hash table (e.g., a data structure that associates keys with values) to store an object that represents a data element passed between API calls, the logic analyzer 116 may receive user input to determine dependencies between API calls because the logic analyzer 116 may otherwise interpret multiple possible dependencies between API calls when in fact only one or a finite set of valid dependencies exists.

[059] In another implementation, the logic analyzer 116 may analyze the data flow paths (e.g., edges between nodes discussed below) (e.g., link heuristics) of the potentially relevant applications 210, and potentially relevant API call logic 212 to determine the connectivity rankings 214 of each connection between potentially relevant API calls 212. In one implementation, the heuristic relevance ranking engine 114 may determine the application heuristic relevance ranking 216 for the potentially relevant application 210, shown in Figure 7 as potentially relevant application A, based on the total number of API calls 'n' represented by nodes 712-720 that represent different potentially relevant API calls 212 found in the potentially relevant application 210, the total number of connections between the potentially relevant API calls 212 (e.g., edges 712-720) equal to $n(n-1)$ (e.g., data flow edge count 614), the quality of the connections (e.g., strong connectivity or weak connectivity), and the type of link (e.g., loop link, single link, or no link) between the potentially relevant API calls 212.

[060] The applications metadata 608 and API metadata 610 may describe the data flow paths between the different potentially relevant API calls 212 (e.g., nodes 702-710) within the potentially relevant application 210. For example, the logic analyzer 116 may determine common API parameters 626 and logic branches (e.g., if-then-else) found within the potentially relevant application 210 and potentially relevant API

calls 212 to generate the data flow graphs 432, control flow graphs 433 and API call graphs 434. The logic analyzer 116 may, as Figure 7 also illustrates, identify the function (e.g., $K(x)$, $J(x)$, $S(y)$, $P(y)$, $F(x)$, and $G(z)$) of each potentially relevant API call 212 (e.g., 702-710, and 722) to determine the connectivity rankings 214.

[061] In one implementation, the logic analyzer 116 may assign a weight W_i (e.g., connectivity ranking 214) to each connection between the potentially relevant API calls 212 (e.g., nodes 712-720). The logic analyzer 116 may assign weak connections a weight of 0.5 and strong connections a weight of 1.0 depending on multiple factors. For example, edge 712, edge 716 and edge 720 may represent weak connections between potentially relevant API calls 212 represented by node pairs 702 and 710, 702 and 704, and 706 and 708 (e.g., function pairs $K(x)$ and $F(x)$, $K(x)$ and $J(x)$, and $S(y)$ and $P(y)$, respectively). Following the above example, where functions $K(x)$ and $F(x)$ share a common API parameter 626, but neither function generates the value of the common API parameter 626 then the logic analyzer 116 may assign the connectivity ranking 214 between node pair 702 and 710, represented by edge 720, a weak connection weight of 0.5. A weak connection assigned to a node pair (e.g., 702 and 710) may indicate a low relative probability (e.g., in comparison to the connectivity rankings of other node pairs) that the node pair implements the high-level processing concept represented by the original query 202. The logic analyzer 116 may use other heuristic analysis methods and tools to determine whether to assign a weak connection to a connectivity ranking 214.

[062] Alternatively, edge 714, and edge 718 may represent strong connections between potentially relevant API calls 212, represented by node pairs 702 and 708, and 704 and 706 (e.g., function pairs $K(x)$ and $P(y)$, and $J(x)$ and $S(y)$, respectively). The logic analyzer 116 may determine that where function $J(x)$ produces variable y , which both $J(x)$ and $S(y)$ share then the node pair 704 and 706, represented by edge 714, may be assigned a strong connectivity ranking 622. A strong connection assigned to a node pair (e.g., 704 and 706) may indicate a high relative probability (e.g., in comparison to the connectivity rankings of other node pairs) that the node pair implements the high-level processing concept represented by the original query 202. The logic analyzer 116 may use other heuristic analysis methods and tools to determine whether to assign a strong connection to a connectivity ranking 214.

[063] The logic analyzer 116 may also assign a link value L (e.g. link value 620) to each connection between potentially relevant API calls 212. For example, Exemplar system 124 may assign a link value L equal to 1 where a loop link (e.g., edges 712-718 form a loop) exists between potentially relevant API calls 212 (e.g., nodes 702-708). Exemplar system 124 may assign a link value L equal to 0.5 where a single link (e.g., edge 720) exists between potentially relevant API calls 212 (e.g., nodes 702-708). In another implementation, Exemplar system 124 may assign a link value L equal to 0 where no link exists between potentially relevant API calls 212 (e.g., node 722 represents a potentially relevant API call 212 that does not have a connection with other potentially relevant API calls in a potentially relevant application 210). Additional, different, or fewer weights may be used. The heuristic relevance ranking engine 114 may use the connectivity rankings 214 and link values 620 assigned to each connection between potentially relevant API calls 212 to determine the application heuristic relevance ranking 216 according to:

[064]
$$\sum_{i=1}^{n(n-1)} \frac{W_i L_i}{n(n-1)}$$
 where i ranges from 1 to $n(n-1)$. (Equation 1)

[065] The logic analyzer 116 may determine an API call heuristic relevance ranking 218 for a potentially relevant API call 212 based on the connectivity ranking 214 and link value 620 assigned to each edge that includes the potentially relevant API call 212. For example, where m represents the number of node pair including a particular node (e.g., 702 and 704, 702 and 708, and 702 and 710) and the number of edges (e.g., 712, 718 and 720) that include the node equals $m(m-1)$, and the assigned value for each connectivity ranking 214 and link value 620 for each edge that includes the node represent W and L , respectively, the API call heuristic relevance ranking 218 for the node may be determined according to Equation 1 above where m substitutes for n :

[066]
$$\sum_{j=1}^{m(m-1)} \frac{W_j L_j}{m(m-1)}$$
 where j ranges from 1 to $m(m-1)$.

[067] Figure 8 shows the acts that the Exemplar system 124 may take to obtain a relevant application list 210. The Exemplar system 124 may use the interface logic 510 to receive an original query 202 representing a high-level processing concept (e.g., "send secure XML") (802). The help content processor 104 may execute an

original search using the original query 202 (804) to obtain an original search result that includes basis logic results 516 (e.g., basis API call list 204 or basis application list 520). The basis logic results 516 may include a basis logic (e.g., API call) identifier (806). The Exemplar system 124 may combine the original query 202 with any part of the basis logic list to form an expanded query (808). The expanded search logic 514 may execute the expanded query 206 to obtain an expanded search result 208 that includes a potentially relevant applications 210 list and potentially relevant API calls 212 list (810). The potentially relevant logic list may identify potentially relevant logic. The analyzer logic 606 may analyze the potentially relevant logic identified by the potentially relevant application identifiers with respect to the logic repository 112 (812). The heuristic relevance ranking logic 602 may use the connectivity rankings 214, and link values 620 to determine the application heuristic relevance rankings 216 and API call heuristic relevance rankings 218 for the potentially relevant applications 210 and potentially relevant API calls 212 using Equation 1 or another formulation (814). The heuristic relevance ranking logic 602 may apply the application relevance threshold 428 and the API call relevance threshold 430 using the threshold logic 604 to determine whether the potentially relevant applications 210 and the potentially relevant API calls meet the application relevance threshold 428, and API call relevance threshold 430, respectively (816). The heuristic relevance ranking logic 602 may add the potentially relevant application 210 to the relevant application list 220 where the potentially relevant application 210 meets the application relevance threshold 428 (818). The heuristic relevance ranking logic 602 may determine the application heuristic relevance ranking 216 and API call heuristic relevance ranking 218 for each potentially relevant application 210 and potentially relevant API call 212 included the expanded search result 208 (820).

[068] Figure 9 shows the acts that the Exemplar system 124 heuristic ranking engine 114 may take to assign an application heuristic relevance ranking 216 to a potentially relevant application 210. The analyzer logic 602 may analyze the potentially relevant application 210 and potentially relevant API calls 212 found in the logic repository (904). The analyzer logic 602 may generate and analyze a data flow graph 432 and control flow graph 433 (906) used to generate and analyze an API

call graph 43 (908). The analyzer logic 602 may use the data flow graph 432, control flow graph 433, and API call graph to determine the link values 650 for the potentially relevant API calls included in a potentially relevant application 210 and assign a connectivity ranking 214 to each connection between potentially relevant API calls 212 (910). The heuristic relevance ranking logic 602 may determine an application heuristic relevance ranking 216 for each of the potentially relevant applications 210 (912) (e.g., the application heuristic relevance ranking 216 and an API call heuristic relevance ranking 218 may be determined according to Equation 1, as discussed above). The heuristic relevance ranking logic 602 may determine the application heuristic relevance ranking 216 and API call heuristic relevance ranking 218 for each potentially relevant application 210 and potentially relevant API call 212 included the expanded search result 208 (914).

[069] Figure 10 shows the processing that the selection logic and application creation logic may take to generate a new application 440. EPAD project area 400 may use the selection logic 624 to detect selection of a relevant API call Identifier (e.g., as indicated by arrows 426 drawn from the relevant API call identifier-1 418 and the relevant API call identifier-2 420 to the relevant API call logic-1 422 and the relevant API call logic-2 424) from the relevant API call list 222 (1002). The EPAD project area 400 may present the relevant API call logic (e.g., the relevant API call logic-1 422 and the relevant API call logic-2 424) that implements the relevant API calls (1004). The EPAD project area may use the selection logic 624 to detect selection of the relevant API call logic (e.g., the relevant API call logic-1 422 and the relevant API call logic-2 424) and the other relevant application logic 436 to generate a new application 440 (1006). The EPAD project area may provide an option to generate the new application 440, using the application creation logic 626, following selection of the relevant API call logic (e.g., the relevant API call logic-1 422 and the relevant API call logic-2 424) and the other relevant application logic 436 to generate a new application 440 (1008). Exemplar system 124 may also identify requirements 444 for the high-level processing concept represented by the original query 202. In one implementation, Exemplar system 124 may generate requirements documentation and end user documentation based on the help content 106 related to the other relevant application logic 436, the relevant API call logic-1 422, and the

relevant API call logic-2 424 combined to generate the new application 440 and identify the requirements 444.

[070] The Exemplar system 124 greatly reduces the time, cost, and other resource expenditures associated with implementing a new application. The Exemplar
5 system 124 produces relevant results starting with high-level processing concepts. A software developer may use the relevant applications to rapidly build new application prototypes.

[071] A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made.

WHAT IS CLAIMED IS:

1. A method comprising:

accessing a first search result comprising at least one application programming
5 interface (API) call identifier, wherein said first search result is associated with a first
search performed using a first query, and wherein said first search is performed on a
first database;

combining said first query with said at least one API call identifier to form a
second query; and

10 performing a second search using said second query to obtain a second
search result comprising at least one application identifier of at least one application,
wherein said second search is performed on a second database.

2. The method of Claim 1 further comprising:

15 analyzing said at least one application to determine an application relevance
ranking for said at least one application; and

ranking said at least one application with respect to at least one other
application based on said application relevance ranking.

3. The method of Claim 1 further comprising:

20 generating a list of applications including said at least one application identifier
of said at least one application.

4. The method of Claim 2 further comprising:

25 comparing said application relevance ranking to an application relevance
threshold; and

if said application relevance ranking is greater than said application relevance
threshold, adding said at least one application identifier of said at least one application
to a list of applications.

30

5. The method of Claim 1 further comprising:

detecting a selection of said at least one application identifier of said at least one application, wherein said at least one application identifier is displayed in a first region of a user interface; and

5 responsive to said selection, displaying code associated with said at least one application in a second region of said user interface.

6. The method of Claim 1 further comprising:

10 detecting a selection of said at least one API call identifier, wherein said at least one API call identifier is displayed in a third region of a user interface; and

responsive to said selection, displaying code associated with said at least one API call identifier in a second region of said user interface.

7. The method of Claim 2, wherein said analyzing further comprises

15 determining said application relevance ranking based on at least one attribute associated with said at least one application.

8. The method of Claim 7, wherein said at least one attribute is selected from a group consisting of:

20 a number of nodes of an API call graph associated with said at least one application;

a respective weight associated with each connection between a plurality of nodes of said API call graph; and

25 a respective link value associated with each connection between a plurality of nodes of said API call graph.

9. The method of Claim 1 further comprising:

30 accessing information associated with said at least one API call identifier, wherein said information is accessed from a source selected from a group consisting of a database, website, knowledge exchange and document repository, and

wherein said combining further comprises combining said first query with said at least one API call identifier and said information to form said second query.

10. A non-transitory computer-readable medium having computer-readable
5 program code embodied therein for causing a computer system to perform a method comprising:

accessing a first search result comprising at least one application programming interface (API) call identifier, wherein said first search result is associated with a first search performed using a first query, and wherein said first search is performed on a
10 first database;

combining said first query with said at least one API call identifier to form a second query; and

performing a second search using said second query to obtain a second search result comprising at least one application identifier of at least one application, wherein
15 said second search is performed on a second database.

11. The computer-readable medium of Claim 10, wherein said method further comprises:

analyzing said at least one application to determine an application relevance
20 ranking for said at least one application; and

ranking said at least one application with respect to at least one other application based on said application relevance ranking.

12. The computer-readable medium of Claim 10, wherein said method further
25 comprises:

generating a list of applications including said at least one application identifier of said at least one application.

13. The computer-readable medium of Claim 11, wherein said method further
30 comprises:

comparing said application relevance ranking to an application relevance threshold; and

if said application relevance ranking is greater than said application relevance threshold, adding said at least one application identifier of said at least one application
5 to a list of applications.

14. The computer-readable medium of Claim 10, wherein said method further comprises:

detecting a selection of said at least one application identifier of said at least
10 one application, wherein said at least one application identifier is displayed in a first region of a user interface; and

responsive to said selection, displaying code associated with said at least one application in a second region of said user interface.

15 15. The computer-readable medium of Claim 10, wherein said method further comprises:

detecting a selection of said at least one API call identifier, wherein said at least one API call identifier is displayed in a third region of a user interface; and

responsive to said selection, displaying code associated with said at least one
20 API call identifier in a second region of said user interface.

16. The computer-readable medium of Claim 11, wherein said analyzing further comprises determining said application relevance ranking based on at least one attribute associated with said at least one application.

25

17. The computer-readable medium of Claim 16, wherein said at least one attribute is selected from a group consisting of:

a number of nodes of an API call graph associated with said at least one application;

30 a respective weight associated with each connection between a plurality of

nodes of said API call graph; and

a respective link value associated with each connection between a plurality of nodes of said API call graph.

5 18. The computer-readable medium of Claim 10, wherein said method further comprises:

accessing information associated with said at least one API call identifier, wherein said information is accessed from a source selected from a group consisting of a database, website, knowledge exchange and document repository, and

10 wherein said combining further comprises combining said first query with said at least one API call identifier and said information to form said second query.

15 19. A system comprising a processor and a memory, wherein said memory comprises instructions that when executed by said processor implement a method comprising:

accessing a first search result comprising at least one application programming interface (API) call identifier, wherein said first search result is associated with a first search performed using a first query, and wherein said first search is performed on a first database;

20 combining said first query with said at least one API call identifier to form a second query; and

performing a second search using said second query to obtain a second search result comprising at least one application identifier of at least one application, wherein said second search is performed on a second database.

25

20. The system of Claim 19, wherein said method further comprises:

analyzing said at least one application to determine an application relevance ranking for said at least one application; and

30 ranking said at least one application with respect to at least one other application based on said application relevance ranking.

21. The system of Claim 19, wherein said method further comprises:
generating a list of applications including said at least one application identifier
of said at least one application.

5

22. The system of Claim 20, wherein said method further comprises:
comparing said application relevance ranking to an application relevance
threshold; and

if said application relevance ranking is greater than said application relevance
10 threshold, adding said at least one application identifier of said at least one application
to a list of applications.

23. The system of Claim 19, wherein said method further comprises:
detecting a selection of said at least one application identifier of said at least
15 one application, wherein said at least one application identifier is displayed in a first
region of a user interface; and
responsive to said selection, displaying code associated with said at least one
application in a second region of said user interface.

20 24. The system of Claim 19, wherein said method further comprises:
detecting a selection of said at least one API call identifier, wherein said at least
one API call identifier is displayed in a third region of a user interface; and
responsive to said selection, displaying code associated with said at least one
API call identifier in a second region of said user interface.

25

25. The system of Claim 20, wherein said analyzing further comprises
determining said application relevance ranking based on at least one attribute
associated with said at least one application.

30 26. The system of Claim 25, wherein said at least one attribute is selected from

a group consisting of:

a number of nodes of an API call graph associated with said at least one application;

5 a respective weight associated with each connection between a plurality of nodes of said API call graph; and

a respective link value associated with each connection between a plurality of nodes of said API call graph.

27. The system of Claim 19, wherein said method further comprises:

10 accessing information associated with said at least one API call identifier, wherein said information is accessed from a source selected from a group consisting of a database, website, knowledge exchange and document repository, and

wherein said combining further comprises combining said first query with said at least one API call identifier and said information to form said second query.

15

28. The method of Claim 1, wherein said first database is a help content database, and wherein said second database is a logic database.

29. The computer-readable medium of Claim 10, wherein said first database is

20 a help content database, and wherein said second database is a logic database.

30. The system of Claim 19, wherein said first database is a help content database, and wherein said second database is a logic database.

1/10




Figure 1

2/10




Figure 2

3/10




Figure 3

4/10




Figure 4

5/10




Figure 5

6/10




Figure 6

7/10




Figure 7




Figure 8

9/10




Figure 9

10/10




Figure 10

100

