



(19) **United States**

(12) **Patent Application Publication**
Minogue et al.

(10) **Pub. No.: US 2007/0033168 A1**

(43) **Pub. Date: Feb. 8, 2007**

(54) **AGENT RANK**

Publication Classification

(76) Inventors: **David Minogue**, Palo Alto, CA (US);
Paul A. Tucker, Mountain View, CA (US)

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

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

(57) **ABSTRACT**


The present invention provides methods and apparatus, including computer program products, implementing techniques for searching and ranking linked information sources. The techniques include receiving multiple content items from a corpus of content items; receiving digital signatures each made by one of multiple agents, each digital signature associating one of the agents with one or more of the content items; and assigning a score to a first agent of the multiple agents, wherein the score is based upon the content items associated with the first agent by the digital signatures.

Correspondence Address:

FISH & RICHARDSON P.C.
PO BOX 1022
MINNEAPOLIS, MN 55440-1022 (US)

(21) Appl. No.: **11/200,406**

(22) Filed: **Aug. 8, 2005**






FIG. 1

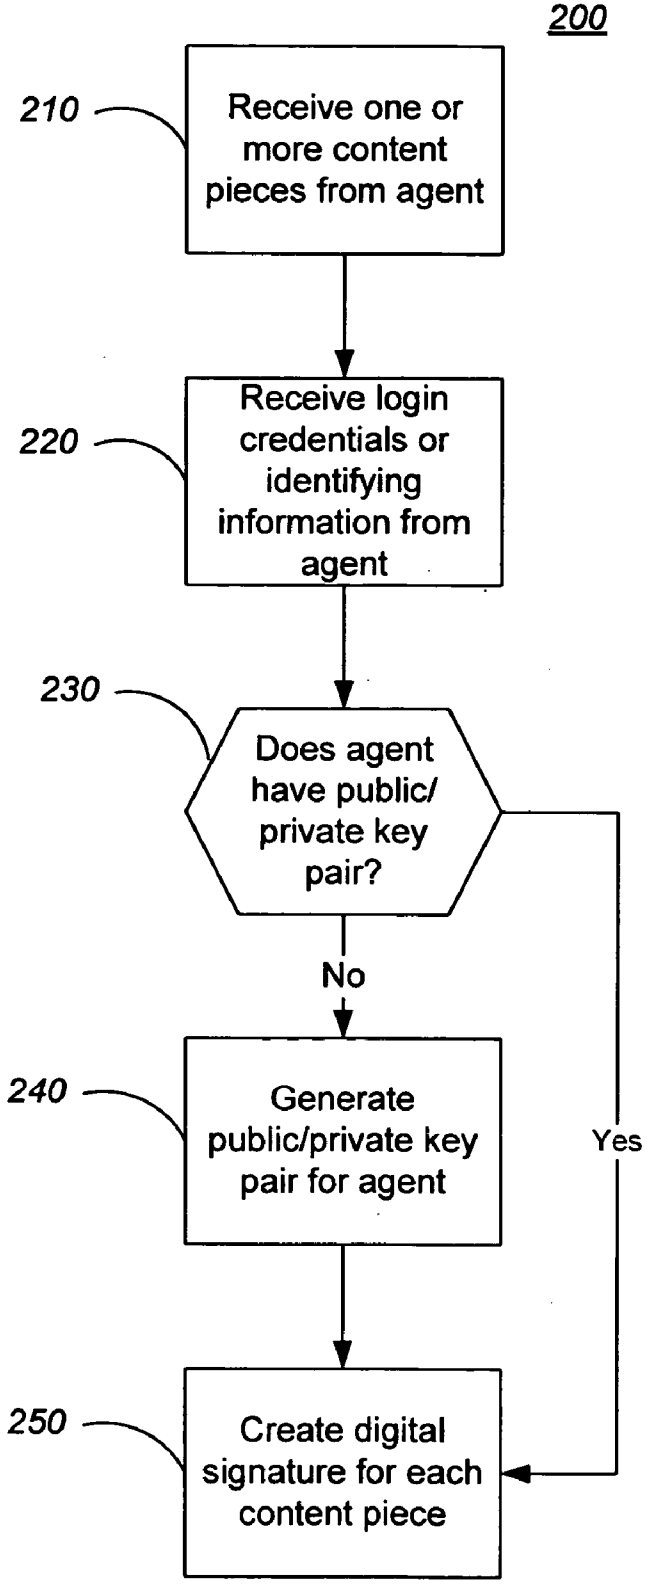


FIG. 2




FIG. 3

AGENT RANK

[0001] This invention relates to the searching and ranking of linked information sources.

[0002] Information Retrieval (IR) is concerned with locating desired elements of information among a large corpus. A search engine is a one example of an IR system that enables documents (usually but not necessarily limited to text) to be retrieved from a large corpus on the basis of their degree of relevance with respect to a compact query presented by a user. The order in which documents are retrieved or presented is the ranking created by the search engine: the highest ranked documents, with respect to the query, are returned or presented first. Search engine ranking may be affected by both query-dependent and query-independent criteria. Query-dependent criteria generally attempt to identify the degree to which a document is semantically related to the query. An example is the correspondence of word frequency distributions. Query-independent criteria often seek to identify the degree to which a document is generally "good", e.g. authoritative, intelligible, not fraudulent or not deceptive. An example of a query-independent criterion is the score computed by the PageRank algorithm, or similar algorithms that examine the link structure of a corpus of documents.

[0003] As mentioned above, query-independent criteria can provide a way of measuring the authoritativeness of a specific information source. For example, the more information sources that point to a specific information source, the higher the search rating score the information source gets, and the more authoritative it is judged to be. In some instances, the search rating algorithm is recursive, meaning that a information source's search rating score is based not only on the number of information sources that reference the original information source, but also on the search rating scores of the referencing information sources. In other words, the search rating score of an information source is based on both the number and quality of the referencing information sources.

[0004] For some information sources, all of the content is under the control of a single agent. In such cases, the reputation of the agent can be directly correlated with the content of the information source. In other cases, however, control may be delegated among several agents, each controlling a partition of the information source. To the extent that these partitions can be identified, agent reputation can be calculated at the partition level.

[0005] In general, however, it is difficult to correlate content on an information source with the agents responsible for creating or publishing that content. For example, an individual author may contribute content to multiple information sources, content within a single information source may originate from multiple agents, or ownership and control of information sources may change over time. As another example, a single web page can contain content controlled by multiple agents, such as advertisements which appear alongside news articles.

SUMMARY

[0006] The present invention provides methods and apparatus, including computer program products, implementing techniques for searching and ranking linked information sources.

[0007] In one aspect, the techniques include receiving multiple content items from a corpus of content items; receiving digital signatures each made by one of multiple agents, each digital signature associating one of the agents with one or more of the content items; and assigning a score to a first agent of the multiple agents, wherein the score is based upon the content items associated with the first agent by the digital signatures.

[0008] Implementations of the invention can include one or more of the following features. The techniques may further include determining the validity of the digital signatures. If no digital signature associates an agent with a specific content item, the content item is associated with an owner of a location where the specific content is found and a score is assigned to the owner based on the specific content item. The content items associated with the first agent include a content item that contains a digital signature associating the first agent with the content item. The content items associated with the first agent include a content item that includes a link to a digital signature associating the first agent with the content item. The content items associated with the first agent include a content item that is a web site or a portion of the web site. A second agent is associated by a second digital signature with a second content item with which the first agent is associated, and the second agent makes an assertion about the content item. The first agent makes an assertion with a digital signature that the first agent is an author of the second content item. The second agent makes an assertion with the second digital signature that the second agent is a reviewer of the second content item. The second agent makes an assertion with the second digital signature that the second agent is an editor of the second content item. The second agent makes an assertion with the second digital signature that the second agent is a publisher of the second content item. Assigning a score to the first agent can include assigning the score based on unsigned content items that the first agent is associated with as an owner of one or more locations where the unsigned content items are found. Assigning a score to the first agent can include assigning the score based on one or more assertions made by one or more other agents about content items associated with the first agent. Assigning a score to the first agent can include assigning the score based on one or more assertions made by the first agent. The score is used in ordering results of a search of the corpus.

[0009] Particular embodiments of the invention can be implemented to realize one or more of the following advantages. The identity of individual agents responsible for content can be used to influence search ratings. The identity of agents can be reliably associated with content. The granularity of association can be smaller than an entire web page, so that agents can disassociate themselves from information appearing in proximity to information for which the agent is responsible. An agent can disclaim association with portions of content, such as advertising, that appear on the agent's web site. The same agent identity can be attached to content at multiple locations. Multiple agents can make contributions to a single web page where each agent is only associated to the content that the agent provided.

[0010] Query-independent rankings of content and authors can be calculated. A query-independent ranking can be, but need not be, calculated offline, prior to accepting a user query of the content, and then used to calculate a query-

dependent ranking used for presentation of results. An algorithm used for calculating a query-independent ranking could also be used within the context of a specific query, with minimal modification, to calculate a ranking specific to that query. For example, the corpus can be limited to the query-relevant content.

[0011] Particular embodiments implement techniques for computing agent ranks on the basis of a corpus of content signed by those agents, where the corpus optionally contains explicit links among documents and signed content. The agent ranks can optionally also be calculated relative to search terms or categories of search terms. For example, search terms (or structured collections of search terms, i.e., queries) can be classified into topics, e.g., sports or medical specialties, and an agent can have a different rank with respect to each topic.

[0012] One implementation of the invention provides all of the above advantages.

[0013] The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the invention will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] FIG. 1 is a schematic representation of a web page.

[0015] FIG. 2 is a flow chart of a method for generating and managing the public and private keys necessary to generate signatures.

[0016] FIG. 3 illustrates a linked database.

[0017] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0018] FIG. 1 shows a mock-up of a sample web page 100. The web page 100 can include various content pieces, including main body text 105, comments 110, 115, image 120, and advertising 125. Each of these content pieces can be created by a different agent. For example, in one implementation, the main body text 105 is created by the owner of the web page 100, the comment 110 is authored by a first agent of the web page 100, and comment 115 is authored by a second agent of the web page 100. An agent is any individual or entity that either provides content pieces, edits existing content pieces, or reviews existing content pieces on a web page. An owner of a web page is the agent that has ultimate control over the web page, including control over all of the content pieces of the web page, including content pieces provided by other agents. Agents that are not owners generally have limited control over content pieces on the web page. For example, in one implementation, a non-owner agent can place a comment on the web page, but not edit or delete the comments of other agents that are included within the web page.

[0019] In an alternative implementation, the comment 110 is authored by the first agent of the web page 100, and the comment 115 is authored by the owner of the web page 100. The advertising 125 can be provided by a third-party advertising service, and the contents of the advertising may or may not be under the control of the owner of the web page

100. In other words, even though the web page 100 may be owned by a single agent, it is possible for content pieces within the web page 100 to have been created or supplied by agents other than the owner of the web page 100.

[0020] Each content piece can be signed with a digital signature, either directly by the agent or indirectly on behalf of the agent. The digital signature identifies the agent that actually created each content piece on the web page 100. In one implementation, each individual content piece on a web page is signed separately. In an alternative implementation, one or more content pieces on a web page is signed while other content pieces on the same web page remain unsigned. In another implementation, a digest or hash of the content piece or content pieces can be generated, and the digest or hash of the content piece is signed. Any suitable protocol for creating and validating digital signatures can be used, e.g., XML Digital Signatures. Additional information about XML Digital Signatures may be found in the XML-Signature Syntax and Processing Recommendation of Feb. 12, 2002, available from the World Wide Web Consortium (W3C) at <http://www.w3.org/TR/xmlsig-core/> and incorporated here by reference.

[0021] The agent signing each content piece can claim various roles relative to the content, e.g., author, publisher, editor, or reviewer. The signature provides evidence that a particular agent has asserted its role with respect to the signed content piece, as the agent has exclusive access to the private key used to sign the content piece. In one implementation, the digital signature can include within the scope of the content signed other metadata such as creation date, review score, or recommended keywords for search.

[0022] In one implementation, agents have the ability to sign a subset of a web page, and exclude content for which the agent does not claim any responsibility. For example, an agent can sign a document while excluding any ads which are being served alongside the document. Signatures can be applied to anything from an individual hyperlink to an entire document. Signatures can also be applied to text, images, audio, video, or any other digital content. The signature allows anyone to verify that the content that is signed has not been materially altered since the signature was generated.

[0023] Signatures can be portable or fixed to a particular web page or uniform resource locator (URL). For example, a syndicated columnist may wish to sign a column once upon creation, and have the signature follow the document wherever it is published. In other cases, the agent signing the content may wish to prevent their reputation from being used to draw traffic to sites they do not control. In either instance, the metadata associated with the digital signature can indicate whether or not the reputation associated with the signing agent is portable or not. For example, in one implementation, the signature is linked to the URL of the site where the content is located by including the URL as metadata within the signed content.

[0024] In one implementation, multiple agents can sign content on a single web page. For example, a message board or web log can allow each post to be signed by its respective author. In an alternative implementation, unsigned content pieces can be attributed to a synthetic agency identified by the host, site, or URL on which the content piece appears. In another implementation, a single agent that controls multiple websites can sign the content on each of the multiple

website, indicating that the single agent is responsible for the content on all of the multiple websites.

[0025] Each digital signature is tied to the content piece that is signed. In one implementation, the digital signature can be appended to the content piece, or otherwise located in immediate proximity to the content piece. In another implementation, the content piece can contain a link to the digital signature, e.g., a uniform resource identifier (URI) identifying the digital signature. In yet another implementation, the digital signature is located in a central file or directory separate from the content piece, and some portion or all of the content piece covered by the digital signature is the target of a link from the central file or directory. In any of these implementations, the digital signature can be used to verify that the content piece has not been modified since the content piece was signed by the agent.

[0026] The digital signatures can be used to influence the ranking of web search results by indicating the agent responsible for a particular content piece. In one implementation, the reputation for an agent is expressed as a numerical score. A high reputational score indicates that the agent has an established positive reputation. The reputational scores of two or more agents can be compared, and the agent having the higher reputational score can be considered to be more authoritative. In an alternative implementation, multiple scores can be computed for different contexts. For example, an agent might have a first score for content that the agent has written, and a second score for content that the agent has reviewed. In another example, an agent that is responsible for an entertainment magazine could have a high reputation score for content related to celebrity news, but a low reputation score for content related to professional medical advice.

[0027] Assuming that a given agent has a high reputational score, representing an established reputation for authoring valuable content, then additional content authored and signed by that agent will be promoted relative to unsigned content or content from less reputable agents in search results. Similarly, if the signer has a large reputational score due to the agent having an established reputation for providing accurate reviews, the rank of the referenced content can be raised accordingly.

[0028] A high reputational score need not give an agent the ability to manipulate web search rankings. In one implementation, reputational scores are relatively difficult to increase and relatively easy to decrease, creating a disincentive for an agent to place its reputation at risk by endorsing content inappropriately. Since the signatures of reputable agents can be used to promote the ranking of signed content in web search results, agents have a powerful incentive to establish and maintain a good reputational score.

[0029] In one implementation, an agent's reputation can be derived using a relative ranking algorithm, e.g., Google's PageRank as set forth in U.S. Pat. No. 6,285,999, based on the content bearing the agent's signature. Using such an algorithm, an agent's reputation can be determined from the extrinsic relationships between agents as well as content. Intuitively, an agent should have a higher reputational score, regardless of the content signed by the agent, if the content signed by the agent is frequently referenced by other agents or content. Not all references, however, are necessarily of

equal significance. For example, a reference by another agent with a high reputational score is of greater significance than a reference by another agent with a low reputational score. Thus, the reputation of a particular agent, and therefore the reputational score assigned to the particular agent, should depend not just on the number of references to the content signed by the particular agent, but on the importance of the referring documents and other agents. This implies a recursive definition: the reputation of a particular agent is a function of the reputation of the content and agents which refer to it.

[0030] In this manner, the reputation of a particular agent can be calculated by an iterative procedure on a linked database. A linked database (i.e. any database of documents containing mutual citations, such as the world wide web or other hypermedia archive, a dictionary or thesaurus, and a database of academic articles, patents, or court cases) can be represented as a directed graph of N nodes, where each node corresponds to an agent along with all of the content pieces associated with that agent, and where the directed connections between nodes correspond to links from a content piece of one agent to a content piece of another agent. A given node has a set of forward links that connect it to children nodes, and a set of backward links that connect it to parent nodes.

[0031] FIG. 3 illustrates a linked database 300. A first agent 310 is associated with content pieces 312, 314, a second agent 320 is associated with content pieces 322, 324, and a third agent is associated with content piece 332. Content piece 312 associated with the first agent 310 is linked (350) to content piece 322 associated with the second agent 320, and content piece 324 associated with the second agent 320 is linked (352) to content piece 332 associated with the third agent 330. Content piece 314 associated with the first agent 310 is linked (354) to content piece 332 associated with the third agent 330, in addition to content piece 332 being linked (356) back to content piece 314.

[0032] In this implementation, the rank of a particular agent A , $r(A)$, is calculated as follows:

$$r(A) = \frac{\alpha}{N} + (1 - \alpha) \left(\frac{r(B_1)}{|B_1|} + \dots + \frac{r(B_n)}{|B_n|} \right),$$

where B_1, \dots, B_n are the agents that link to A , $r(B_1), \dots, r(B_n)$ are their ranks, $|B_1|, \dots, |B_n|$ are the number of forward links in content signed by the agent, α is a constant in the interval $[0, 1]$, and N is the total number of agents in the database. The constant α is interpreted as the probability that a user will jump randomly to any content piece instead of following a forward link.

[0033] In an alternative implementation, a seed group of trusted agents can be pre-selected, and the agents within this seed group can endorse other content. Agents whose content receives consistently strong endorsements can gain reputation. In either implementation, the agent's reputation ultimately depends on the quality of the content which they sign.

[0034] In another implementation, a set of trusted signing authorities can make additional assertions such as establish-

ing the time when content was signed. This would allow priority to be determined if two agents attempted to sign similar content.

[0035] The use of digital signatures permits the reputation system to link reputations with individual agents, and adjust the relative rankings based on all of the content each agent chooses to associate itself with, no matter the location of the content. For example, the content can be located across multiple websites, or mixed with the content of other agents on a single website. In another implementation, the content can include any sort of digital content, e.g., e-mail, CD-ROMs, or DVDs, and the content need not be located on the Internet. In addition, although the use of digital signatures permits signed content to be associated with a specific agent, it is not necessary to know the actual identity of the agent. Although each agent uses a private key unique to the agent to create each signature, no personal information about the agent is necessary for the signature to be created or for the signature to be used by others. In one implementation, the agent can revoke or otherwise invalidate the private key if the private key is compromised. Once the private key is revoked, the signatures created using the revoked private key will not be accepted as valid, and will not be used to link the agent with the content signed with the revoked private key.

[0036] FIG. 2 shows the use of a content authoring tool to generate and manage the public and private keys necessary to generate signatures. The content authoring tool receives one or more content pieces from an agent (step 210). The content pieces can include text, images, audio, video, or any other static digital content. The content pieces can represent an entire web site, an individual web page, or individual components of a web page. The content authoring tool then receives login credentials or other identifying information from the agent (step 220). The login credentials or other identifying information uniquely identify the agent. In one implementation, personal information can be associated with the login credentials or other identifying information, e.g., billing information. In an alternative implementation, no personal information is associated with the login credentials or other identifying information.

[0037] Next, the content authoring tool determines if the agent already has a public/private key pair for generating digital signatures (step 230). If the agent does not have a public/private key pair, a key pair is generated for use by the agent (step 240). In one implementation, the key pair is generated by an authentication service upon the request of the content authoring tool. In either case, the private key associated with the agent is used to create a digital signature for each of the content pieces (step 250). In one implementation, various metadata can be associated with the digital signature, such as a timestamp indicating the time and date that the digital signature was created, keywords relating to the content piece, or the URL of the website where the content piece is located. In one implementation, the metadata is appended to the content piece, and the content piece containing the metadata is digitally signed.

[0038] In one implementation, the content authoring tool can also be used to determine if the signature associated with a content piece is valid, and identify the agent that signed the content piece in question. Alternatively, any interested entity can use the public key portion of the public/private key pair

to determine if the signature associated with a content piece is valid, and identify the agent that signed the content piece in question.

[0039] In another implementation, the content authoring tool can be used by an agent acting in an editorial or reviewing role to digitally sign a content piece as having been edited or reviewed. The process is similar to that shown in FIG. 2 and described above, with the addition that the digital signature created for each content piece includes metadata indicating that the agent edited or reviewed the digitally signed content piece. In one implementation, the metadata can also include a listing of edits performed or review score.

[0040] Embodiments of the invention and all of the functional operations described in this specification can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structural means disclosed in this specification and structural equivalents thereof, or in combinations of them. Embodiments of the invention can be implemented as one or more computer program products, i.e., one or more computer programs tangibly embodied in an information carrier, e.g., in a machine-readable storage device or in a machine-readable propagated electromagnetic signal, for execution by, or to control the operation of, data processing apparatus, e.g., a programmable processor, a computer, or multiple processors or computers. A computer program (also known as a program, software, software application, or code) can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program does not necessarily correspond to a file. A program can be stored in a portion of a file that holds other programs or data, in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network.

[0041] The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

[0042] Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a processor for executing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or

optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio player, a Global Positioning System (GPS) receiver, to name just a few. Information carriers suitable for embodying computer program instructions and data include all forms of non volatile memory, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0043] To provide for interaction with a user, embodiments of the invention can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input.

[0044] Embodiments of the invention can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the invention, or any combination of such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), e.g., the Internet.

[0045] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0046] Particular embodiments of the invention have been described. Other embodiments are within the scope of the following claims. For example, actions recited in the method claims can be performed in different orders and still achieve desirable results.

What is claimed is:

- 1. A method comprising:
 - receiving multiple content items from a corpus of content items;
 - receiving digital signatures each made by one of multiple agents, each digital signature associating one of the agents with one or more of the content items;
 - assigning a score to a first agent of the multiple agents, wherein the score is based upon the content items associated with the first agent by the digital signatures.

- 2. The method of claim 1, further comprising:
 - determining the validity of the digital signatures.
- 3. The method of claim 1, wherein if no digital signature associates an agent with a specific content item, the content item is associated with an owner of a location where the specific content is found and a score is assigned to the owner based on the specific content item.
- 4. The method of claim 1, wherein the content items associated with the first agent include a content item that contains a digital signature associating the first agent with the content item.
- 5. The method of claim 1, wherein the content items associated with the first agent include a content item that includes a link to a digital signature associating the first agent with the content item.
- 6. The method of claim 1, wherein the content items associated with the first agent include a content item that is a web site or a portion of the web site.
- 7. The method of claim 1, wherein a second agent is associated by a second digital signature with a second content item with which the first agent is associated, and the second agent makes an assertion about the content item.
- 8. The method of claim 7, wherein the first agent makes an assertion with a digital signature that the first agent is an author of the second content item.
- 9. The method of claim 7, wherein the second agent makes an assertion with the second digital signature that the second agent is a reviewer, editor, or publisher of the second content item.
- 10. The method of claim 1, wherein assigning a score to the first agent further comprises assigning the score based on unsigned content items that the first agent is associated with as an owner of one or more locations where the unsigned content items are found.
- 11. The method of claim 1, wherein assigning a score to the first agent further comprises assigning the score based on one or more assertions made by one or more other agents about content items associated with the first agent.
- 12. The method of claim 1, wherein assigning a score to the first agent further comprises assigning the score based on one or more assertions made by the first agent.
- 13. The method of claim 1, further comprising:
 - using the score in ordering results of a search of the corpus.
- 14. A computer program product, tangibly embodied in an information carrier, consisting of instructions operable to cause data processing apparatus to:
 - receive multiple content items from a corpus of content items;
 - receive digital signatures each made by one of multiple agents, each digital signature associating one of the agents with one or more content items;
 - assign a score to a first agent of the multiple agents, wherein the score is based upon the content items associated with the first agent by the digital signatures.
- 15. The computer program product of claim 14, wherein if no digital signature associates an agent with a specific content item, the content item is associated with an owner of a location where the specific content is found and a score is assigned to the owner based on the specific content item.
- 16. The computer program product of claim 14, wherein the content items associated with the first agent include a

content item that contains a digital signature associating the first agent with the content item.

17. The computer program product of claim 14, wherein the content items associated with the first agent include a content item that includes a link to a digital signature associating the first agent with the content item.

18. The computer program product of claim 14, wherein the agent associated with the content item makes an assertion about the content item.

19. The computer program product of claim 14, wherein the instructions to assign a score to the first agent further comprise instructions to assign the score based on one or more assertions made by one or more other agents about content items associated with the first agent.

20. The computer program product of claim 14, further comprising instructions to:

use the score in ordering results of a search of the corpus.

* * * * *