

The Fractal Nature of the Semantic Web

Tim Berners-Lee and Lalana Kagal

■ *In the past, many knowledge representation systems failed because they were too monolithic and didn't scale well, whereas other systems failed to have an impact because they were small and isolated. Along with this trade-off in size, there is also a constant tension between the cost involved in building a larger community that can interoperate through common terms and the cost of the lack of interoperability. The semantic web offers a good compromise between these approaches as it achieves wide-scale communication and interoperability using finite effort and cost. The semantic web is a set of standards for knowledge representation and exchange that is aimed at providing interoperability across applications and organizations. We believe that the gathering success of this technology is not derived from the particular choice of syntax or of logic. Its main contribution is in recognizing and supporting the fractal patterns of scalable web systems. These systems will be composed of many overlapping communities of all sizes, ranging from one individual to the entire population that have internal (but not global) consistency. The information in these systems, including documents and messages, will contain some terms that are understood and accepted globally, some that are understood within certain communities, and some that are understood locally within the system. The amount of interoperability between interacting agents (software or human) will depend on how many communities they have in common and how many ontologies (groups of consistent and related terms) they share. In this article we discuss why fractal patterns are an appropriate model for web systems and how semantic web technologies can be used to design scalable and interoperable systems.*

Many patterns in nature—including clouds, mountains, leaves, and even language—are fractal, in that they have repeating structures at every scale and cannot be modeled with classical geometry (Mandelbrot 1977). Fractals provide solutions to capture and think about complex and irregular patterns. They can be used to solve a variety of problems including predicting weather, designing antennae, and even characterizing metals. Identifying that a pattern is fractal helps us model and analyze it appropriately. The inherent fractal nature of language and culture in human societies leads us to expect the semantic web to demonstrate the self-similar patterns of fractals (Berners-Lee 1998).

An ontology is a description of the concepts and relationships in a domain that is used for modeling and sharing domain knowledge (Gruber 1995). Ontologies allow agents both within and outside a system to have a common understanding of information used in the system and promote interoperability between systems. Large, monolithic ontologies such as Cyc (Lenat 1995) and Suggested Upper Merged Ontology (SUMO)¹ were developed using classical knowledge representation theory. These kinds of large ontologies are usually developed by a small group of people with a lot of time and effort and try to meet the requirements of a much larger group of users. When adopted, they provide greater interoperability between different groups. At the other end of the scale are small, isolated (stovepipe) systems that are developed by larger sets of software engineers.

Desktop applications are an example of these kinds of systems. Each application usually has a specific purpose such as scheduling appointments or analyzing bank statements and cannot reuse information produced by other applications. For example, it is not possible to have a calendar view of your bank statements or to put your pictures into your scheduler. From these two extremes, it is clear that there is a trade-off between size and effort versus reuse and interoperability as illustrated in figure 1. Neither of these approaches is alone useful on the web, and we believe that the solution lies somewhere in between these extremes.

It appears that human society is fractal in that human groups are stable when they have a set of peers and when they have a substructure (Kleinberg 1999). Groups that have a very large number of peers or that are composed of a set of a very large number of subgroups cannot work effectively because of the cost of communication and interoperability. Each subgroup within a stable group develops (or adopts) its own language but needs to share some terms with every subgroup or set of subgroups that it has to work with. As these subgroups overlap, several common shared terms emerge that are understood by everyone in the group.

In a similar manner, we believe that web systems will be fractal and be composed of overlapping communities of all sizes. Each community will have its own ontologies but will also develop shared ontologies with communities it interacts with. Some basic ontologies will be used globally by all communities. Communi-

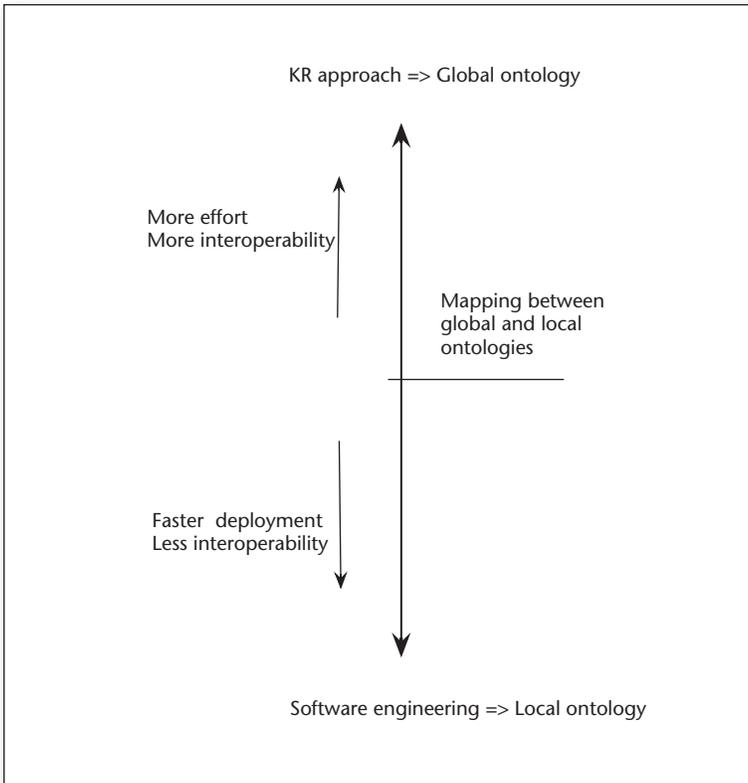


Figure 1. Global Versus Local Ontologies.

cation between agents either in the same or different community will contain terms from both global and local ontologies. The amount of data that will be understood and reused by interacting agents will depend on how many communities they have in common and how many ontologies they share. Semantic web technologies support this kind of ontology development and use. Terms are defined in ontologies and ontologies are defined by communities. A person can be involved in many communities, a message can mix terms from many ontologies, and an operation only requires consistency between parts of ontologies that are in use for that particular operation. This promotes greater harmonization, but does not require the establishment of a global ontology of everything. A single huge ontology of everything is difficult to accomplish, as the effort of getting consensus on it becomes unimaginable. On the other hand, stovepipe systems with only local ontologies lead to interoperability problems and a lack of reuse of data. The fractal patterns of web systems allow a compromise between having large global ontologies and small isolated ones.

In the following sections, we discuss the fractal nature of the semantic web and how we can exploit semantic web technologies to develop effective and interoperable web systems.

Real World Example: Fractal Topology

As an example of a fractal topology of communities, consider a product label as shown in figure 2. The label consists of the bar code for identifying the product, the nutrition section that provides information about the serving size, calorie contents, and the ingredients, and the product logo and name. Each part of this label is generated and used by a different community of users. The bar code is used by the store for pricing and inventory and has to be in one of several universally accepted formats such as Bookland EAN Bar Code or EAN-5. The format and specifications for the nutrition information are provided by the U.S. Food and Drug Administration. The logo and name of the product is from the producer's domain, whereas the series of numbers on the right side are numbers that are printer specific and will not even be visible when the label is pasted on the container. Each section of the label is from a different community and uses a different set of concepts or ontology. However, the label can be used by all communities without misunderstandings because the portions of the label that are irrelevant to a community can be dropped without causing an inconsistency in the relevant portion. Each portion of the label also contains within it an implicit reference to the domain from which the terms were drawn making it easy to ground the terms.

Culture, Boundaries, and the Web

When a group of people communicate among themselves, they develop, to a certain extent, their own language. Sometimes, they pick terms understood by one party, talk enough to develop a shared understanding of the meaning of the term, and adopt it across the group. Sometimes, as discussion proceeds within the group, meanings are adjusted so that they can be used for new concepts that are created or discovered by the group's activity. At other times a group will deliberately and quite specifically make up a new term that is different from any other word or phrase used before. While this is evident in technical groups, this process also happens in all walks of life, legal and political, as well as social and familial.

The result of this process is a new language, a new strain of language, or just a twist in the use of an existing word. The motivating factor is to enable communication within the group. A greater shared vocabulary broadens the scope of common discussion that can be made without misunderstanding. The second, complementary effect of this change is to create a common bond within the group, which at the same time, erects a barrier around the group. In most cases, this is uninten-

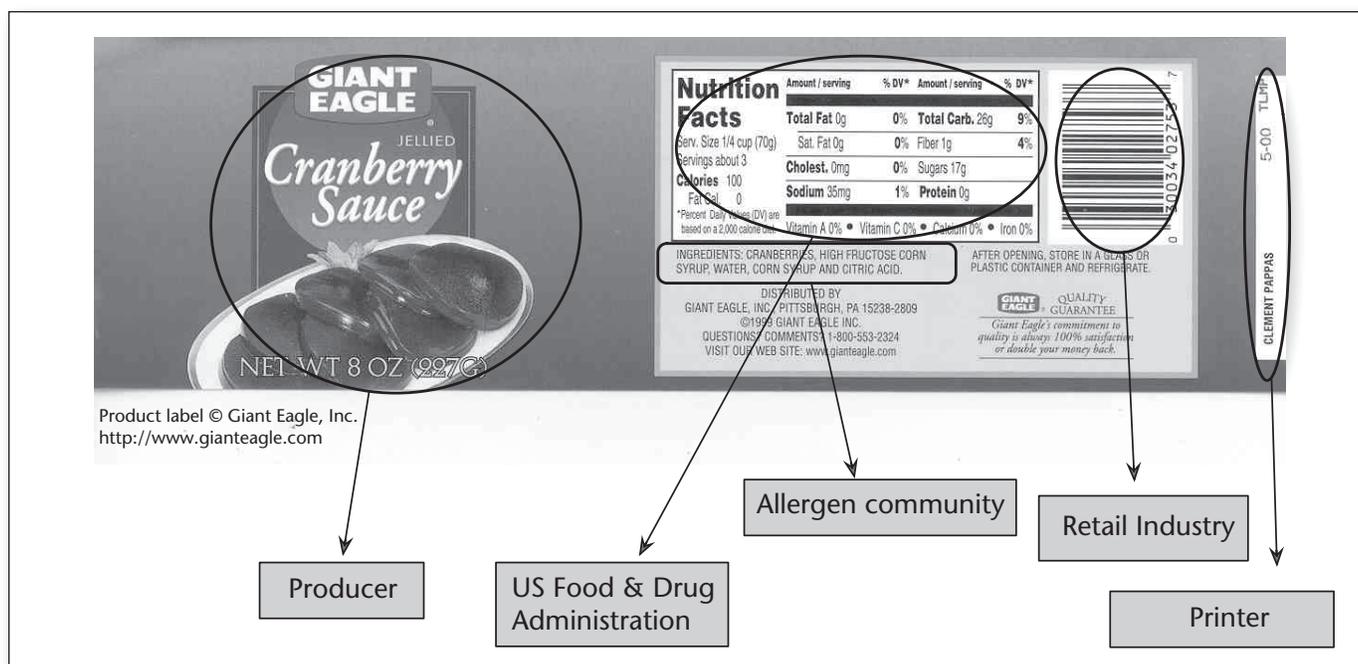


Figure 2. Product Label That Contains Data from Different Ontologies

tional. For every linguistic development that promotes communication within the group, a corresponding step change is made in the difficulty of communicating across the boundary of the group. In order to get wider interoperability, an essential solution is for those involved to consider what those on the other side of the boundary are thinking. In a conversation, this is the job of listening. In a technical setting, it can involve a careful study of the words used in the other's seemingly senseless protestations to logically build up a conclusion of how those words must be related in the other's mind.

There is constant tension between the need to get things done quickly with less effort, by working within a small group, and the need to get a wider understanding between different groups, which takes so much more time. In practice, life is made up of a fractal tangle of overlapping communities, of overlapping cultures. This means that the tension is ever present. It also means that there is always a small amount of common language shared by a very large number of people, a large number of concepts local to an individual, and everything in between. In centuries before this one, geography played an important role in constraining group development, and so a nested two-dimensional pattern existed.

With the Internet and the web, we can connect things without the constraint of these nested geographical areas. We can choose not just to be members of communities such as town, region, state, and country but also to be specialists in a given

field, people with a particular medical condition, or people concerned about a particular global issue. This means that the topology of the communities, and the connectedness by some metrics, may be different and in fact better than before. The topology that emerges on the web will depend on the individual choices of many people, but in accordance with previous studies of web usage, we believe it will be a fractal distribution, emphasizing all scales.

Ontology Development and Usage on the Semantic Web

Starting from the early days of the web, several studies have observed that web use follows a Zipf or fractal distribution of popularity where a small number of websites account for most of the web traffic (Nielsen 1997, Menasce et al. 2002, Dill et al. 2002, Shirky 2003). The analysis of current ontology use has shown that it also follows a similar fractal distribution. Figure 3a illustrates the usage of ontologies (approximately 13,675) based on how many of the total semantic web documents (approximately 2,379,164) they are used in. The graph shows that only 10.4 percent of the ontologies have been used in at least 10 documents and less than 1 percent (about 10 ontologies) have been used in over 100,000 documents. According to Swoogle (Ding and Finin 2006), the most used ontology is Dublin Core,² which has been used at least once in more than one million documents on

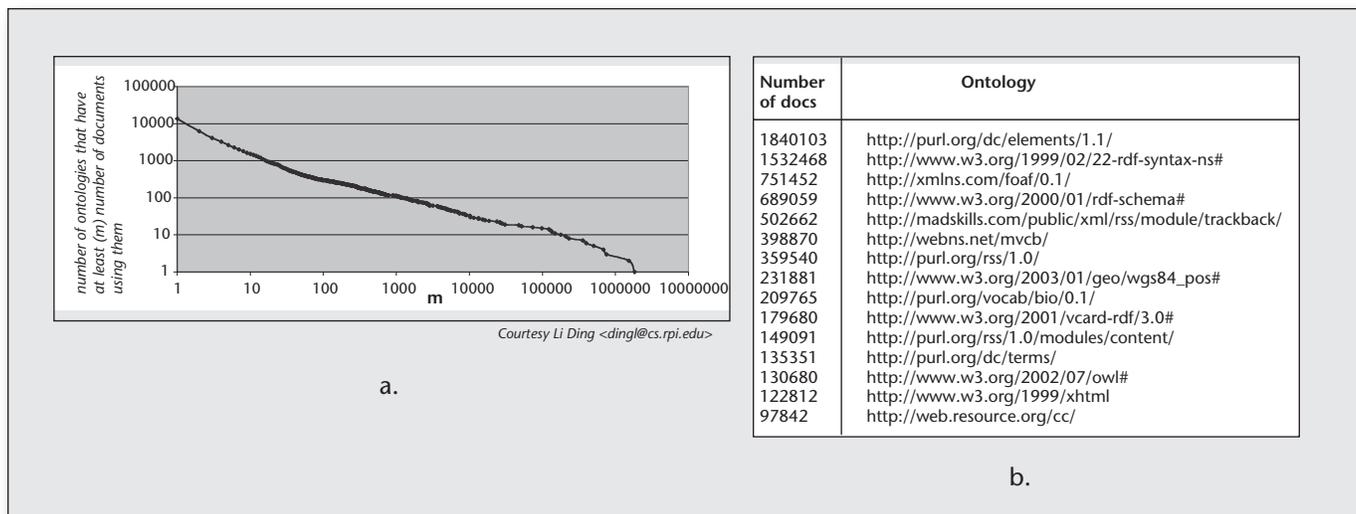


Figure 3. Fractal Use of Ontologies on the Semantic Web.

a. Ontology usage in 2007 as observed by Swoogle. b. Most commonly used ontologies and the number of semantic web documents they are used in.

the semantic web. The next most popular ontologies include Resource Description Framework,³ Friend Of A Friend (FOAF),⁴ RDF Schema,⁵ and Trackback of RSS.⁶ Please refer to figure 3b for a list of the most used ontologies and the number of documents they have been used in.

The cost of ontology development is often miscalculated. Most system developers are under the misconception that the ontologies required by the system will be either all developed by a top-down or a bottom-up approach. In a top-down approach, all ontologies will be created by standards bodies forcing system developers to wait for standard ontologies to be developed that meet all their requirements. In a bottom-up approach all the information required by the system is created by the system developers, making it costly and time consuming. However, in reality, the fractal nature of the semantic web leads to a cost-effective solution, which is a combination of these approaches. Global ontologies such as Dublin Core, Cal⁷, and Geo⁸ are created and maintained by standards bodies; community ontologies are created by groups of organizations that need those ontologies so the cost is shared; and local ontologies that are specific to the system are created by system developers. This implies that though system developers are mainly responsible for their local ontologies, it is beneficial for them to participate in the development of both standard and community ontologies.

Design Considerations for Semantic Web Technologies

Given the natural overlapping of human groups, it is important that semantic web technologies sup-

port a similar overlapping in web systems. These technologies need to enable the development of a global community of distributed but interconnected and interoperable information systems. These technologies should (1) provide global unique identification so users can uniquely identify terms and resources from different communities, (2) allow the free mixing of terms from different ontologies as several ontologies (local and global) will be used for develop knowledge or messages, (3) support extensibility so individuals can create new terms or ontologies without waiting for consensus from a governing body, (4) allow portions of ontologies or data to be dropped or ignored without affecting the meaning of the other concepts, and (5) support mapping between ontologies for communities that use different ontologies but need to communicate.

Semantic web technologies such as the Resource Description Framework vocabulary language (RDFS) (Brickley and Guha 2002) and the Web Ontology Language (OWL) (Bechhofer et al. 2004) provide the required support through the following characteristics:

First, while communicating within a local group, users can use concepts from an agreed upon, shared set of terms. However, for interoperability across diverse communities, a single global identification system is required (Jacobs and Walsh 2004). Uniform resource identifiers (URIs) (Berners-Lee, Fielding, and Masinter 2004) allow web resources and terms to be unambiguously identified. For example, the Person concept defined by the <http://xmlns.com/foaf/0.1/Person> URI is not the same as the Person concept defined by the <http://www.w3.org/2000/10/swap/pim/contact#Person> URI. In case of web resources, the Hypertext

Transfer Protocol (HTTP) provides a way to associate a URI and the resource it represents (also called dereferencing). URIs and HTTP form the foundation of semantic web technologies.

Users can refer to and use terms from different ontologies through groups of URIs known as namespaces. Namespaces provide a way to qualify terms by associating them with URI references. Consider as an example foaf:knows. If the foaf namespace refers to <http://xmlns.com/foaf/0.1/> then foaf:knows is a property defined in the foaf namespace.

The web architecture supports the easy creation of new communities by enabling them to use the Domain Name System (DNS) to register their domains and to host their community information including their ontologies.

Semantic web technologies provide simple mechanisms through RDF and RDFS that allow individuals to easily create new ontologies or add to or modify existing ontologies. For example, in order to create a new concept one would use `rdf:Class` and to create a new property one would use `rdf:Property`. New ontologies do not have to be published on a specific website or at a certain URI in order to be used. Also, modifications to existing ontologies do not have to be published at the same URI as the original ontology. For example, it is possible to add properties to the Person concept defined in the foaf namespace and publish the changes on a different URI. Users can add to or modify terms in existing ontologies, publish the changes to their own URIs, and use them immediately. In general, users do not need to wait for consensus from the community or any governing body to create or modify ontologies.

Semantic web data is organized as RDF graphs, which are made up of triples (subject predicate object). These triples are only conjuncted. This means that if someone doesn't want to use or does not understand a portion of an RDF graph, it can be easily ignored or dropped. Assume an agent accesses a document that contains terms from two ontologies, only one of which it understands. In this case, the agent can ignore the terms from the other ontology and be assured that the remaining portion of the document will be consistent. If RDF graphs were made of both conjunction and disjunction of triples, it would have been difficult to drop the irrelevant portions.

For mapping between terms (both concepts and properties) in different ontologies, OWL provides properties such as `SameAs`, `equivalentProperty`, and `differentFrom`. For example, it is possible to say that `tel` defined in the `vcard` ontology is an `owl:equivalentProperty` to `phone` defined in the foaf ontology. As a consequence, when a document uses `vcard:tel` to state someone's telephone number, it would be understood to mean the same

as `foaf:phone`. RDF has properties that deal with classification such as `subClassOf` and `subPropertyOf` that can also be used to describe the relationship between terms in different ontologies.

Principles for Developing Ontologies and Applications

In this section we suggest some principles for developing scalable semantic web communities.

Ontology Development

It is important to be aware of the community and to use terms from ontologies that have been developed by existing communities. For example, if you are working on a scout troop ontology then you should consider what terms are commonly used by the troop members and whether there are existing troop ontologies already in use.

You should only define terms that are specific to your community of interest and try to leave terms specific to other subcommunities to be defined in those communities. For example, if you are developing a scout troop ontology, you should not try to develop terms that a girl scout troop might require.

You should publish the terms you define on a community web page and demonstrate the trustworthiness of your ontology by maintaining the URIs.

Client Cache

In client-side applications, global ontologies should be cached locally not only because they are accessed often but also to prevent stressing the server that is hosting the ontology; shared/community ontologies should be stored in a persistent store; and local or working ontologies should be read at run time because they might change often.

User Interface

To support fractal communities, special attention should be paid to user interface design, prompting and helping users make the right decisions. As an example, consider how RDF graphs are edited in the Tabulator (Berners-Lee et al. 2006). While changing or adding a predicate, users are first prompted with predicates from global ontologies such as `Friend Of A Friend`. Users are then shown community or domain-specific terms such as those that are specified in a configuration file. After this they have a choice of nearby data such as data that has just been added to the Tabulator. Finally if none of the available data meet their requirements, they can create their own terms. This emphasizes the reuse of existing ontologies and helps create interoperability.

Functionality

Different kinds of ontologies require different functionality. Global ontologies should have specific code because they are used frequently. Domain- or community-specific ontologies should have adaptable code such as plug-ins that can be downloaded as and when required. Local or specific ontologies can have basic functionality such as being viewable in a spreadsheet. For example, within an application, there should be specific code for handling FOAF ontologies, code for handling the cranberry sauce ontology should be downloadable from the cranberry sauce community web page, and information such as printer information on the cranberry sauce label should only be text so it can be cut and pasted into a spreadsheet, if required.

Summary

Human society is made up of a fractal tangle of overlapping communities and cultures. We expect the same fractal patterns to appear in scalable web systems within which information will be composed of terms from different ontologies—global, community specific, and local. There will be some global shared ontologies such as iCal and Geo, but most of the ontologies will be established by smaller communities of different sizes such as a cranberry sauce ontology or a girl scout troop ontology. Semantic web technologies will help achieve scalable and interoperable systems with finite cost and effort by leveraging this fractal distribution and emphasizing global identifiers and the reuse and extensibility of ontologies.

Notes

1. www.ontologyportal.org.
2. purl.org/dc/elements/1.1.
3. www.w3.org/1999/02/22-rdf-syntax-ns.
4. xmlns.com/foaf/0.1.
5. www.w3.org/2000/01/rdf-schema.
6. madskills.com/public/xml/rss/module/trackback.
7. www.w3.org/2002/12/cal/ical.
8. www.w3.org/2003/01/geo/wgs84_pos.

References

- Bechhofer, S.; van Harmelen, F.; Hendler, J.; Horrocks, I.; McGuinness, D. L.; Patel-Schneider, P. F.; and Stein, L. A. 2004. OWL Web Ontology Language Reference, W3C Recommendation 10 February 2004. Cambridge, MA: W3 Consortium (www.w3.org/TR/owl-ref/).
- Berners-Lee, T. 1998. The Fractal Nature of the Web (www.w3.org/DesignIssues/Fractal.html).
- Berners-Lee, T.; Chen, Y.; Chilton, L.; Connolly, D.; Dhanaraj, R.; Hollenbach, J.; Lerer, A.; and Sheets, D. 2006. Tabulator: Exploring and Analyzing Linked Data on the Semantic Web. Paper presented at the 3rd International Semantic Web User Interaction Workshop, International

Semantic Web Conference, Athens, GA, 6 November.

Berners-Lee, T.; Fielding, R.; and Masinter, L. 2004. Uniform Resource Identifiers (URI): Generic Syntax, IETF Standard (gbiv.com/protocols/uri/rev-2002/draft-fielding-uri-rfc2396bis-07.html).

Brickley, D., and Guha, R. 2002. RDF Vocabulary Description Language 1.0: RDF Schema, W3C Recommendation. Cambridge, MA: W3 Consortium (www.w3.org/TR/rdf-schema).

Dill, S.; Kumar, R.; Mccurley, K. S.; Rajagopalan, S.; Sivakumar, D.; and Tomkins, A. 2002. Self-Similarity in the Web. *ACM Transactions on Internet Technology* 2(3): 205–223.

Ding, L., and Finin, T. 2006. Characterizing the Semantic Web on the Web. In *Proceedings of the 5th International Semantic Web Conference*. Berlin: Springer.

Gruber, T. 1995. Toward Principles for the Design of Ontologies Used for Knowledge Sharing. *International Journal of Human-Computer Studies* 43(4–5).

Jacobs, I., and Walsh, N. 2004. Architecture of the World Wide Web, Volume One, W3C Recommendation. Cambridge, MA: W3 Consortium (www.w3.org/TR/webarch).

Kleinberg, J. 1999. The Small-World Phenomenon: An Algorithmic Perspective. Cornell Computer Science Technical Report 99-1776, Ithaca, NY.

Lenat, D. 1995. Cyc: A Large-Scale Investment in Knowledge Infrastructure. *Communications of the ACM* 38(11).

Mandelbrot, B. 1977. *Fractals: Form, Chance, and Dimension*. San Francisco: W. H. Freeman and Company.

Menasce, D. A.; Abrahao, B. D.; Barbara, D.; Almeida, V.; and Ribeiro, F. 2002. Fractal Characterization of Web Workloads. Paper presented at the Eleventh International World Wide Web Conference, Honolulu, HI, 7–11 May.

Nielsen, J. 1997. Zipf Curves and Website Popularity. Alertbox: Current Issues in Web Usability (www.useit.com/alertbox/zipf.html).

Shirky, C. 2003. Power Laws, Weblogs, and Inequality. Clay Shirky's Writings about the Internet (www.shirky.com/writings/powerlaw_weblog.html).

Tim Berners-Lee, a graduate of Oxford University, England, is now with the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology (MIT). He directs the W3 Consortium, an open forum of companies and organizations with the mission to realize the full potential of the web. With a background of system design in real-time communications and text processing software development, in 1989 he invented the World Wide Web, an Internet-based hypermedia initiative for global information sharing, while working at CERN, the European Particle Physics Laboratory. He wrote the first web browser and server in 1990.

Lalana Kagal is a research scientist in the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. Her primary research focus is in usable behavior and knowledge management through policy application in open, dynamic, and distributed environments. After receiving Ph.D. and MS degrees in computer science from the University of Maryland, Baltimore County, she joined MIT as a postdoctoral associate. She has authored more than 45 peer-reviewed publications and contributes actively to the research community.