Interoperability as Desideratum, Problem, and Process

Gary Richmond

City University of New York
garyrichmond@rcn.com

Abstract. In a global environment more and more influenced by the use of Internet technology in distributed settings, issues of interoperability have become crucial for business and governments, and in fact for all individuals and organizations employing it and increasingly dependent on it. We analyze three inter-related but distinct levels of interoperability, the syntactic, semantic and pragmatic, and discuss some of the interoperability issues related especially to the semantic and pragmatic levels. We briefly look at the relationship between philosophical Ontology and ontology as the term is used in AI, suggesting that especially the scientific metaphysics (Ontology) of Charles Peirce might be of value in helping KR workers generally and ontologists in particular to uncover possibly hidden ontological commitments. We then consider John Sowa's Unified Framework (UF) approach to semantic interoperability which recommends the adoption of the draft ISO standard for Common Logic, and touch upon how CGs could be effectively employed within this framework. The promise of Service Oriented Architectures (SOA) in consideration of real world applications is remarked, and it is suggested how in a fast paced and continuously changing environment, loose coupling may be becoming a necessity—not merely an option. We conclude that, while the future of network interoperability is far from certain, our communities at least can begin to act in concert within our particular fields of interest to come to agreement on those "best in class" theories, methods, and practices which could become catalysts for bringing about the cultural change whereas interoperability, openness, and sharing are seen as global desiderata.

Introduction

As Web-based computing increasingly influences the communication of individuals, teams, and organizations, the complexity of the problem of access to relevant information seems itself to be expanding exponentially. In this post-modern environment of diverse and dynamic distributed information sources a variety of interoperability issues was certain to follow just from the extraordinary complexity of the almost "too rich" multicultural, cross-disciplinary, and indeed global environment we seem to be evolving

Anyone following the development of the WWW since the turn of the new millennium as it relates to business, government, education, and other cultural enterprises has surely observed not only that interoperability has become a growing concern for experts in those fields specifically related to its development, but also that, in several of its expressions (and of course interoperability means many things to many people), the 'interoperability question' has implications very generally: for business, of course, but also for many political, economic, and other power relations. For example, the case of the European Union Commission (EC) *versus* Microsoft illustrates rather strikingly the high-stakes economic ramifications of certain interoperability issues. The core argument of the EC has been that Microsoft has abused its power by restricting interoperability with non-MS products and thereby achieving corporate dominance for essential IT network server operating systems. Indeed the EC's tentative solution to the problem has been termed the "interoperability remedy" [33]. It would appear, however, that the actual application of the "remedy" is still very much a work in progress [2]. This seems not at all surprising since the emerging goals of openness, access, and sharing on the one hand—epitomized by the philosophy of the open source movement [25]—continues be opposed by very different and even opposite tendencies which valorize the closed and the proprietary.

Nevertheless, not only in business, but also in the cultures of governments, universities, museums, the media, and indeed nearly all large enterprises employing information technology, there would appear to be "a growing recognition of the need for common 'architectures' within which truly useful applications and services can be constructed" [17]. Yet how are we to go about realizing the goal of interoperability within a wide range of contexts and across myriad services given the complexity of the situation?

One important albeit preliminary consideration involves the very definition of 'interoperability'. Indeed there is a broad spectrum of definitions coming from quite diverse perspectives and from quite different needs and purposes so that one finds both technical and technical-semantic as well as more social and pragmatic definitions and discussions of what 'interoperable' means. It is certainly legitimate to define interoperability solely in terms of the technical considerations involved—tags, file formats, protocols, and so forth—when referring to the technical systems providing/accepting services. But it is also possible and increasingly useful to define interoperability more generally in terms of the tools and agents working to optimize the means available for especially social networks to accomplish common (or closely related) tasks and working towards shared goals. And in fact there are an increasing number of discussions of interoperability that can be seen as emphasizing those aspects of the process utilizing the technology and semantics to purposeful ends. The result is that, as clearly essential as technical-semantic interaction certainly is, semantic-social interaction/interoperability is coming to be seen as an equally vital element in the continuing growth of networks as it becomes clear that it is precisely the goals and purposes of users which provide the context of the semantics. The problem is not so much with the formal representations of meaning (mainly ontologies), but how these meanings ought—or at least could—be put to use.

There remains, however, no consensus yet as to exactly what the optimal relationship between the technical, the semantic, and the social and organizational ought to be, how this balance could be achieved, and how technical and semantic concerns themselves might be balanced with individual, corporate, national, and international interests. It is likely that the means to an optimal balance has not been fully achieved in *any* field including those directly involved in the development of tools and others resources bearing upon the development of networks serving enterprises employing these tools and services. But, for now, and to assist us in a preliminary analysis, let's consider somewhat abstractly the general goal of 'interoperability'. Speaking from the standpoint of information management, Paul Miller has written that "to be interoperable, one should actively be engaged in the ongoing process of ensuring that the systems, procedures and culture of an organization are managed in such a way as to maximize opportunities for exchange and re-use of information, whether internally or externally" [18].

From this vantage point the interoperability terrain can be seen to be vast, as indeed it is, involving almost all network technology and its possible use. Consequently in this short essay it will be possible to touch upon only the most general features as we try in particular to contextualize the semantic realm, this contextualization representing the very essence of the pragmatic layer. In section 2 we distinguish three kinds of interoperability, the syntactic, semantic, and pragmatic. In section 3 we consider the philosophical notion of 'ontological commitment' especially as it relates to the semantic level. In section 4 we briefly consider John Sowa's proposal—based on perhaps two decades of sound software engineering principles—to enhance semantic interoperability through a Unified Framework (UF). In section 5 we examine the promising notion of interoperability as it appears in Service Oriented Architecture (SOA) and especially in its principle of 'loose coupling." In section 6 we glance for a moment at the possible future of interoperability.

2. Interoperable Three Ways

A concise definition of 'interoperability' will be helpful at this point. The IEEE defines interoperability as "the ability of two or more systems or components to exchange information and to use the information that has been exchanged" [14]. At the very least this definition will allow us rapid entry into the technical level of interoperability, while it already hints at the meaning and purposeful use of the information-sharing operations and services that are built upon the technology.

If the efficient exchange of accurate and relevant information for specific creative and collaborative purposes is taken as the goal of the Semantic Web, then perhaps its best hope of success lies in encouraging "the emergence of communities of interest and practice that develop their own consensus knowledge" [29]. Network architecture can no sooner be seen to end at the semantic level than Peirce's philosophical pragmatism can be imagined to culminate in his Critical Logic, or in his Existential Graphs (EGs) which John Sowa has transmuted into Conceptual Graphs (CGs) for an information age. Knowledge representation is a means to an end and not the end itself. The 'end itself', at least for such thinkers as Peirce and Engelbart, would be something like the evolution of consciousness itself [CP 6.289, 6.490, 7.515].

At ICCS01 at Stanford in far-ranging informal discussions, Aldo de Moor and I began to imagine that we had conceived the idea of a Pragmatic Web (move over Al

Gore!) Well, perhaps we had if only in the sense that even if we weren't exactly the first to use the phrase "Pragmatic Web"—we later noted that it seemed to be "in the air" at the time—we seem to be the first to associate it with Peircean pragmatism, as we would make explicit in a paper we co-authored with Mary Keeler, "Towards a Pragmatic Web" [6]. Since then de Moor has been in the vanguard of those seeking to legitimize the idea of a Pragmatic Web and realize it in fact. Recently he and his co-chairs announced the First International Pragmatic Web Conference to be held this September in Stuttgart, Germany [28]. Less than a year before this he had presented at ICCS06 his "Patterns for the Pragmatic Web" [5] upon which the immediately following analysis is based. He conceptually divides the World Wide Web into a trichotomy of Webs (following Peirce's semeiotic trichotomy[1]).

> The Syntactic Web consists of interrelated syntactic information resources, such as documents and web pages linked by HTML references. These resources describe many different domains.
>
> The Semantic Web consists of a collection of semantic resources about the Syntactic Web, mainly in the form of ontologies. The ontologies contain semantic networks of concepts, relations, and rules that define the meaning of particular information resources.
>
> The Pragmatic Web consists of a set of pragmatic contexts of semantic resources. A pragmatic context consists of a common context and a set of individual contexts. A common context is defined by the common concepts and conceptual definitions of interest to a community, the communicative interactions

---

[1] Peirce's TRICHOTOMIC *category theory* is discussed in [26] and applied to several analyses here. The author's diagrammatic transmutation of this applied science is outlined in a PowerPoint presentation [27]. A trikon, symbolized |> and looking something like an equilateral "forward" button, gives the categorial structure of a *genuine triadic relationship* of categorial "firstness" "secondness" and "thirdness" (1^{ns} , 2^{ns}, and 3^{ns}) placed around the triangle thus:

1^{ns} [*can* *be* ═ possible]
|> 3^{ns} [*would* necessarily *be* ***if*** x ═ necessary]
2^{ns}[actually *is* ═ actual]

Peirce associated each of the three branches of his science of *logic as semeiotic* with one of the 3 categories: 1^{ns} with theoretical grammar (syntax), 2^{ns} with critical logic (semantics), and 3^{ns} with methodology, theory of inquiry, and the pragmatic maxim (pragmatism):

structure (syntactic)
|> purpose (pragmatic)
meaning.(semantic)

A synthetic extension of this analytic part of trikonic involves what I call *trikonic vector analysis* (or, reflection on the 6 possible paths of movement through the three categories and their possible patterns of inter-relating), a topic of my ICCS06 paper. |>*k vector analysis-synthesis has recently been applied to a software interoperability case [32]. However, a vast amount of work lies ahead in applying Peirce's category theory (not to be confused with mathematical category theory, btw) and especially the vector part of this to knowledge representation and knowledge management (not to mention business applications!)

in which these concepts are defined and used, and a set of common context parameters (relevant properties of concepts, joint goals, communicative situation, and so on). Each community member also has an individual context, consisting of individual concepts, definitions of interest, and individual context parameters. [5]

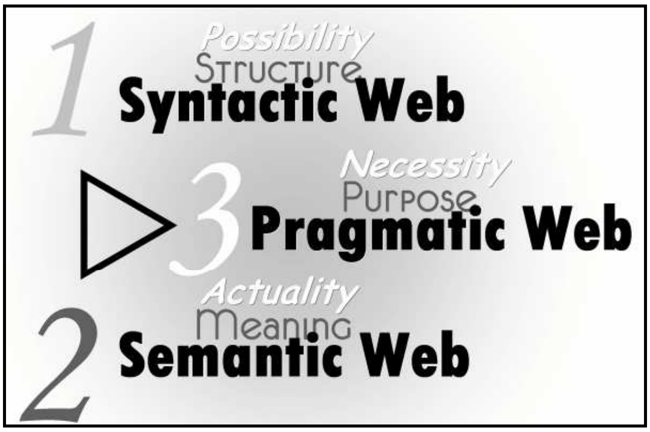

Fig. 1.

There are undoubtedly any number of other ways of conceiving the general structure of the Web[2], but a trichotomic division will be employed here as best representing the deep structural inter-relations of the various levels of networks. The sense that there are indeed several distinct while interpenetrating facets of web development leads us to a consideration of that to which we will refer here as syntactic, semantic, and pragmatic information architectures.

Syntactic architecture: At this level we are concerned with technical considerations such that individual systems are the principal focus. This level details specifications for components of the system as a whole as well as standards and protocols to be used for communication between components. It is not intrinsically concerned with the meaning or purpose of the system, nor of its users and their interactions. In relation to the next level, the semantic, it is concerned with tags as identifiers for any and all purposes (anything can be identified with a URI).

Semantic architecture: Here considerations of meaning and function dominate. It becomes a question of the functions that the system is designed to fulfill that a user or group of users may want to utilize. This level concerns what the tags refer to and how this metadata can be meaningfully related, for example, what ontologies are needed for specific uses. Semantic architecture includes at least the following: "mark up" of documents with information-extending HTML tags, common metadata organization schemata (principally ontologies), automated agents allowing these schemata to be put to use, and web services to supply users with the specific information which they

[2] Miller, for example, identifies six not altogether distinct architectures: Technical, Semantic, Political/Human, Inter-community, Legal and International interoperability. In our analysis, the Pragmatic architecture incorporates the last four of these.

seek. The common denominator here is, of course, that the *meaning of the data* (metadata) is emphasized

Achieving semantic interoperability is proving to be a most difficult challenge for a number of reasons. One of these is illustrated by the UK's High Level Thesaurus Project (HILT) [13] which reveals "a wide range of existing, incomplete, terminologies, category lists, classification schemes and thesauri" which are being used and which users are loathe to give up whatever the limitations of a given schema. This would suggest that even in terms of the currently existing classification schemas and ontologies, while the long-range possibility of developing efficient and effective ways to cross-search this wide range of resources certainly needs to be considered, short-range solutions will also and especially be valued in what may eventually come to be seen as a transition period.

Many other vexing problems exist as well. For example, the question of the clear need for a high-level ontology which can point to the general kind of information that is sought in a specific inquiry (distinguishing, for example, "apple" in horticulture, computing, folklore, cooking, etc.) has been partly conflated with the putative need for an upper ontology—highest level—in the architecture (see Section 4 below). One can observe that this last matter has hardly been resolved. Yet, these issues are crucially important because the Semantic Web is also uniquely concerned with "standards for knowledge representation" as Eugene Kim has commented [16].

There is also the question of ontology management and the tools used to map, merge, align, and integrate ontologies. It has been suggested by [10] that, since no management tool handles all these tasks, there is the need for a suite of tools assembled as a workbench to determine such terminological, taxonomic, and conceptual mismatches as occur, as well as to integrate tools and methods affording users enhanced interoperability in actual and conceivable projects. However, the full resolution of this problem will probably require considerably more research which, fortunately, is underway. Guy Mineau has shown [19] how problems of semantic interoperability could be automatically detected using CG-based software agents. In the same paper he demonstrates how information exchanged between communication systems could be controlled semantically by implementing semantic filters. This idea is picked up and developed by Dan Corbett [3] toward the goal of making knowledge bases semantically interoperable through the use of an ontology developed by the user involving a formal definition of ontology which is simplified by omitting the class/object boundary (employing CG theory which fully supports first-order logic). On the basis of CG theory Corbett demonstrates a method for automated comparisons of ontologies "represented as concept type hierarchies."

> Note that this hierarchy is not necessarily a taxonomy, in that a type may have multiple supertypes. Further note that there is no point on the hierarchy where we must make a distinction between a type and an instance. Every concept on the hierarchy is treated as a type. A type may have subtypes and supertypes, but there is no need to distinguish these from instances of the type. [Corbett then distinguishes this from the OO objective of objects inheriting properties of a class and gives as counter-example "treating a kitchen as you would any generic room."]

> The ontology (as a concept type hierarchy) acts as the framework, with conceptual graphs that conform to the hierarchy used to instantiate concepts in the domain. The ontology is populated by creating conceptual graphs which represent actions, ideas, situations of states in the domain. Recall, though, that a conceptual graph need not be a complete description, and will always be treated in the same manner as any other type [3].

But at this point it becomes important to consider that semantic interoperability also and especially involves sensitivity to the practices of users of information. Semantic interoperability necessarily involves the meanings that are embedded in information and the interpretation of these meanings, the interpreters usually being human users and not machines. This is the case both at the beginning and the end of the process—for example, the creation of ontologies and their utilization by individuals and organizations. So, in a word, semantic interoperability can be seen as essentially involving an *interpretive practice*. This brings us to our third and final architectural consideration.

Pragmatic architecture: This has been called by Miller a "landscape architecture" serving "to bound the realm of possibilities, to define what is 'in' and what is 'out, and to (ideally unambiguously) describe the relationships between users, resources, and technical systems"[17]. Its purpose in relation to the semantic sphere is to assist human collaboration by modeling and evolving optimal applications of the other two architectures as this involves ontology negotiation and interaction so that, for example, Harry Delugach [4] has set forth a vision of the possible development of practical active knowledge systems using semantics in operational systems within real contexts intending to solve complex practical problems. The concerns here naturally go far beyond compatible software, hardware, and ontologies and associated tools, all of which are necessary but not sufficient for such social uses as collaboration. Further, the pragmatic level ought to include what de Moor has termed "goal alignment" which, however, cannot be considered anything less than a formidable challenge. In this connection I argue in this year's conference paper that such extremely difficult socio-intellectual challenges might be approached through a reflection on Peirce's idea of *critical common sense* (CCS) and facilitated by certain forms of diagram observation developed for the purposes of goal alignment and other forms of consensus formation (for example, our trikonic vector analysis as applied to software engineering architecture) [32]. CCS would seem to represent the *sine qua non* of high-level consensus formation, and the need *to think critically together* in such matters has perhaps never been so pressing.

In another sense, the Pragmatic Web viewed as the goal and perhaps *telos* of information architecture relates all three as it "scopes the systems, data models, content, machine—machine (m2m) and machine-user interactions, and the environment within which interactions and transactions occur" [17]. Another way of saying this is that the Pragmatic Web *involves* the other two Webs, while it does not and cannot sublate them. The syntactic and semantic architectures are not *Aufgehoben* in the Hegelian sense, but the relation of the three can be analyzed involutionally in Peirce's sense: That is, the Pragmatic Web *involves* the Semantic Web which in turn *involves* the Syntactic Web. But at the end of the day, all remain in tact, and as equal partners.

3. Ontological Commitment in the Semantic Realm

Nicola Guarino asks us to "consider the distinction between "Ontology" (with the capital "O"). . . and "ontology" (with the lowercase "o")." He writes:

> In the philosophical sense, we may refer to an ontology as a particular system of categories accounting for a certain vision of the world. As such, this system does not depend on a particular language: Aristotle's ontology is always the same, independently of the language used to describe it. On the other hand, in its most prevalent use in AI, an ontology refers to an engineering artifact, constituted by a specific vocabulary used to describe a certain reality, plus a set of explicit assumptions regarding the intended meaning of the vocabulary words. This set of assumptions has usually the form of a first-order logical theory, where vocabulary words appear as unary or binary predicate names, respectively called concepts and relations. In the simplest case, an ontology describes a hierarchy of concepts related by subsumption relationships; in more sophisticated cases, suitable axioms are added in order to express other relationships between concepts and to constrain their intended interpretation [9].

From the standpoint of AI this seems accurate enough. But some philosophers might object to the somewhat limited characterization of philosophical Ontology. For example, Peirce sees metaphysics, or general ontology, as a *critical scientific discipline* falling exactly between logic as semeiotic (which has as its final branch methods of scientific inquiry involving the pragmatic maxim[3]) and the special sciences such as physics on the one hand and economics on the other. Peirce maintains that since there's no way really to avoid holding metaphysical views (so no way to avoid ontological commitments) that it is important that especially the scientist thoroughly criticize them.

> Find a scientific man who proposes to get along without any metaphysics -- not by any means every man who holds the ordinary reasonings of metaphysicians in scorn -- and you have found one whose doctrines are thoroughly vitiated by the crude and uncriticized metaphysics with which they are packed. We must philosophize, said the great naturalist Aristotle-- if only to avoid philosophizing. Every man of us has a metaphysics, and has to have one; and it will influence his life greatly. Far better, then, that that metaphysics should be criticized and not be allowed to run loose [CP 1.129].

So, someone states that he may dispense with metaphysics as he is a "practical" man. He apparently does not see that he is ontologically committed to his "practicalism" and a thousand and one presuppositions lie hidden behind that unstated commitment. Had he critically reflected on his *practical world view* in a scientific

[3] Consider what effects that might conceivably have practical bearings you conceive the objects of your conception to have. Then, your conception of those effects is the whole of your conception of the object.(CP 5.438)

spirit he might have eventually even come to a position not unlike authentic philosophical pragmatism (be it Peircean, Habermasian, Apelian, etc.)

So what are the kinds of questions that might be termed ‘metaphysical’ taking the term in this critical sense? Perhaps there’s no way to summarize Peirce’s thinking on this, but it seems significant that he prefaces his “short list” of pressing metaphysical questions with a comment suggesting that a scientific pragmatism might actually help settle them.

> There are certain questions commonly reckoned as metaphysical, and which certainly are so, if by metaphysics we mean ontology, which as soon as pragmatism is once sincerely accepted, cannot logically resist settlement. These are for example, What is reality? Are necessity and contingency real modes of being? Are the laws of nature real? Can they be assumed to be immutable or are they presumably results of evolution? Is there any real chance, or departure from real law? CP 5.496

Scientists who have studied Peirce’s metaphysics, for example, Nobel Prize winner Ilya Prigogine, have found him inspiring in the sense of encouraging them towards making promising abductions such as Prigogine’s of ‘dissipative structures’ in thermodynamic systems far from equilibrium [23]. Prigogine has written that “Peirce’s metaphysics . . . appears to be a pioneering step towards the understanding of the pluralism involved in physical laws” [24] and, indeed, it is even possible to see the universe to be of the nature of a complex argument [CP 5.119] involving not only deduction, but induction and abduction as well. Be that as it may, but returning now to Guarino’s analysis, we see that, while it may provide a good summary of ontology as “engineering artifact,” he goes on to make a rather radical terminological proposal.

> The two readings of ontology described above are indeed related [to] each other, but in order to solve the terminological impasse we need to choose one of them, inventing a new name for the other: we shall adopt the AI reading, using the word conceptualization to refer to the philosophical reading. So two ontologies can be different in the vocabulary used (using English or Italian words, for instance) while sharing the same conceptualization." [9]

This proposal to terminologically limit ‘ontology’ to its use in AI while renaming what has hitherto been designated ‘Ontology’ in philosophy as ‘conceptualization’ seems unlikely to be acceptable to philosophers nor to critical thinkers generally, and for the simple reason that the meaning of being would seem to far exceed the meaning of AI. In addition, ‘conceptualization’ seems too general and vernacular a word whereas a reconstituted and truly scientific metaphysics requires an exact and precise terminology. Finally, and perhaps most importantly, it is not necessary to see these two uses of the term as essentially problematic or in conflict and, indeed, the reverse ought be the case—they ought to be seen as mutually informing each other.

Peter Øhrstrøm, Jan Andersen, and Henrik Scharfe have also examined the idea of ontology as it appears in philosophy and computer science in “What Has Happened to Ontology” [20] emphasizing that the metaphysical notion of “ontological commitment” to true statements about reality ought to find its proper place in computer science. They suggest that such a move might actually prove mutually instructive and beneficial for both philosophical metaphysics and computer science.

> The introduction of conceptual ontology into computer- and communication science is an example of how knowledge from humanities can be imported and utilized in modern information technology. This emphasizes the importance of an ongoing dialogue between computer science and the humanities. Some of the new findings in computer science may in fact be generated by the use of research imported from the humanities [20].

The distinction between *Ontology* as a philosophical discipline and *ontology* as information practice is certainly important as it avoids conflating the problem of the communication of knowledge with the different, although no doubt related problem of existence as such, Reality in the Peircean sense of that which is what it is whatever anyone or any group at any given time may think it to be, or describe or classify it as being [CP 5.408]. It is, after all, what we think we know of some facet of some *Reality* that we wish to communicate. For even given Peirce's principle of 'fallibility' (that we cannot be absolutely certain that what we believe is true might not be proven false in further inquiry at a later date), it would appear that some positions are "truer" (i.e., less false) than others.

It would therefore seem highly desirable for computer scientists to consider the implications of philosophical Ontology in relation to the process of building ontologies. And especially as there may indeed be some hidden ontological commitments that those creating ontologies would be wise to bring to the surface and reflect upon, the study of philosophical Ontology *in a scientific spirit* may help make explicit what has been implicit regarding individual and communal ontological commitments.

A serious discussion of the relationship between philosophical Ontology and ontology creation has at best only begun. In this connection, Peirce's decidedly scientific metaphysics in consideration of an evolutionary Reality leading directly to the growth of consciousness and ultimately to the co-evolution of man and machine as Doug Engelbart has imagined it [7] may prove instructive. It seems to me that it is possible that the metaphysics of Peirce points precisely to the kind of ontological conditions (leading to commitments) ontologists ought to be reflecting upon.

The authors of "What Has Happened to Ontology" conclude that an ontology "ought to be based on a coherent and consistent theory which deals with Reality in a satisfactory manner. It is an obvious obligation on the developer of an ontology to discuss and defend his choice of theory and the ontological commitments to which it gives rise" {20]. Because there is one Reality--one Universe--however we may slice it up for the purposes of ontology building (while assuming an infinite number of possible theories), there is no need nor compelling reason to have one meaning for ontologists and a different one for other purposes. Perhaps distinguishing the two uses of the term by upper case 'O' and lower case 'o' will have to suffice for now and perhaps for some time to come[4].

---

[4] To help bridge the gap between "O" and "o" I would suggest that 'ontological commitment' might use the lowercase "o".

This is not to suggest, however, that there aren't some very significant related distinctions which we could and probably should be making, for example that between the tacit knowledge of human communities as opposed to the explicit knowledge which represents this tacit knowledge formally. It is perhaps of the essence of the pragmatic level that it attempts to reconnect tacit and explicit knowledge, striving to combine the (formal) reasoning power of machines with the (informal) interpretive power of humans and human communities[5].

4. Semantic Interoperability and Its Discontents

Semantic interoperability involves the meanings embedded in electronically exchanged information as well as the interpretation of these meanings. These interpretations are characteristically made by users, not by the electronic devices themselves, while the approach to the meaning-embedding has itself proved to be a thorny issue. The Semantic Web as conceived by Tim Berners-Lee and the W3C group [1] mixes some quite excellent features (such as URIs, Unicode, and XML) with some rather disappointing ones (for example, RDF has been shown [8] to be inadequate for logic, as an analysis of OWL—itself not fully supporting first-order logic [3]—in relation to RDF makes clear enough). It seems likely, however, that we may have to live with this situation involving RDF for some time to come unless both ***we and our machines*** begin to negotiate rather than to legislate such crucial decisions [30].

Yet this suboptimal condition is not in its nature a *fate* determining the further development of the Semantic Web (which after all is essentially only the *idea* of the same). John Sowa has recommended the adoption of the draft ISO standard for Common Logic [15] to serve as foundation for a more flexible, thus reinvigorated, Semantic Web. There are already mappings of RDF and OWL into Common Logic which also includes XCL, an XML notation, all of which adds to the feasibility of the proposal. Indeed, Sowa articulates what seems a reasonable and quite practical approach to semantic interoperability on the Web, what he calls the Unified Framework (UF).

> The point is the UF is primarily intended as a framework for communication among potentially (or actually) incompatible systems. The major inconsistencies arise at the level of axioms, which none of these systems would accept from one another. But they can usually accept lower-level facts without creating any conflict.
>
> Therefore, UF should be very rich in types, but very, very poor in axioms. Still, any serious inferencing (which may be logic-based, statistical, computational, or whatever) will require much more. Yet, every system that adds more does so in ways that are incompatible with some other system.
>
> Any axiom that causes a conflict with any major system shall be deleted from UF, but there may also be a large number of microtheories (as in Cyc) or modules. . .which could be very rich in axioms expressed in very rich versions of

[5] I am indebted to Aldo de Moor for this insight.

> logic. . . Any axiom that is deleted from UF will not go away, but it will be available in modules or microtheories that could be used as needed by various systems. In effect, the topmost levels of most ontologies are the most controversial. Therefore, UF should have a highly impoverished top level. . .with all the complexity moved to modules or microtheories.
> [31].

While this proposal seems not only sound but even rife with possibilities, Sowa has not yet been able to convince the ontology community at large of the efficacy of embracing this general ontology (or Unified Framework) that is yet capable of specialization. In one sense it is hardly surprising that, for example, the Standard Upper Ontology (SUO) group has so far not fully embraced this approach as it is precisely the upper levels of ontologies which have been controversial and Sowa is proposing that there be a *very* "underspecified" upper level. He argues that considerable progress could be made in the field if all perspectives sought to come to agreement on a precise definition of the "minimal assumptions" at the top. This would result in something "like a precisely defined and corrected WordNet extended with many additional vocabularies" which "would be 100% correct *all the time* for what it says, but it would not make any commitment on any controversial issue." One observes that WordNet has been combined with many ontologies, most notably Cyc and SUMO, so that this "defined and corrected" version of it is potentially very powerful as a place where all controlled vocabularies might meet [31].

Further, through a principle of *neutrality*, "all existing ontologies would be on an equal footing, and nobody who has an application that uses any of them would have to make any major adjustments" [31].We agree that such a Unified Framework, in providing only "placeholders" for the vocabulary, could facilitate an optimizing of semantic interoperability and thus help bring about a thorough-going strategy for defining and using ontologies. Summarizing this approach Sowa writes:

> Basic principle: it should be possible to import any subset of UF into any of the major systems such as Cyc, SUMO, Dolce, etc. without any fear of inconsistencies caused by conflicting axioms. It may be necessary to do some alignment of terminology, but after that has been done, the axioms should not conflict [31**].**

The "alignment of terminology" however, may not be achieved as easily as that phrase was written. Further, the knowledge representation community will have to become fully convinced of the reasonableness and efficacy of a UF if it is to be implemented. Meanwhile, one awaits stronger arguments against it, which so far seem to be centered on the "popularity" of RDF (and given RDF's deficiencies this seems hardly a good enough reason to endorse it in perpetuity).

5. Interoperability and SOA: Loose Coupling as a Way of Life

Einstein's remark that "Things should be made as simple as possible, but no simpler" has been suitably paraphrased for software systems architecture by Hao He as "Artificial dependencies should be reduced to the minimum, but real dependencies should not be altered" [11]. This paraphrase could serve as a succinct summary of the design principle of *loose coupling* wherein an attempt is made to minimize unnecessary constraints while recognizing that there are indeed some real dependencies that cannot be circumvented and that indeed need to be respected. The goal of achieving loose coupling would seem to constitute something of a *summum bonum* of computational interoperability in relation to the needs of users.

Interacting agents in loosely coupled relations already exist of course in all sorts of readily recognizable forms. Hao He gives as familiar example a CD which can be played on any number of machines (however varying in quality of sound, portability, etc., that is in the service provided) thus following a central principle of Service Oriented Architecture (SOA). The alternative Object Oriented (OO) architecture would in the present example require that each CD have its own player. There are organizations whose communication is hindered by intractable systems based on such inflexible architectures. In this connection, Simon Polovina has discussed the limitations of object-orientation and suggests how an "organizational semiotics" employing CGs could be employed to help overcome the "object-data divide" [22].

From our perspective loose coupling would appear to be the most promising and potentially powerful approach to achieving interoperability on the Web at this level, but SOA now fully realized in turns of flexibility and providing an increased potential for real organizational learning to occur. At the semantic-pragmatic juncture it may even have the potential for contributing to our creatively rethinking ways to reshape organizational design and behavior away from the tightly coupled and characteristically top-down approaches still dominant in many settings, towards flexibility and increased learning in the interest of healthy organizational development.

Yet what does this mean in terms of the architectures needed? First, SOA requires such design constraints as platform independence (= cross-platform structuring), extremely well-defined interfaces and "Coarse-grained, self-describing and self-contained messages" [21]. The essential design requirements in SOA are:

- simple/ubiquitous interfaces
- common semantics encoded at the interfaces
- interfaces available for all service providers/clients
- descriptive—mainly not prescriptive—messages delivered through the interfaces
- interfaces constrained by an extensible schema limiting the vocabulary and structure of the messages
- messages restricted in format, structure, and vocabulary (yet cognizant that increasing restriction decreases extensibility and vice versa)
- since extensibility is yet crucial, establishing a balance of restriction and extensibility [21]

If information and business managers increasingly come to appreciate its potential value, SOA could become a catalyst for changing more hard-wired and inflexible approaches. Perhaps it will even prove to be something of a necessity in a rapidly changing world where flexibility and growth would no longer appear to be merely options.

6. Conclusion: Catalyzing Architectural Change

Enabling the full interoperability of knowledge management tools in distributed information systems will no doubt require changing a number of assumptions about what this would entail. In truth, the kind of interoperability being argued for here would require something of a radical change in attitudes towards information. This re-visioning, however, holds the promise of transforming how organizations work and how people productively collaborate.

Organizational resource managers will need to see that increased openness and interoperability bring with them many advantages both to organizations internally by facilitating organizational communication, but also by potentially making them more visible in the market place—and for KR and KM especially "the market place of ideas"—and by contributing to the growth of a sophisticated, subtle, more fully informed digital community which, while perhaps finding it yet necessary to employ proprietary tools in transition, will less and less find such constraints on its potential for creativity and productivity acceptable. The fear that an evolved interoperable environment will not allow for the unique, the personal, for unique branding, and so forth, will need to be countered by the argument that by opening up the portals for exchanging information, the distinctiveness of products and services will be communicated to a larger audience.

It would appear that the UK has already begun to see and act upon the value of interoperability on a fairly large scale. Indeed, in some ways and at some levels, interoperability is already being widely achieved; for example, in the UK government at all levels, in the administration systems of universities, in the catalogues of museums major publishing houses, to name just the principal examples, interoperability has become more than just a goal—its very fruits are becoming evident. The man who once held the job title of "Interoperability Focus" has written that "undeniably valuable information is being made available to a wide range of users, often for the first time" [18].

Of course much more is needed in the UK and elsewhere in terms of harnessing methods, tools and procedures that can contribute to increasing interoperability. Certainly the great hope is that interoperability of the kind sketched here would lead to fundamental changes in the way that organizations operate, and this to their own advantage and to the advantage of society generally. So, the design and implementation of the common architectures which we hope to be developing need to be deeply informed not only by user requirements, but also by the worthiest goals of individuals, organizations and institutions. In this regard, Miller has written:

> "A truly interoperable organization is able to maximize the value and reuse potential of information under its control. It is also able to exchange this information effectively with other equally interoperable bodies, allowing new knowledge to be generated from the identification of relationships between previously unrelated sites of data" [17].

As Peirce succinctly put it, "Symbols grow," and in such ways that the knowledge representation community has the potential for contributing mightily to the development of robust and truly fully interoperable networks. These contributions have of course already begun. Yet bringing about the changes in principles and practices that might lead to the growth of healthy and fully empowered communities using technologies to achieve significant and beneficial cultural goals is unquestionably a difficult and unpredictable process for which there can be no recipe for or certainty of success. As Larry Hickman (paraphrasing another extraordinary Pragmatist, John Dewey) has commented, "technological revolution is not a matter of distinguishing technological artifacts from the ways in which we use them, because our technological artifacts *are* the ways we use them." [12] Perhaps this is why it has also been suggested that the problems we are facing are not new, but rather that they have to do with the really age-old problem of getting people to agree to those things which, were they to do so and truly commit to them, would benefit all concerned. Yet, because of the intense competitiveness often present in organizations, many enterprises are, as Hickman notes, using network services mainly as a means for integration and have hardly begun to consider advanced implementations.

There is not much that members of the knowledge representation community can do to bring about these broader kinds of cultural change *except* to try to come to better agreement among ourselves on such crucial matters as standards regarding some of the semantic and pragmatic issues discussed here, while continuing to develop the theories, methods, technologies and—of no less importance—the communities that help support complex and mutable processes, especially business processes. In this way, and as the title of one of Miller's papers has it, members of the knowledge representation community may yet hope to more fully realize what could be at the heart of our vocation, to become "architects of the information age."

Acknowledgments

I thank Aldo de Moor, Harry Delugach, and Simon Polovina for many things, the useful references they've sent me, their helpful comments on a draft of the paper, but especially their many *cotary* challenges. I also want to thank Ben Udell for his fine editing of the paper and for producing the graphic of the three Webs.

References

[CP] *Collected Papers of Charles Sanders Peirce*, 8 vols. Edited by Charles Hartshorne, Paul Weiss, and Arthur Burks (Harvard University Press, Cambridge, Massachusetts, 1931-1958).

1. T. Berners-Lee (2002). "Web Architecture from 50,000 Feet" (updated version of a presentation to the W3C chairs forum and the W3C Advisory Committee in 1998) http://www.w3.org/DesignIssues/Architecture.html
2. *Computer Technology Law Review.* http://subscript.bna.com/SAMPLES/ctl.nsf/04e3efa01c7130cf85256b57005a98b4/0ccc42ea9ef664bf8525707d006f437f?OpenDocument
3. D. Corbett (2004). "Interoperability of Ontologies Using Conceptual Graph Theory." In *Proc. of the 12th International Conference on Conceptual Structures (ICCS 2004), Huntsville, AL, USA*, Lecture Notes in Artificial Intelligence, No. 3127, Springer-Verlag, Berlin.
4. H. S. Delugach (2003). "Towards Building Active Knowledge Systems with Conceptual Graphs." In *Proc. of the 11th International Conference on Conceptual Structures (ICCS 2003), Dresden, German,* Lecture Notes in Artificial Intelligence, No. 2746, Springer-Verlag, Berlin.
5. A. de Moor (2005). "Patterns for the Pragmatic Web." In *Proc. of the 13th International Conference on Conceptual Structures (ICCS 2005), Kassel, Germany,* Lecture Notes in Artificial Intelligence, No.3596, Springer-Verlag, Berlin.
6. A. de Moor, M. Keeler, and G. Richmond (2002). "Towards a Pragmatic Web." In *Proc. of the 10th International Conference on Conceptual Structures (ICCS 2002), Borovets, Bulgaria*, Lecture Notes in Artificial Intelligence, No. 2393, Springer-Verlag, Berlin.
7 D. C. Engelbart (1962). "Augmenting Human Intellect: A Conceptual Framework." http://www.invisiblerevolution.net/engelbart/full_62_paper_augm_hum_int.html
8. P. F. Patel-Schneider (2005). "A Revised Architecture for Semantic Reasoning." *Third Workshop on Principles and Practices of Semantic Web Reasoning*, Dagstuhl, Germany. LNCS 3703, Springer Verlag.
9. N. Guarino (1998). "Formal Ontology and Information Systems." In *Proc. of FOIS '98, Trento, Italy*. IOS Press, Amsterdam (Amended version at http://www.loa-cnr.it/Papers/FOIS98.pdf).
10. A. Hameed, D. Sleeman, and A. Preece. "OntoManager: A Workbench Environment to facilitate Ontology Management and Interoperability." In *Proc. of the EON-2002, Workshop on Evaluation of Ontology-based Tools at the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW-2002),* Sigüenza, Spain.
11. H. He (2003). "What Is Service-Oriented Architecture." O'Reilly Media http://www.xml.com/pub/a/ws/2003/09/30/soa.html
12. L. Hickman. *Philosophical Tools for Technological Culture: Putting Pragmatism to Work*. Indiana University Press. Bloomington. 2001.
13. *HILT: High Level Thesaurus Project: Final Report To RSLP & JISC* (2001). http://hilt.cdlr.strath.ac.uk/Reports/Documents/HILTfinalreport.doc
14. "Interoperability" in *IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries*, Institute of Electrical and Electronics Engineers, 1990.
15. ISO/IEC FDIS 24707 – "Common Logic (CL) – A framework for a family of logic- based languages." http://cl.tamu.edu
16. E. E. Kim (2004). "A Manifesto for Collaborative Tools." *April 02, 2004* Dr. Dobb's Portal. http://www.ddj.com/showArticle.jhtml;jsessionid=2WVAZP0RVZRJ4QSNDBCSKHSCJUMEKJVN?articleID=184405649
17. P. Miller (2001). "Architects of the Information Age." *Ariadne Issue 29*. http://www.ariadne.ac.uk/issue29/miller/intro.html

18. P. Miller (2000). "Interoperability: What is it and Why should I want it?" *Ariadne Issue 24*. http://www.ariadne.ac.uk/issue24/interoperability/
19. G. W. Mineau (2002). "A First Step toward the Knowledge Web: Interoperability Issues among Conceptual Graph Based Software Agents Part I." In *Proc. of the 10th International Conference on Conceptual Structures (ICCS 2002), Borovets, Bulgaria,* Lecture Notes in Artificial Intelligence, No. 2393, Springer-Verlag, Berlin.
20. P. Øhrstrøm, J. Andersen, H. Scharfe (2005). "What Has Happened to Ontology." In *Proc. of the 13th International Conference on Conceptual Structures (ICCS 2005), Kassel, Germany*, Lecture Notes in Artificial Intelligence, No.3596, Springer-Verlag, Berlin.
21. D. Orchard (2004). "Achieving Loose Coupling." In *Dev2Dev* http://dev2dev.bea.com/pub/a/2004/02/orchard.html
22. S. Polovina, D. Strang (2003). "Facilitating Useful Object-Orientation for Virtual eBusiness Organisations Through Semiotic." http://www.polovina.me.uk/publications/os6-full-paper.pdf In *Proc. of The 6th International Workshop on Organisational Semiotics: Virtual, Distributed and Flexible Organisations, Reading, UK.*
23. I. Prigogine. *From Being to Becoming: Time and Complexity in the Physical Sciences*. W. H. Freeman. New York. 1980.
24. I. Prigogine and I. Stengers. *Order Out of Chaos.* Bantam. New York. 1984
25. Raymond, E. S. (2000) "The Cathedral and the Bazaar" http://www.catb.org/esr/writings/cathedral-bazaar/cathedral-bazaar/
26. G. Richmond (2005). "Outline of *trikonic* |>*k: Diagrammatic Trichotomic." In *Proc. of the 13th International Conference on Conceptual Structures (ICCS 2005), Kassel, Germany,* Lecture Notes in Artificial Intelligence, No.3596, Springer-Verlag, Berlin.
27. G. Richmond (with B. Udell). *trikonic*, slide show in ppt format of presentation at *PORT Workshop, ICCS 2004 Huntsville, Alabama.* http://members.door.net/arisbe/menu/library/aboutcsp/richmond/trikonicb.ppt
28. M. Schoop, A. de Moor, and J. L. G. Dietz (2006). "The Pragmatic Web: A Manifesto." In *Communications of the ACM*, 49, 5. New York.
29. M. P. Singh (2002). "The Pragmatic Web: Preliminary thoughts." In *Proc. of the NSF-OntoWeb Workshop on Database and Information Systems Research for Semantic Web and Enterprises.*
30. J. Sowa (2002). "Negotiation Instead of Legislation." Slides for a keynote presentation given at the Knowledge Technologies Conference, Seattle, Washington. http://www.jfsowa.com/talks/negotiat.htm
31. J. Sowa. (Nov. 24 & 25, 2005) in "Semantic Interoperability" thread, 2 posts to ONTAC-WG General Discussion http://colab.cim3.net/forum/ontac-forum/2005-11/msg00115.html
32. C. Spence-Hill, S. Polovina (2006). "Enhancing Software Engineering with Trikonic for the Knowledge Systems Architecture of CG Tools." Submission paper to *Conceptual Structures Tool Interoperability Workshop* of 14th International Conference on Conceptual Structures (ICCS 2002) Aalborg University Press, Aalborg, Denmark.
33. *Techworld.* http://www.techworld.com/opsys/features/index.cfm?featureid=1095
34. "Trikonic" article in *Wikipedia*. http://en.wikipedia.org/wiki/Trikonic