KNOWLEDGE-BASED SYSTEMS AND THE INTERNET: A FUTURE PERSPECTIVE

David M. Dutton, Robert W. Amor, David P. Bloomfield

ABSTRACT: This position paper details views of the UK Building Research Establishment with respect to future research concerning the Internet and Knowledge-Based Systems. It contains some view-points on what the future (both near and far) may hold for the construction industry in these areas. Discussion of these topics is presented with the intention of stimulating debate, and promoting consideration of them as possible research directions.

KEYWORDS: Knowledge-Based Systems, Internet, Industry Knowledge Base, Case-Based Reasoning, Machine Learning, Knowledge Discovery, Data Mining.

1. INTRODUCTION

Information Technology (IT) encompasses many techniques, methodologies and paradigms which have considerable potential for improving the management of information within the construction industry. Indeed, better and increased usage of IT has been identified as a prerequisite for the desired improvements in efficiency and quality within UK construction (Latham 1994).

Within IT there exist fields such as artificial intelligence (AI), which includes knowledge-based systems (KBSs), and the burgeoning Internet. As tasks become ever more information- and knowledge-laden, the requirement for *intelligent* decision support systems continues apace. The continued increase in the number and size of information resources indicates that better management of them is a necessity. Some commentators hold that the amount of data doubles every 20 months; others describe society as becoming increasingly data-rich but knowledge-poor. Progress toward remote working and virtual teams will further accentuate the need for fast, efficient, peer-to-peer communications to facilitate unfettered supply of information and knowledge as and when it is required. In these respects, the combination of AI and the Internet has much to offer the construction industry.

2. THE CURRENT SITUATION IN THE FIELD OF KBS

The Internet is unparalleled in its scope for providing information on a global scale. Such is the growth, that the search for, identification and retrieval of, relevant data/information has become problematic. This trend will continue but there are currently a number of initiatives within the construction industry aiming to structure and provide better access to such information (e.g. the UK Industry Knowledge Base, the Construction Information Gateway, see below). Moreover, there is potential in the longer term to encompass many disparate software systems, such that remote users, through to virtual corporations, might access whatever *expertise* they require *on-line*. For instance, consider the types of information required in *design* (Vanier and Turk 1994): building codes, vocabularies, national and international standards, national specifications, contract documents, as-built and working drawings, manufacturers drawings, specifications and instructions. An *ideal* design tool would support multi-media and would store, manipulate and communicate all of these.

World Wide Web browsers such as Netscape[1] and HOTJAVA[2] provide the means to download and execute mini-application programmes (applets), and thus it would appear that the ability to execute and interrogate remote information systems is approaching. Currently, the JAVA language (in which the HOTJAVA browser is written) allows the down-loading of applets from remote sites, to be executed *locally*. This represents a step forward in how users may wish to use the Internet, but fails to address the situation where users may wish to interrogate large, complex systems which precludes their down-loading before execution can take place. Far better in such cases that bandwidth is conserved and applications are executed in situ. Alternatively, the desired response times may dictate which option is chosen. Below, a number of largely disparate paradigms are introduced, followed by a discussion of how these might be integrated to facilitate more coherent and complete information and knowledge repositories.

KBSs offer ‘expert’ advice in narrow domains of application. The intention is to capture the knowledge of domain experts, and to represent this in an explicit form, whereupon it is used for reasoning on some complex task. KBS technology is sufficiently mature to form part of more widespread information systems, i.e. many organisations (not necessarily in construction) utilise KBSs as part of everyday business (e.g. Hayes-Roth and Jacobstein 1994, Allen 1994). Perhaps more frequently, KBSs constitute stand-alone decision support systems which are called upon by relatively isolated tasks, working groups, disciplines, etc. on an ad hoc basis. This state of affairs, together with the undoubted expense associated with developing and maintaining KBSs, implies there are many further subject areas which could benefit from KBSs. Possible reasons for the slow uptake are that users are unaware of the benefits; purchase/development is not economically justifiable; interface issues are insufficiently considered; or even that past KBSs have not met user needs, resulting in a general decline in user confidence. However, ‘successful’ KBSs can improve the efficiency and quality of information systems by an order of magnitude and their advantages are well documented (Hayes-Roth and Jacobstein 1994). Current KBS usage in construction includes design synthesis, cost estimation, fault diagnosis, failure prediction, and indeed, most of the construction process. However, it is hard to assess just how many of these are ‘successful’ commercial systems, applied beyond the prototype stage. To a certain extent, similar comments apply to other more conventional technologies (e.g. databases), especially where small to medium enterprises are concerned (SMEs). Therefore, the provision of such information services *across the Internet*, on a ‘pay-per-use’ basis constitutes one area of potential development, and would improve technology diffusion. The development of so-called ‘Internet computers’ corroborates the view that software will need to be remotely accessible.

In a marriage of the two technologies (amongst others, e.g. DBMSs), the feasibility of a UK Industry-wide Knowledge Base (IKB) is currently being studied by BRE and others. The UK Department of the Environment (DOE) has sponsored this study, and its aim is to improve the UK construction industry’s efficiency and quality, whilst reducing costs. The need for an IKB was identified in the Construct-IT report produced by British Telecom and Andersen Consulting (Construct-IT 1995). The feasibility study is being carried out by BRE and P-E Consulting and addresses not only the technical issues, but also the all-important commercial issues. The basic concept of the IKB is that of a ‘Construction Information Gateway’ acting as a single entry point to a set of distributed databases. A prototype system was developed using the Internet and can be accessed via URL http://cig.bre.co.uk. This prototype demonstrates that product catalogue data (for doors and windows), regulatory information, expert systems, etc. could all be easily accessed

[1] http://home.netscape.com/
[2] http://java.sun.com/

remotely. There are already many on-line systems offering similar functionality. What is novel about the CIG demonstrator is its use of STEP-compliant (Standard for Exchange of Product Data) technology to define the structure of the product data and thus to allow direct downloading into a 3D CAD system to update a building design with real components. This system has excited considerable interest but represents only one way in which the objectives of an IKB could be met. The DOE feasibility study is considering issues such as security, control of quality, liability, charging structure etc.. The latter two issues are clearly of crucial importance as such a system will not be sustainable unless all the parties - information originators, information providers and information users - can obtain the services they require at reasonable cost/price. This study should be complete by the time this paper is presented. An eventual IKB proposes to provide access to BRE (and others') information, such as Digests, reports, research results, etc., in electronic form, via Internet gateways. The aim is to preserve extensive and scarce expertise in a manner which may be more efficiently and effectively distributed to *those who need it*. Moreover, an organisation's Intranet (an 'internal' Internet) will also benefit from these trends, and may well be the first enabling step which organisations take in order to realise more ambitious IKB-type systems.

The representation of data/information/knowledge is of crucial importance for any KBS and as the level of semantics increases so do the difficulties in finding adequate descriptions. It is still a subject of debate how far it is necessary or practical to standardise definitions of building components and related data. The work of ISO 10303 (STEP), amongst others, has been ongoing for a considerable time, with more recent concentration of resources in the Buildings field. Progress has been made recently and useful and usable results are now beginning to emerge from, for example, EU projects COMBINE, CIMSTEEL, ATLAS and the STEP Building and Construction Core Model (BCCM). The recent formation of the Industry Alliance for Interoperability with the intention of agreeing 'Industry Foundation Classes' (and membership by many companies in a number of countries), is evidence of the perceived need to allow interoperability through shared definitions. These common definitions should allow higher quality KBSs to be built with greater in-built intelligence, and allow modular development leading to practical delivery of systems for wider markets.

Other less well known sub-fields of AI, viz. machine learning (ML), and its application (Knowledge Discovery in Databases, or KDD), and case-based reasoning (CBR) also have great potential. With CBR, past problems and their solutions are stored as individual case histories, and reasoning is based upon the retrieval and use of similar problem descriptions. In other words, a user describes the current situation (a problem) in terms of its salient characteristics, and the system matches this against a base of past problem descriptions. The best match is retrieved and its associated solution now describes one potential solution for the current problem. This paradigm is intuitively appealing in that humans often base solutions on past experiences. That is, reasoning is *analogical* and so is based upon context-sensitive information. In this manner, CBR systems in their simplest form can be thought of as external memories for humans, i.e. provide decision support. CBR offers a different slant to the more conventional KBSs, e.g. rule-based systems, and may overcome some problems associated with rule based systems (e.g. knowledge acquisition, and encoding this knowledge in rule form). Current examples include architectural design (Pearce et al. 1992), building defects diagnosis (Watson and Abdullah 1994), and timber selection for design (Dutton and Maun 1996). Nevertheless, fully-functional CBR systems typically cannot rely on cases only and require further domain expertise to enable past solutions to be adapted to current needs. This is often realised through more general rules, models etc., and may also be needed to provide solutions where sufficiently similar cases do not exist. Thus it would appear that CBR offers the potential to develop true 2nd generation expert systems which employ multiple representation and reasoning methods (David et al. 1993). The aim is to realise more robust, more general, more efficient and more effective KBSs which perform more complete reasoning cycles and cope with deficient data and other anomalies.

Machine learning is concerned with, in a practical sense, automating the acquisition and structuring of knowledge, and so is capable of alleviating a number of problems associated with KBS development (e.g. knowledge acquisition, refinement, and extension). ML most commonly involves the 'induction' of general rules from examples. An example typically describes one domain 'object' in terms of its salient characteristics. For example, if one wanted to (automatically) acquire knowledge on how to distinguish between different causes of dampness, then examples of past problems situations are required. (Note that experts are often much happier thinking in terms of example problems and solutions than general methods and rules.) It should be apparent that there are considerable similarities between CBR cases and ML examples. However, ML may require many examples, which are typically, significantly less complex than CBR cases.

More recently, together with other more conventional technologies (e.g. data warehousing), ML is being used for 'data mining', or knowledge discovery. This entails the semi- or fully-automated interrogation of an organisation's databases in order to uncover information and knowledge pertinent to business needs, such as trends, anomalies, errors, and so on. Such information can then simply be used by analysts, or form the knowledge base of a KBS. Data-warehousing can be considered an enabling technology in that it aims to 'fuse' numerous, separate data sources into a coherent whole. The use of 'middleware' permits the removal of added complexity inherent in querying such systems. In such situations, rows of tables (in a relational DBMS) may form the example sets required for automatic learning, and structured, even previously unknown knowledge can result.

In the following section, these areas and their (possible) inter-relationships are discussed with a view to mapping out what may occur in the construction industry in the future, how an IKB might be realised, and also to suggest what might form part of current and future research efforts. Moreover, the discussion constitutes the viewpoints of BRE's IT R&A Section, and are presented with the intention of promoting discussion. That is, feedback is welcomed, especially if concerning the possible future collaboration in the development of prototype systems.

3. A VISION OF FUTURE REQUIREMENTS

The development of multiple KBSs each attacking small sub-problems in the design field demands a system to allow the KBSs to communicate and pass control between each other to generate the required results. It is beneficial to distinguish between 'smaller' KBSs in the usual sense of the phrase and the larger all-encompassing IKB-type systems. That is, the typical expert systems which are tailored to narrow application areas are henceforth referred to as *micro-KBSs*, whilst the (potentially) multi-modal, multi-media IKB is referred to as a *macro-KBS*. This is intended to signify that the macro type may encompass disparate databases, multi-media, search engines, micro-KBSs, etc., all (possibly) accessed through hypertext. In addition, some or all of these constituents may be remotely located. For instance, *commercial* on-line information servers are already in operation in the UK, such as On-Demand, Building On-line, and Alpha-DIDO[3]. The addition of true knowledge interrogation will be a significant step.

---

[3] On-Demand Information: +44 171 291 1003

Purists in the AI field might not agree that the IKB represents a true knowledge-based system in that inherent knowledge is not necessarily explicitly represented (e.g. as in DBMSs and other such conventional software). Whatever one's definition of knowledge, it seems vital that the definition of KBSs will need to be reworked to encompass macro-KBSs. As (micro-) KBSs themselves will contain both knowledge and data (e.g. during an on-going consultation), it follows that larger collections of knowledge and data will also constitute a KBS. However, would a collection of disparate (micro-) KBSs and databases also constitute a (macro-) KBS? On the whole it would seem prudent to accept this only if there is some concrete inter-relationship between sub-systems, e.g. between a given micro-KBS and a database (etc.). In other words, one uses the other in providing its primary output(s), therefore precluding unrelated, unstructured groups of information systems. KBSs may currently utilise multi-media (etc.) to display relevant background information etc. and so some connections/relations are less concrete.

Within the KBS community, *blackboard systems* have been in existence for a number of years. Here, separate knowledge bases (KBs) cooperate to solve a problem. Typically each KB will address a sub-problem, and will communicate with each other via a central 'blackboard'. Blackboard systems allow more modular KBSs (and even hierarchical organisation), and increased flexibility (e.g. alternative control schemes as sub-systems are clearly separated). It is not hard to envisage such sub-systems as being physically distributed, i.e. not within one system, but accessible across the Internet. Therefore, one might also refer to the IKB as a blackboard system.

Such a scenario raises questions of coordination: perhaps one should view the World Wide Web (WWW) as an integrating technology, permitting access to specialised KBSs across the globe. For example, consider the following scenario. Multi-media software resides on a user's machine, and performs a coordinating role: keeping track of results to-date (working memory), and specifying what type and quantity of information is required to solve the current problem. HTML links (standard WWW interface description language) provide access to remote sites. The client tool accesses these sites; up-loads situation specific details (e.g. results of previous queries, held in working memory), and entering other details via forms interfaces or ftp (file transfer protocol). A user is able to view and execute remote information systems, down-load results onto the remote host and update working memory. This process would be re-iterated to fill-in all the required details. Unsatisfied goals, queries, etc. might be passed onto remote sites for further refinement and so on. Thus any one sub-system should have knowledge of the kinds of knowledge held elsewhere which might be of use to it. This could be implemented through information/knowledge classification, ontologies, and general *knowledge maps*, detailing types, quantities and alternative sources of knowledge, even mirror sites (for robustness). Any one site would on the whole only require knowledge of sites immediately applicable to itself. Thus requests for knowledge (etc.) which it is unable to answer may be passed onto other 'sites'. So for instance, in terms of rule-based micro-KBSs, the inference engine would be able to:-

- use rules to deduce new information;
- ask questions of users;
- search working memory for current conditions;
- and access remote sources, e.g. a DBMS, a KBS, or even a human expert.

Building On-line Ltd.: +44 171 837 5072
Alpha-DIDO: hr@cix.compulink.co.uk, +44 1483 203221

This is not as far-fetched as it might seem: current software (e.g. the Harvest system, http://harvest.cs.colorado.edu) enables the automatic searching and interrogation of say, construction related Internet sites, and builds an index of 'key' words. A user selects one or more words as the basis of a search and can access the related documents via the associated Universal Resource Locators (URLs). The database of URLs and the indices are incrementally updated over time. Knowledge maps would appear to be a relatively straight-forward extension of this idea. Indeed, this Section is currently formulating plans to use ML and KDD to classify documents (so aiding the indexing), and ultimately identify and scope remote knowledge sources.

The definition of a common language to allow systems to communicate meaningfully is still an open issue. Standard definitions are currently being defined for small domains (e.g. STEP). However, there is a realisation that it is not possible to develop a complete model for AEC, for all life-cycles. Parallel approaches contemplate dynamic schemes which expand and restructure as required by the attached applications. STEP, BCCM, Industry Foundation Classes and KBS ontologies may all offer means to represent and structure information and knowledge better, and to remove ambiguity, redundancy and so on. Other efforts in the classification of information in the construction industry will also play a part, as will other communication 'standards' such as CORBA and EDICON; document structuring standards such as SGML, even the likes of Agent Control Languages (Genesereth and Ketchpel 1994) for specifying communication between intelligent agents (see below) may play a part. On the whole, these initiatives can be seen as specifying the underlying technical detail required to fulfil these ideas. This discussion takes more of a conceptual view.

The classification of, and searching for, useful data are becoming longer and harder tasks. CBR might be thought of as an integrating technology, particularly if based upon a fully object-oriented representation language. If this is so, then cases of differing structure can be 'housed' under one system. Carried to the extreme, one can envisage a system which consists of several subsystems, some remote. An example might be a defects 'database' aimed at domestic housing (see diagram).

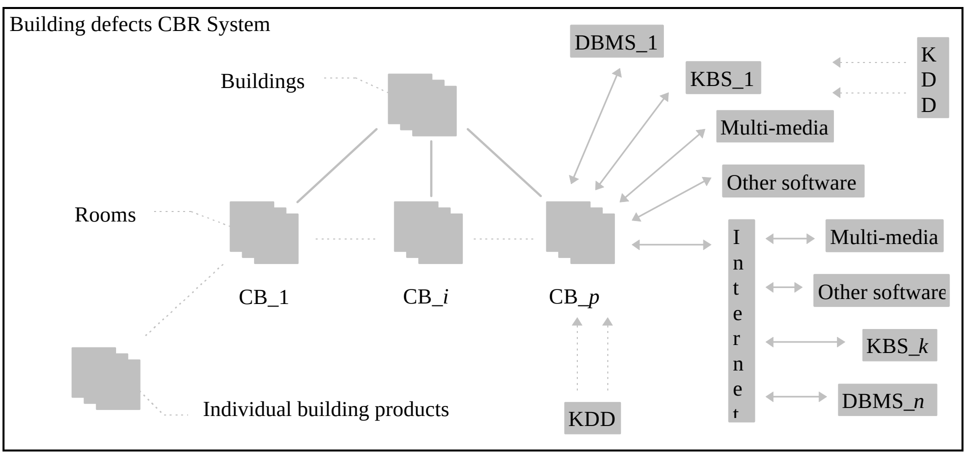

The many different types of fault would require substantially different case structures, and these might be arranged in a general-to-specific hierarchy. Each element of the hierarchy could consist of a 'cluster' or mini-case-base specific to the problem area (*CB_i*). Each cluster might be augmented with a KBS, for instance the extant BREDAMP which diagnoses causes of dampness. In this manner, each 'element' of an object hierarchy would be populated by its own case base of context-sensitive problem solutions, backed up by KBSs, DBMSs, multi-media display tools, and so on. General-to-specific hierarchies of ‘concepts’ would permit the abstraction of knowledge and data, so removing unnecessary detail, and specifying relationships between elements. Users could traverse the hierarchy as needs required, and as problem descriptions evolve and become more specific (i.e. as a consultation progressed). Alternatively, user interrogation may be via one of the more general top levels, whilst the KBS itself had access to a number of more specific sub-systems. Indeed, some CBR systems already utilise a similar structure, built up by successively generalising cases, in order to aid indexing and retrieval of the most useful cases in a given situation. Eventually this may allow users to retrieve all documents related to their current choice.

Searching the Internet currently can be an overwhelming experience. Words like ‘construction’ and ‘architecture’ have meanings in many contexts or domains such that their use becomes almost worthless. Narrowing the huge number of responses to a search can be a long and tedious task with no guarantees that the required information will still remain at the end, or where it was missed. CBR might be used in a number of other contexts as well, i.e. not simply as an alternative to rule-based systems. Current work includes the CAIRNS project (Watson 1996), which aims to use CBR to aid retrieval of information across the Internet. The objective is to remove some of the complexity and inherent variability of search engines and so forth. CBR can be extremely useful due to its capability for inexact matching (including natural language) and the scoring and ranking of retrieved solutions. Indeed, CBR as an *access* technology seems intuitively appealing. That is, CBR systems require cases, and cases may be ‘derived’ from conventional databases, e.g. one row or record per case. Such is the need for this kind of information acquisition that many CBR development tools provide means to access standard databases, e.g. through ODBC and SQL queries. Populating a case-base from a database does not, in certain cases, seem the most sensible approach. Where considerable redundancy may result, e.g. through duplication of data, (after all that is what DBMSs are designed to minimise), then retaining cases within the donor DBMS, but accessing them through CBR as required, might seem a more efficient route. As mentioned above, CBR’s flexible matching (also dealing with unknown, incorrect or imprecise data), and its ability to adapt solutions through domain knowledge, represents considerable potential. Thus an access-oriented CBR system would be populated by *virtual cases*. Extending this train of thought further, one realises that physically removed databases and even a number of disparate systems might be used to ‘donate’ virtual case bases. Again, this then implies the use of the Internet as a possible source for *distributed* CBR. Whether this occurs dynamically as a user interrogates a system, or complete but temporary case-base instantiation occurs before use, is a moot point. Such a system may be necessary if information is relatively volatile and must be updated regularly by its provider. This approach may be facilitated through data-warehousing, where numerous sources of data may be integrated into one DBMS.

In a similar manner, intelligent agent technology is currently being used to facilitate searching and retrieval from the Internet (and so forth). The general idea is that relatively simple but separate software ‘agents’ cooperate to fulfil allotted tasks. Thus communities of such agents are also distributed problem solving systems. This might be construed as an alternative arrangement for Internet access and even distributed KBSs, i.e. multi-agent systems (MASs) depend upon strong inter-communication capabilities and are thus adept at passing information between the various entities involved. Indeed, the modular composition of such systems could improve flexibility in both implementation and use, and distributed/virtual KBSs might in fact be implemented in MAS languages. In such a system, requests for information could be agent-based, with other agents residing at remote ‘donor’ sites. Such systems could also be capable of adapting to specific user’s needs, e.g. by building up a profile of macro-KBS usage over time. Perhaps this might improve information through-put by automatically reloading and caching the information and knowledge used the last time the system encountered a similar type of problem.

Other scenarios concern the above mentioned knowledge discovery (KDD). Due to the vast quantities of data pertaining to any one domain, the identification of relevant data may not always be enough. In many cases too much information will be identified and what is often required are summaries, trends and other inherent ‘knowledge’. Automating knowledge acquisition is of considerable benefit to industry, and offers a method to analyse and manage the plethora of information produced at an ever increasing rate. As the underlying ‘donor’ systems are again DBMSs and so forth, one again is faced with the prospect of ‘distributed’ information which can be accessed over the Internet. Thus automatically discovering information and knowledge from a number of sources simultaneously is possible. For instance, consider data-warehousing and enterprise information systems which aim to make better use of corporate data. Side-effects of this process could be to corroborate information and knowledge provided or discovered elsewhere, and also to integrate sources into a whole (e.g. IKB or warehousing). Again, CBR may well prove useful for mining and discovery due to the ability to add considerable background knowledge. Thus for a construction industry KDD tool, one could make use of contextual information which would constrain the discovery process and enhance the quality of the resulting information and knowledge. KDD and ML can be viewed as enabling technologies which improve the ability to sift large amounts of data, thereby instantiating data- and knowledge-bases. Removing this form of basic work from the ‘human-only’ realm will eventually be necessary to fully populate the IKB, and to maintain, debug, and refine it.

4. DISCUSSION

The above comments might seem to be mere pipe-dreams, but current projects can be construed as indicating such a ‘direction’ anyway, and would indeed provide a firm basis for further work. For instance, STEP and the BCCM aim to facilitate the inter-communication of data between disparate systems. It would seem prudent to extend this notion by including *knowledge*. Multi-agent systems are capable of just that as a by-product of their main purpose (distributed problem solving). Object-oriented systems aim to modularise system functionality such that objects may be re-used and ‘repackaged’ as and when required, into new systems, even communicated (CORBA). Extending this conceptual view to KBSs could enhance their (re-) usability, i.e. one does not need to reinvent the wheel; but also might encourage their integration into 2nd generation expert systems and other more conventional systems (in fact integrated ‘toolboxes’ are beginning to emerge, at least as research prototypes). Furthermore, intelligent CAD systems have emerged which integrate design knowledge with domain ‘objects’ (e.g. doors, windows, walls, bricks). The result is that objects ‘know’ how they relate to each other and can automatically reconfigure their properties (etc.) as a user alters or moves them about a design drawing.

KBSs with any country-specific, i.e. geographically sensitive knowledge may cause problems. For example, building regulations (and their interpretation) would differ across the world, as do climate, laws, work practices, and so on. On the other hand, multi-national companies may require up to date information on another country’s codes of practice etc.. Moreover, many faults (etc.) may well have international applicability. An analogous situation is observed with the CIB International Construction Database which contains abstracts of a diverse range of technical information, such as periodicals, books, reports, theses and conference proceedings. Thus international information can be, and is, ‘stored’ and used across the globe.

With respect to the structuring and coordination of macro-KBSs, domain ontologies could suffice, at least at a conceptual level. Ontologies aim to record domain terminology, concepts, relations, etc., and thus might perform the role of a data-dictionary and facilitate communication between KBS subparts. Indeed, a general to specific hierarchy of ontologies might be appropriate in such a system. Again, current projects within the industry aim to classify common terminology and information (e.g. UC/CI and others). As above, STEP and BCCM might provide the means for structuring information, with ontologies providing the necessary superstructure for knowledge.

Of course, realising such potential systems may never come to pass. One only has to consider the real problems associated with components of the preceding ‘framework’ to realise that distributed KBSs are probably a long way off. For example, distributed and federated databases are difficult to construct and maintain due to the proprietary systems which must be accommodated. Moreover, finalising the data structures in STEP standards is a lengthy and complex process. If such fundamentals are lacking, then one might be forgiven for being pessimistic. However, the current IKB project will hopefully prove this wrong. Other initiatives include IAI (Industry Alliance for Inter-operability) and IFC (Industry Foundation Classes) which have similar aims, albeit from a more practical viewpoint, and are intended to be usable in the short-term. Much thought will be required on the incompatibilities between competing software and hardware. However, open standards *are* a trend, and the JAVA language has made significant progress in this respect, i.e. it is not platform specific and so software written in it is highly portable. Afterall, the Internet shows how many different forms of hardware and software can cooperate.

Client-server architectures, where computing power is devolved to individual workstations instead of being provided as a central resource, are currently en vogue. However, many also recognise the continued usefulness of mainframe based systems, and both ‘architectures’ are suitable for different tasks. It may transpire that client-server technology is not suitable for multi-user interrogation of KBSs and DBMSs, and so larger scale systems may be required, e.g. data warehouses. On the other hand, suitably powerful workstations may ‘simply’ require distributed knowledge in order to function and thus only need to access ‘small’ packets of information on an as-needed basis. Current speculation on ‘Internet computers’, with less storage but stronger links to the Internet is also relevant. Such systems would facilitate the remote execution and/or interrogation of information systems.

Discussions concerning business over the Internet will also play a part. That is, if physically distributed KBSs are to become a reality, then issues such as performance, multi-user access and security will need to be addressed. Furthermore, information and knowledge maintenance and integrity could prove problematic. How can the completeness and correctness of provided information and knowledge be ensured? Access to sensitive material, and/or the potential for misuse and subsequent litigation may be reflected in pricing structures, which may preclude some users.

Finally, it is worth noting that general purpose KBSs are definitely not envisaged. The representation of common-sense type knowledge is notoriously difficult, if not impossible, and so any macro-KBS will more than likely still be constrained to one ‘domain’, e.g. an industrial sector. In such a scenario, the Internet would be of considerable benefit as many thousands of contributors would be able to provide and access information and knowledge as appropriate. Some interested parties could cooperate to provide some kind of structure (e.g. macro-KBS), whilst others could begin to populate the necessary sub-systems.

5. CONCLUSION

To paraphrase Vanier and Turk (1994), the Internet is growing bottom-up, i.e. it is demand led, which should encourage manufacturers, suppliers, consultants, builders and owners to follow suit. “Soon we will have .... this information at our fingertips” (Vanier and Turk 1994).

KBSs and the Internet are both interesting and exciting areas, and are vital for future progress in information dissemination. Current discussions in the UK on requirements for an IKB raise questions of structure, and the types of tools to be included. This paper discusses one conceptual viewpoint of a future ‘macro-KBS’ and how it may be realised. It is likely that of the paradigms mentioned here, e.g. KBS, CBR, KDD, STEP, BCCM, all will prove useful, and many will prove necessary. What is certain is that any successful realisation of an IKB will need to encompass these techniques.

Macro-KBSs can cover a large number of disparate information systems, which need to be coordinated in order to add value to the existing constituent systems. In addressing such aims, this vision consists of:-

- macro-KBSs which detail technical information, and consist of general principles, methods, prototypes, contextual examples, etc.;
- micro-KBSs which can learn and automatically adapt to circumstances;
- intelligent, fast, accurate and complete retrieval and filtering of information and knowledge according to (dynamic) user needs;
- effective intercommunication of information and knowledge between sub-parts, whilst retaining the inherent flexibility of including any system a user requires;
- effective decision support systems which facilitate knowledge-intensive tasks;
- virtual teams, remote working and peer-to-peer communications.

Only then will the full potential of AI and IT in construction be close to reality. This vision of a structured and coherent source of quality information and knowledge, which is easily accessible as and when required, should be encouraged. Sharing this on a global scale *will* facilitate a more efficient and effective industry and so should be an objective worth pursuing. BRE is already actively involved in these areas and is ready to partner others in securing their advancement.

6. REFERENCES

Allen, B. (1994). Case-Based Reasoning: Business Applications, Communications of the ACM, vol. 37, no. 3, March, pp. 41-2, ACM Press.

Construct-IT (1995). Construct-IT: Bridging the Gap, British Telecom, Andersen Consulting, DoE, HMSO.

David J., Krivine J., Simmons R. (1993). Second Generation Expert Systems, Springer Verlag.

Dutton, D. and Maun, K. (1996). CBR for the Selection of Appropriate Timber in Design, Proceedings of the 2nd UK CBR Workshop, Salford, April.

Geneserith, M. and Ketchpel, S. (1994). Software Agents, Communications of the ACM, vol. 37, no. 7, pp. 48-53, ACM Press.

Hayes-Roth, F. and Jacobstein, N. (1994). The State of Knowledge-Based Systems, Communications of the ACM, vol. 37, no. 3, March, pp. 27-39, ACM Press.

Latham M. (1994). Constructing the Team, HMSO.

Pearce, M., Goel, A., Kolodner, J., Zimring, C., Sentosa, L., and Billington, R. (1992). Case-Based Design Support: A Case Study in Architectural Design, IEEE Expert, October, pp. 14-20.

Vanier, D. and Turk, Z. (1994). Internet and the construction industry, Construct Canada : Computer Integrated Design and Construction, Canada, pp. 1-8.

Watson, I. & Abdullah, S. (1994). Developing CaseBased Reasoning Systems: A Case Study in Diagnosing Building Defects, Proc. IEE Colloquium on CaseBased Reasoning: Prospects for Applications, Digest No: 1994/057, IEE.

Watson I. (1996). Case-Based Information Retrieval in Construction project, personal communication.