Capturing Tacit Knowledge in P2P Networks

Paul Fergus, Anirach Mingkhwan, Madjid Merabti, Martin Hanneghan

Networked Appliances Laboratory
School of Computing and Mathematical Sciences
Liverpool John Moores University,
Byrom Street, Liverpool L3 3AF, UK.
Email: {cmppferg, cmsaming, M.Merabti, M.B. Hanneghan}@livjm.ac.uk

Abstract - **Encapsulating explicit knowledge within large scale organisations is well understood and several commercial off the shelf applications exist that enable the implementation of efficient knowledge base management systems. This provides the organisation with added value and enables knowledge components to be re-used and freely transferred throughout the organisational enterprise, however a large amount of implicit knowledge is directly unobtainable and often lost when employees leave the organisation. This is known as tacit knowledge and is generally deep routed within the employee's memory and communicated through face-to-face human interactions. The challenge must be to devise indirect mechanisms to extract and represent tacit knowledge using distributed peer-to-peer networks, further enhancing the intellectual capital within the organisation. This paper describes the socio-technical requirements paramount to capturing tacit knowledge using peer-to-peer enabled technologies. We describe the Conceptual Query (ConQue) algorithm and the Knowledge Source Extraction (KeSEn) algorithm and illustrate the functionality of our working model.**

I. INTRODUCTION

Imagine a situation whereby you resign from an organisation and part of the process is to give back all the know-how (tacit) knowledge you have accumulated during your time at the company. As well as returning your security cards, keys and office materials you have to download the know-how knowledge you have into an explicit digital format - this knowledge is then distributed within the organisational enterprise and used by employees to solve business requirements. To date this is unachievable – it is impossible to directly access the knowledge within a person's head and codify it into an explicit representation. However our research investigates the possibility of indirectly accessing this knowledge by amalgamating the social interactions between people and peer-to-peer collaborative systems.

Much of the information exchanged in our everyday lives is transferred via personal contacts and word of mouth. Technological advances in communication infrastructures and protocols have led to the exploitation of this behaviour as in peer-to-peer trading. This is an area that is well understood and has been implemented in applications such as Napster [9] and KaZaa [8]. The difficulty lies in capturing information that is private to the outside world – knowledge contained within our heads. Apart from ourselves, no one else is able to gain direct access to this information therefore we require indirect techniques to expose the tacit knowledge we are interested in. The challenge for the management of the emerging knowledge is particularly important when applied within a peer-to-peer environment. This challenge can be divided into two main tasks: a technology advance and knowledge capture. The former needs to address issues of modelling, scalability, operational deployment and management. The latter requires us to capture, manipulate, and exchange tacit knowledge.

The first task has in part been addressed and social interaction in distributed applications is becoming more widespread in the form of email exchange, messenger systems and the distribution of digital content within peer-to-peer networks. These technologies exist in isolation and in large are unexplored by organisational enterprises, however integrated collaborative applications such as Groove [11] and SWAP [3], aim to integrate these disparate technologies and are generating a great deal of interest. The second task is more problematic and requires us to understand the social interactions involved with knowledge dissemination and the technical requirements needed to capture and codify tacit knowledge.

We highlight a key theoretical solution for capturing tacit knowledge within peer-to-peer networks. We show how ad hoc systems can collate social human behaviour with technological advances and capture tacit knowledge within peer-to-peer networks. Our solution builds on our previous work [5, 6] and attempts to implement several new algorithms that enable peer level knowledge bases to evolve and self-organise by monitoring the collaborative interactions between the peer user and the peer network. The knowledge base itself is a collection of concepts, which can be classed as the peers ontological view point, which attempts to capture the tacit knowledge contained within the users head and conceptually map knowledge in the knowledge base with digital content and services the peer node provides.

ISBN: 1-9025-6009-4

II. BACKGROUND AND RELATED WORK

KnowledgeMail [10] addresses the challenge of capturing tacit knowledge by monitoring collaboration streams. This enables the organisation to capture what people know, their interests and the job function activities they are directly involved with. This is an automated process that relieves the employee from having to manually publish digital content, however the employee has control over what information is exposed. KnowledgeMail suffers from a number of drawbacks most notably that it is centralised and based on keyword search – resulting in irrelevant as well as relevant information being returned.

The concept of "emerging knowledge" is currently gathering pace as a research field. For example, SWAP (Semantic Web and Peer-to-Peer) [3] focuses on sharing information within and across organisations. Their solution aims to overcome the limitations with traditional centralised solutions that experience bottleneck problems and high maintenance costs by distributing information sources across a network of interconnected peers. An interesting SWAP feature that will be used in our research is the provision of a semantic layer enabling concept based processing in order to alleviate the inherent problems associated with keyword-based search found in current P2P applications.

Groove [11] supports synchronous and asynchronous collaborative tools such as whiteboards, threaded forums, text chatting, voice communication, shared calendars and cooperative web browsing. Groove integrates currently used technologies such as Instant Messaging, file-sharing and several other collaborative systems, creating an interesting solution that enables collaboration between people within the organisation. The application enables employees to form direct one-to-one communications as well as group-based virtual meeting rooms. This enables the flow of tacit knowledge between employees ensuring knowledge replication, which ensures that information is more readily captured and contained within the organisation. It is clear that Groove bridges the gap between social interaction and technological advances and enables the distribution of tacit knowledge, however Groove lacks the ability to capture and codify tacit information in machine understandable semantics. The challenge is to explicitly capture tacit knowledge and store these representations within knowledge base systems to enable intelligent integration of know-how knowledge accumulated over time.

An approach taken by Aberer, K. et al., [1] is to capture this knowledge through gossiping. Their approach aims to interconnect peers within a P2P network via user-defined schemas to share and incrementally evolve the search capabilities within the network. Their approach assumes a large amount of data exists and that it has been organised and annotated according to local schemas, which is not always the case in tacit knowledge. Using schemas in this way requires high maintenance and requires the replication of any changes throughout the rest of the system.

III. SOCIAL AND TECHNICAL REQUIREMENTS

In large the gap between what we do socially and what we do technically is significant and often performed in isolation. The challenge is to bridge the gap between these activities enabling employees within the organisation to collaborate via technological advances in distributed peer-to-peer systems. This to a certain extent is emerging as an embedded part of our everyday lives as we further rely on Instant Messenger systems to talk to friends and loved ones in countries far and wide as well as sharing and discovering distributed digital content. Furthermore the way we use current Internet technologies to perform a plethora of e-commerce tasks as well as using Internet sources to discover information and experiences relative to our everyday job functions, has made it difficult to imagine how we would cope if the Internet ceased to exist.

These technical advances have, in most cases, been ignored within large scale organisations who label technologies such as peer-to-peer and Instant Messaging systems as disruptive. However it is becoming apparent that these types of distributed systems have more to offer than just exchanging illegally produced digital content such as MP3's and the latest movies. Organisations are now questioning whether these technologies are capable of providing a vehicle for the dissemination of tacit knowledge within and across the organisation.

A. *Social Requirements*

Kim et al. [7], state "currently, we lack the technical mechanisms to fully bridge the gulf between the social and the technical worlds. Furthermore, many existing interactive systems are developed without the involvement, input, or advice of trained psychologists or human factors experts, who are best equipped to forecast the way that users will respond to new or different designs." We believe that large scale organisations overlook the social requirements for effective knowledge transfer and exhibit an element of distrust resulting in centralised control of how we conduct our day-to-day job functions. Technologies that reside outside the organisations control such as Instant Messenger systems and peer-to-peer networks are blocked and perceived as disruptive technologies. As Kim et al. highlight, management within organisations fail to understand the psychological issues involved with regard to the complex social interactions that take place within our day-to-day working lives.

Before collaborative technologies can be fully implemented we believe that a number of social requirements need to be met and that management need to:

- be totally committed; believe and understand the benefits associated with peer-to-peer collaborative systems.
- educate employees about how to share and discover tacit knowledge using peer-to-peer technologies.
- convince employees why it is important to contribute knowledge to the knowledge management system and how they will benefit from this in the long term.
- allow employees time to reflect on and codify the tacit knowledge they have.
- reduce the level of control enforced on the employee's technological interactions and allow communities to freely evolve over time.
- ensure that the organisation is treated as a living organism that can evolve and self-organise under its own dynamics – limiting centralised control as much as possible.

The organisation needs to enable an environment that guides and encourages employees to combine their everyday social interactions with distributed collaborative systems. Meeting these requirements will prove difficult for large organisations because this level of freedom is perceived as uncontrollable and difficult to monitor. However this level of freedom will allow us to indirectly access the plethora of tacit knowledge contained within the heads of employees within the organisational enterprise.

B. *Technical Requirements*

Implementing a technical-based social environment is well understood and evident in the form of Instant Messaging systems, white boards, file-sharing and forums. The challenge is to develop applications that can, in part, monitor these collaboration streams enabling tacit knowledge to be extracted and codified. Below we describe three primary issues we are currently addressing within our research in an attempt to provide a viable solution.

- Creating distributed collaborative environments and knowledge extraction algorithms: Tacit knowledge is inherently communicated via face-to-face interactions; therefore we need to integrate these social activities within a technological environment. Merging the social with the technical allows us to develop knowledge extraction algorithms that attempt to gain a conceptual understanding of these interactions in order to extract tacit knowledge and codify it in a knowledge management system. The challenge is to realise such an environment and develop algorithms that effectively extract and codify tacit knowledge. The primary motivation is to manage knowledge more effectively through sharing, reducing the requirement to re-discovering information that has already been discovered.

- Semantic interoperability within P2P networks: Explicit as well as tacit information distributed within the organisation will be inherently heterogeneous and in most cases human centric - the difficulty lies in how we address interoperability between different knowledge representations and extract formal meaning. For example, email messages have one structure, directory structures have another and communications within Instant Messaging systems are inherently represented in natural language and are subjective resulting in different ways of describing the same concept.

- Ontology construction and maintenance: Researchers within the Semantic Web community generally agree that ontologies enable semantic interoperability within distributed networks [2], however there is disagreement with regards to how ontologies should be constructed. Traditionally ontologies have been inherently centralised and maintained by a consortium of knowledge engineers. The main problem is getting everyone to agree on how conceptual information should be represented, and furthermore as the ontology grows it becomes increasingly difficult to maintain [4]. In our previous work [5] we have attempted to address these problems by completely automating the ontology construction process in a distributed peer-to-peer network and enabling the self-organisation of knowledge structures based on general consensus.

This paper concerns itself with capturing tacit knowledge in peer-to-peer networks, therefore it is paramount that these requirements are met in order to access knowledge sources and attempt to codify these interactions between employees to dynamically create self-organised knowledge management systems. The remainder of this paper provides a theoretical model of how we interact within this collaborative environment and explicitly codify tacit knowledge.

IV. THEORETICAL IMPLEMENTATION

In part we have developed several simple prototypes to evaluate and ground our concepts [5, 6] however there are several areas, which we have identified, that will be addressed in future work. This section describes how we intend to capture tacit knowledge – the algorithms we have identified to extract information from documents found in target locations – and the algorithms required to conceptually map queries with knowledge structures and digital content.

A. *Querying Tacit Knowledge*

In this section we describe the Conceptual Query (ConQue) algorithm, which attempts to determine whether queries conceptually match any of the knowledge structures contained in the peer's knowledge base and to discover the conceptual closeness of documents retrieved from the target location. Within our approach we make no assumptions regarding syntactic similarities between queries and concepts within the knowledge base, but rather attempt to determine conceptual relationships using inference based on semantics.

The Conceptual Query (ConQue) algorithm receives queries from the peer network and directly interacts with the peer's knowledge base. This algorithm uses the inferential capabilities of the knowledge base and tries to determine whether the peer contains structural information that conceptually matches the query received. The primary challenge is to extend the semantic queries we developed in [6] to provide a more generic means of creating and issuing queries as well as resolving syntactic and semantic differences. A theoretical solution for the ConQue algorithm is described below.

Documents contained in the peer's target directory are conceptually mapped to concepts contained in the peer's knowledge base. Fig. 1 describes the process of conceptually retrieving documents based on information contained in the query and information contained in the document. The initial step involves submitting the query to the ConQue algorithm, which extracts the keywords – in this instance 'Italian', 'Hotel' and 'Restaurant'. These keywords are used to query the knowledge base in order to determine any other keywords that are semantically related. The result of this is a table which contains the keywords extracted from the query and any keywords retrieved from the knowledge base. The ConQue algorithm uses this list of keywords to discover what documents conceptually relate to the query by looking up keyword configurations in the Document Mapping table.

Using the subsumption architecture of the concept, keyword configurations are generated that enable the ConQue algorithm to determine words that are syntactically different but semantically the same. For example Fig. 1 illustrates that we have a query, which contains three keywords, 'Italian', 'Hotel' and 'Accommodation'. The assumption is that the query is syntactically different from the information contained within the mapped document, but the two semantically match. For example the document contains information about 'Accommodation' and makes no reference to the word 'Hotel' and our query contains information about 'Hotel' and makes no reference to the word 'Accommodation'. It is clear that these two words are syntactically different, however semantically, in part, they mean the same thing. We can resolve this difference because within our concept we have defined 'Hotel' as a subclass of 'Accommodation', and we can prove this by using the inferential capabilities of the knowledge base to process the structure using subsumption and infer that 'Accommodation' and 'Hotel' mean the same thing.

The ConQue algorithm uses the inferential capabilities of the knowledge base to determine the semantic relationship between keywords in a document and the keywords in the query allowing the algorithm to generate keyword configurations that are semantically the same. The ConQue algorithm then uses these keyword configurations to search the Document Mapping table to find documents based on a conceptual understanding between the query and the document. This process is described in more detail in our working model in section 4.3 below.

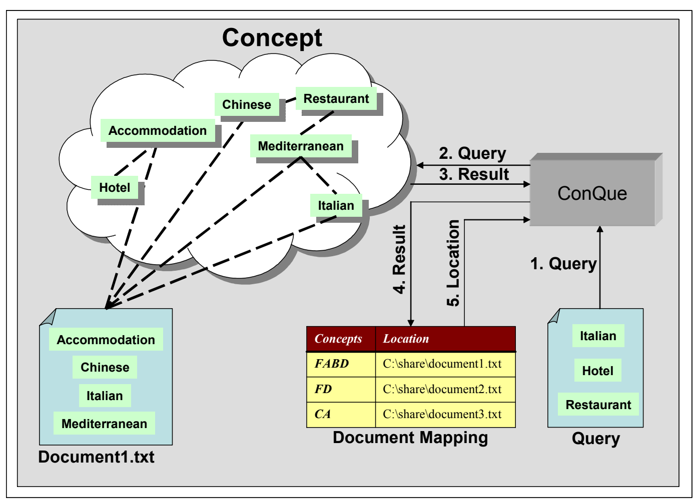

Fig. 1 Conceptually querying knowledge structures.

We have carried out simple experiments with this idea in our past work [6], however the development of queries were pre-determined and carried out within a controlled experimental environment. The queries we developed are able to dynamically discover Web Services based on semantics using the inferential capabilities provided by the knowledge base. The idea is to further enhance the query construction process using the ConQue algorithm, which will provide a more flexible and scalable means of query knowledge bases within the peer network.

B. *Conceptually Capturing Knowledge*

In this paper we describe several requirements needed to capture tacit knowledge in organisational enterprises and conceptually map documents on the peer node. We define documents as any information that has been generated by the user and which has automatically been codified within a file – this can be the result of codifying collaboration streams such as E-Conversations, white boards or forums or it may be the result of explicit text based documentation generated by the user.

The primary problems associated with knowledge extraction from unstructured documents relates to the way information is represented. In our previous work we processed information structures that are deterministic and structured using subsumption. This allows us to predetermine how to process knowledge structures and extract required information. This paper concerns itself with tacit knowledge, which is inherently communicated in a person-to-person manor. It is possible to process email messages and conversations that take place via Instant Messaging and e-group collaborations, however the information exchange is inherently represented in natural language. This format to a large extent is only readable by humans and proves difficult for machines to understand and process. This section explains how we propose to overcome this limitation and conceptually understand unstructured documents by describing the Knowledge Source Extraction (KeSEn) algorithm.

In this section we have identified unstructured documents, as semantically rich knowledge sources that allow us to capture tacit knowledge. The challenge is to develop an algorithm that automatically determines the concept of a particular document and link it to concepts contained within the peer's knowledge base. If the peer does not contain any concepts relating to the information extracted from the document then it must attempt to use the peer network to formulate concepts and extend its knowledge base to reflect newly acquired knowledge structures.

Fig. 2 describes the process of conceptually mapping knowledge sources on the peer node with concepts contained in the peer's knowledge base. The KeSEn algorithm initially begins by processing documents contained within a target location on the peer's workstation. The trigger for this process is based on the detection of any changes within the target directory – this may be event-driven or occur off-line as a house keeping routine.

The first step as depicted in Fig.2 is to retrieve the document and extract all the recognised words from the text. The keyword extraction process uses a knowledge base dictionary (KBDictionary) to extract words from the document that appear in the dictionary and disregard all others. The KBDictionary is a distributed representation of all concept names that exist within the peer network. When a new concept is added to a peer's knowledge base each word in the concept is added to the KBDictionary. Each peer maintains its own dictionary, which is periodically propagated to the peer network enabling peers to update their own KBDictionary to capture the concepts that exist within the peer network. The KBDictionary itself does not contain a particular concept but rather the keywords that appear in the concept. This enables peers to view the conceptual capabilities of the peer network as a whole.

Fig. 2 illustrates the format of the KBDictionary, which contains two columns representing the keywords and the location of keywords. The location of a keyword can reside either on the peer's local workstation or on another peer within the peer network. If the concept relating to the keyword resides locally then the peer contains the concept within its knowledge base, however if the concept relating to the keyword resides on another peer within the network, then the peer does not contain that concept but it knows that it exists.

Once all keywords have been extracted from the document they are sent to the DistrES [5] protocol along with location information about a concept. When a keyword is found which states a concept resides locally, the peer retrieves the concept from its knowledge base. If the keyword is found and the location is not located on the peer, the keyword is propagated to the peer network and all responses received are collected – this enables the DistrES [5] protocol to collect the conceptual interpretation of keywords from other peer nodes. Concepts retrieved from the peer's local knowledge base are combined with all concepts received from the peer network and evolved to determine the most optimal structure of a concept that captures the general consensus for structural categorisations.

Once an optimal structure is evolved the DistrES [5] protocol merges this structure with the peer's current structure in the knowledge base. When this process is complete a response is returned to the KeSEN algorithm notifying it that the concept has been created. This results in the KeSEn algorithm creating a conceptual mapping between the document and the knowledge base and adding a record to the Document Mapping table – this table contains all the keywords for relative concepts and the physical location of the document.

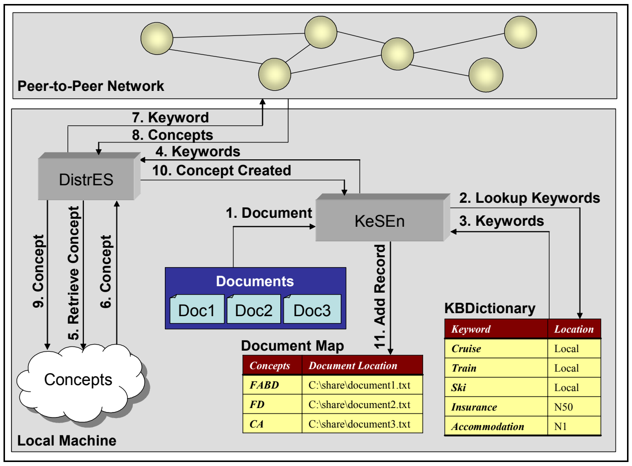

Concepts	Document Location
FABD	C:\share\document1.txt
FD	C:\share\document2.txt
CA	C:\share\document3.txt

Keyword	Location
Cruise	Local
Train	Local
Ski	Local
Insurance	N50
Accommodation	N1

Fig. 2 Capturing tacit knowledge and mapping knowledge sources.

As described above, this process periodically runs as a house keeping task to determine if any changes have occurred within knowledge sources on the peer node or as an event-driven model. It is envisaged that this algorithm will run using an event-driven model to dynamically ensure that the peer's knowledge base remains up to date.

C. *The Working Model*

Fig. 3 illustrates the theoretical results, which describes how queries received from the peer network are conceptually related to documents mapped in the target location on the peer node. This model shows that we can receive queries and extract keywords from those queries using the KBDictionary. Once the keywords have been extracted from the query they are used to retrieve keywords from the local knowledge base that semantically match, along with the fitness values of keywords. Fitness values are determined based on the taxonomical location of the keyword in the knowledge base. A direct map – this is where a keyword extracted from the query matches a word found in the knowledge base – results in a base fitness value being assigned to the keyword, which in this theoretical experiment equates to 10. We define this class as the base keyword, which enables us to evaluate semantic similarities between surrounding keywords. Any subclasses of this keyword are extracted and a fitness value is incrementally assigned dependent on the location of the sub classed keyword. Keywords that semantically distance themselves from the base keyword are assigned decremented fitness values dependent on the location of the keyword.

In Fig. 3 we can see that the keyword 'Mediterranean' has a fitness value of 10 because it appears in the query and appears in the concept – we can say that this is a direct mapping and therefore becomes the base keyword. As described above all the subclasses of the keyword 'Mediterranean' are retrieved and in this instance they appear one level below the base keyword so we assign a fitness value of 11 to the keywords 'Italian' and 'Spanish' and add them to the Semantic Keywords table. Once all the keywords from the query have been processed the Semantic Keywords table will contain all the words that are semantically related to the keywords found in the query. In this instance the keywords 'Cruise', 'Hotel' and 'Mediterranean' found in the query have yielded the following words 'Cruise', 'Caribbean', 'Hawaii', 'Hotel', 'Mediterranean', 'Italian' and 'Spanish' as a result of processing the concept, which indicates they are semantically related in this context.

Once the ConQue algorithm has processed all the query keywords a complete Semantic Keywords table is produced and is used to evaluate what documents mapped in the target location conceptually match the query. This is achieved by evaluating the Document Mapping table and calculating how many of the words found in the Semantic Keywords table appear in the concept column. For each word that appears the fitness value of the document is incremented based on the fitness value assigned to the keyword in the Semantic Keywords table. For example the Semantic Keywords 'Cruise', 'Hawaii', 'Hotel', and 'Spanish' appear in the solution document therefore the fitness value accumulated is 41. If a solution document cannot be found, the ConQue algorithm uses the Semantic Keywords table to determine alternative meanings for each word and attempts to find documents that most closely relate to the query conceptually. For example if a document cannot be found which contains any of the keywords that appear in the Semantic Keywords table, then the ConQue algorithm iterates through the semantic keywords and queries the knowledge base in order to determine if there are any other keywords that are similar. Fig. 3 illustrates that the first related document in the Query Result contains the keyword 'Restaurant' therefore we can determine that this is conceptually similar to 'Mediterranean' in this context by using the keyword 'Mediterranean' to extract similar keywords from the knowledge base using inference. As we infer up the tree structure we eventually reach the 'Restaurant' node. We can now retrieve this keyword and semantically match it with the keyword 'Restaurant' in the document and assign a lower scoring fitness value because it is conceptually similar to the keyword 'Mediterranean'. We perceive this as a necessary requirement because it provides a greater deal of flexibility – for example the solution document, although conceptually equivalent to the query in this instance, may not satisfy the query initiator or their may be a problem with the solution itself. In this situation the peer node can present alternate documents that are conceptually similar.

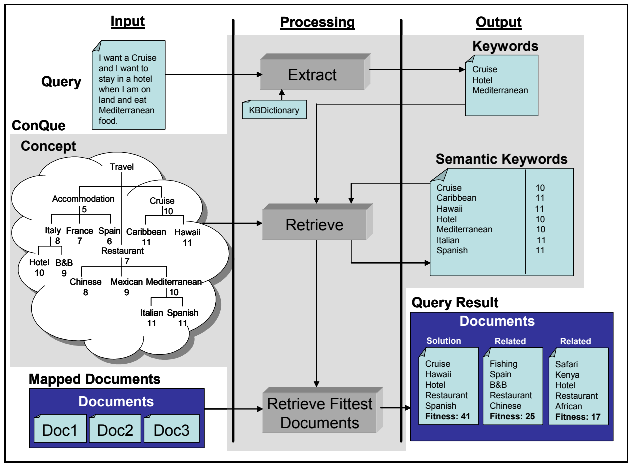

Fig. 3 Capturing tacit knowledge and mapping knowledge sources.

Our theoretical results prove that this is a viable solution and illustrates, in detail the requirements involved. In our future work we intend to take this solution past the design phase and implement it within the DiSUS [6] framework and directly link it into the DistrES [5] protocol.

V. CONCLUSION AND FUTURE WORK

In this paper we have defined the social and technical requirements required to enable the extraction and codifying of tacit knowledge. We explain that tacit knowledge is inherently deep routed within the minds of employees and is lost when they leave the organisation. We illustrate that we can overcome this loss of intellectual capital by developing collaborative applications that integrate the social with the technical, which allow us to capture and codify tacit knowledge. We describe that the gap between the social and technical is significant and often performed in isolation and that our solution describes how we need to bride this gap using distributed peer-to-peer collaborative systems.

We have developed a theoretical model in which we describe two algorithms respectively known as KeSEn and ConQue, which plug directly into the DiSUS framework we developed in a previous work [6]. We describe how the KeSEn algorithm attempts to conceptually understand the content within unstructured documents and create a mapping between this content and concepts contained within the peer's knowledge base. We further describe how the ConQue algorithm attempts to conceptually match incoming queries with mapped documents using keywords extracted from the query and keywords found in the Document Mapping table. We believe that these two algorithms will effectively codify tacit knowledge and provide a semantic layer, which will conceptually relate queries with documents found in the target location. This will enable the organisation to capitalise on the intellectual capital possessed by employees and ensure that this know-how information remains within the organisational enterprise even if the employee leaves. Furthermore we realise that the amount of information generated over time will be significant and lead to information overload. We understand that as well as collecting information we need to manage it and believe that our solution realises this by ensuring that information is processed conceptually and only matches mapped documents with queries if they are conceptually equivalent.

In our future work will complete the development and integration stages and produce the experimental results, which will enable us to determine the success of the two algorithms mentioned above. We have a sound conceptual understanding of the two protocols and have the detailed requirements needed for implementation.

References

[1] K. Aberer, Cudre-Mauroux, P., Hauswirth, M. The Chatty Web: Emergent Semantics Through Gossiping, Twelth International World Wide Web Conference. 2003. Budapest, Hungary: Springer,

[2] S. Decker, Melnik, S., Harmelen, F. V. Fensel, D., Klein, M., Broekstra, J., Erdmann, M., Horrocks, I., The Semantic Web: The Roles of XML and RDF. IEEE Internet Computing, 2000. 4(5): pp. 63 - 74.

[3] SWAP: Ontology-based Knowledge Management with Peer-to-Peer Technology. 2002, SWAP, Accessed: 05-03-2003, http://swap.semanticweb.org/public/publications/ehrig03swapb_wiamis-3.pdf.

[4] F. Fensel, Ontologies: Dynamic Networks of Formally Represented Meaning, Vrije University: Amsterdam, 2001.

[5] P. Fergus, Mingkhwan, A., Merabti, M., Hanneghan, M., Distributed Emergent Semantics in P2P, John Moores University: Liverpool, Technical Report, TR-CMS-2003-04-002, 2003.

[6] P. Fergus, Mingkhwan, A., Merabti, M., Hanneghan, M. DiSUS: Mobile Ad Hoc Network Unstructured Services, to appear in Personal Wireless Communications (PWC2003). 2003. Venice, Italy: Springer,

[7] K. Kim, Isenhour, P. L., Carroll, J. M., Rosson, M. B., Dunlap, D. R. TeacherBridge: Knowledge Management in Communities of Practice, Home Oriented Informatics and Telematics - The network home and the home of hte future (HOIT2003). 2003. Irvine, California, USA

[8] KaZaa. 2003, Sharman Networks, Accessed: 7-02-03, http://www.kazaa.com/us/index.php.

[9] R. Stern, Napster: A Walking Copyright Infringement. IEEE Micro, 2000. 20(6): p. 4 - 6.

[10] KnowledgeMail. 2003, Tacit Knowledge Systems, Inc., Accessed: 22-04-2003, http://www.tacit.com/products/knowledgemail/.

[11] Groove pushes sharing ideas on the Net. 2000, CNET News, Accessed: 20 April 2003, http://news.com.com/2100-1001-247492.html.