Context as a tool for organizing and sharing knowledge

Paul Warren[1], Ian Thurlow[2], John Davies[2], Igor Dolinsek[3], Carlos Ruiz[4], Jose-Manuel Gomez-Perez[4]

[1]Eurescom GmbH Wieblinger Weg 19/4, D-69121 Heidelberg, Germany
[2]BT Innovation & Development, Adastral Park, Martlesham Heath, Ipswich IP5 3RE, U.K.
[3]Hermes SoftLab d.o.o., Litijska 51, 1000 Ljubljana, Slovenia
[4]iSOCO, Pedro de Valdivia, 10, Madrid, Spain
warren@eurescom.eu, {ian.thurlow, john.nj.davies}@bt.com, igor.dolinsek@comtrade.com, {cruiz,jmgomez}@isoco.com

Abstract. This paper describes how tools and technologies have been developed within the ACTIVE project to enable user context to be employed to organize and share knowledge. The approach taken is a combination of top-down, where the user explicitly creates and manages context on the system, and bottom-up, where machine intelligence techniques support the user. The work of ACTIVE is being validated in three case studies, and the paper describes in particular how the approach to context is being validated in a case study in BT.

Keywords: context, information overload, knowledge sharing, machine learning, context detection, context discovery

1 Introduction

Information overload is a significant problem for all who perform knowledge work and the use of context is seen as an important technique for information delivery [1]. This paper describes work within the ACTIVE project (http://www.project-active.eu) which is using context both to combat information overload and to facilitate knowledge sharing. The paper describes the tools being developed and how they are being validated in a case study in BT.

ACTIVE is a European integrating project. Fundamental to the project are three interrelated research themes:

- Facilitating knowledge sharing through the synergy of a Web 2.0 approach and a more formal ontological approach.
- The provision of tools and technologies for describing, sharing, learning and optimizing informal, working processes.
- The use of context for combating information overload and for knowledge sharing.

This paper concentrates on the last of these themes. However, before discussing ACTIVE's work with context in detail, section 2 provides some background information about the project. Section 3 then talks specifically about the work in the project relating to context, whilst section 4 describes the validation of this work in the

BT case study. Section 5 discusses related work and section 6 summarizes some conclusions and discusses next steps.

2 The ACTIVE project

2.1 The ACTIVE research themes

Knowledge sharing – a synergy of formal and informal approaches

The first of these themes involves two complementary approaches. The Semantic MediaWiki (http://semantic-mediawiki.org/wiki/Semantic_MediaWiki) [2] is being extended on the basis of requirements elucidated from the project's case studies. The Semantic MediaWiki uses an informal approach to knowledge representation. For example, users are free to create their own relations rather than select from a pre-defined ontology. Where different relations are used to express the same semantics, these can be equated. The result is an approach to ontology and knowledge creation which offers less of a barrier to entry than a more formal approach. Whilst restricted in power it offers most or all of what is required in a knowledge management application. The second approach to knowledge sharing also emphasizes ease of use through the use of tagging to describe information objects. Here again, the simplicity of the approach arises from the fact that the user is not restricted to a pre-defined vocabulary. The user can choose whatever tags he or she likes. In addition, we are using machine intelligence techniques to suggest tags; the suggestions can be accepted or rejected. The suggestions are based on an analysis of file contents and also on tags used for similar files. Similarity is computed using a technique based on cosine similarity [3].

Informal processes – describing, sharing, learning and optimizing

Work on informal processes is motivated by the insight that, besides a relatively small number of well-defined business processes which every organization has, we all create our own informal processes to achieve our everyday working goals [4]. Because these processes are rarely written down they are not easily shared between colleagues. Hence they are continually being reinvented and not subject to the informal peer-review which would lead to their improvement. In ACTIVE we combine top-down and bottom-up approaches to the management of informal processes. A top-down approach means that users are given tools to describe these informal processes. In the bottom-up approach, technology is being developed to monitor user actions and learn the frequently-used processes. Once learned, they can be shared and compared with what is thought to be the ideal. Technology is also being developed to optimize these processes.

Our work on context also combines a top-down and a bottom-up approach. On the one hand, users are free to define their contexts and to specify their current context, that is the context in which they are currently working. On the other hand, machine intelligence techniques are available to discover the user's contexts from his or her activities and to detect the user's current context. The development of this approach, and its validation in the BT case study, is the theme of this paper and more information is provided in sections 3 and 4.

Work on each of these themes does not take place in isolation. ACTIVE is an integrating project, and the three themes all interact. This is achieved through the ACTIVE knowledge workspace (AKWS). The AKWS integrates the tools developed in ACTIVE, and ensures interoperability with an extensively enhanced Semantic MediaWiki and extended Microsoft Office tools.

2.2 ACTIVE architecture

ACTIVE uses a service-oriented approach. ACTIVE services perform user, information, context and process management as well as process mining and metadata recommendation. Each organization deploys the ACTIVE services on a set of their integrated servers or places them in a cloud. The complete functionality of ACTIVE services is accessible through a rich set of web service interfaces and through the ACTIVE event bus. End-users access the ACTIVE services through the ACTIVE web portal; through a dedicated ACTIVE taskbar; and through ACTIVE-extended Microsoft Office tools, and Windows File Explorer.

2.3 ACTIVE case studies

Three case studies are being used within ACTIVE to validate the research and development activity. In two of these, with Accenture and BT, we are using ACTIVE technology to support generic knowledge work. In the other, with Cadence Design Systems, the technology is being used to support the specific activity of electronic circuit design.

Accenture (http://www.accenture.com) is a global management consulting, technology services and outsourcing company, with more than 186,000 people serving clients in over 120 countries. Accenture consultants heavily rely on knowledge sharing and reuse to accomplish various tasks, for example the creation of customer proposals. This case study is making use of context, e.g. to present and prioritize search results.

The BT case study is similar to that of Accenture in being focused on people who work very much with customers and perform tasks such as customer proposal writing. However, whereas in the Accenture case study, ACTIVE technology is being used to enhance a pre-existing knowledge management system, in BT ACTIVE is deploying tools and user interfaces developed within the project. More detail about the BT case study is provided in section 4.

In Cadence, as with Accenture, ACTIVE technology is being used to enhance an existing suite of tools. As already noted, the goal is to support the electronic design process. The complexity of the design processes is such that it is often not obvious to an inexperienced designer what is the next appropriate step to take. This will often depend, e.g. on the type of integrated circuit being designed, and hence this kind of knowledge develops with experience. A goal of ACTIVE is to capture this tacit knowledge so that it can be shared with less experienced designers. The complexity of the process is also such that it is difficult to understand a designer's current position in a design process, i.e. how far through the current process he or she has got. This makes management of the process difficult. ACTIVE's process learning technology is being used to identify the designer's current position in the overall process and thereby make such design processes more manageable.

2.4 Validating the benefits

ACTIVE has a number of activities with the goal of validating the project benefits. Usability evaluation is being performed to ensure that the tools developed are right for their users. For the BT case study, this is discussed in more detail in section 4.

The business benefits of our approach are being evaluated within each of the case studies. This is being done by particular reference to the key goals of each case study. For example, in the Accenture and BT case studies a key goal is to develop higher quality customer proposals. To gauge the impact of our approach, customer proposals can be evaluated along a number of dimensions, e.g. the extent to which they address the customer's stated requirements, and internal consistency. Work is also being undertaken to understand the costs and benefits of creating and using lightweight ontologies, as in the Semantic MediaWiki, and folksonomies.

3 Context in ACTIVE

3.1 Using context

In ACTIVE we are concerned with using context to combat information overload and to facilitate knowledge sharing. At any given time, a user needs information relating to the context of his or her current work. A key, therefore, to combating information overload is to prioritize that information over information which relates to other contexts.

To take some concrete examples, we imagine that for a customer-facing person the set of contexts will include all the customers, or proposals, with which he is currently working, plus perhaps a few additional contexts such as 'admin' and 'personal'. For a lawyer, his or her contexts are likely to be the particular cases currently being worked on. However, we are not obliging our users to adopt any particular set of contexts. We are providing them with a toolset which they can adopt to their particular way of working.

What 'prioritized' means will depend on the particular application. For search it means that search results related to the user's current context appear higher up the search list than unrelated results. When opening a file within an application such as word-processing, it means that instead of seeing the most recently opened files, the user can opt to see the most recently opened files related to the current context. This is illustrated in figure 1, where the ‘open from context’ feature is being used, as shown highlighted.

We also believe that the use of shared contexts helps the sharing of information, as for any given context, a user can see what information his or her colleagues have available.

Note that ACTIVE enhances the user’s existing applications, rather than replacing them. In the BT case study described in section 4, the applications used are Microsoft Outlook, Word, Excel, PowerPoint, Internet Explorer and Windows File Explorer. Moreover, ACTIVE adds to existing functionality, it does not detract from it. So, in the word-processing example, if the user wants to see all his recently opened files, rather than just the ones in context, he is free to do so. Referring to figure 1, the user would simply click on the ‘open’ command further up the menu.

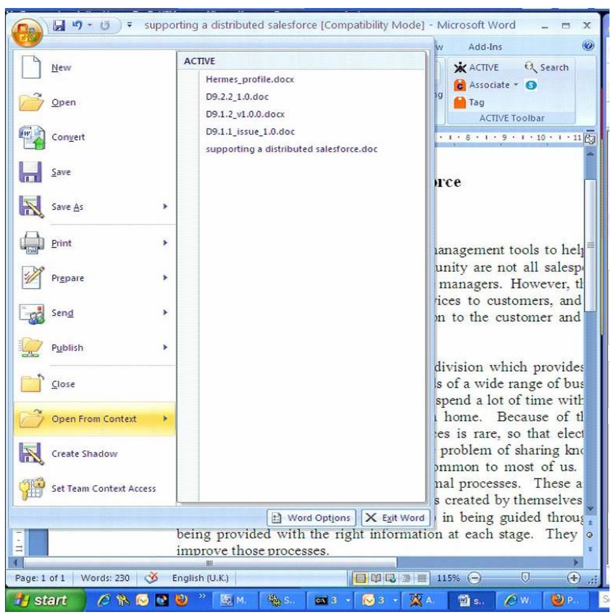

Fig. 1. A user in Word opting to select a file from those associated with his current context

In ACTIVE the relationship between a specific process and a specific context is for the user to determine. Processes and contexts can be independent. For example, in the BT case study a typical user might create a number of contexts corresponding to the customer with whom he or she regularly deals. A customer proposal process might then be created and be applicable to each of these customer contexts. In addition, a process can be assigned to a context, just like any other information object. So a user might have a context concerned with the financial outgoings incurred in his work, and a process for creating a travel expenses claim.

3.2 Top-down – creating and using context explicitly

Our starting point was to provide users with the ability to:

- create context
- assign themselves to contexts, and de-assign from contexts
- associate information objects with contexts
- set the current context

Note that the need for the user to assign him or herself to a context, rather than the context being automatically assigned when it is created, arises for two reasons. Firstly, contexts can be shared. So after one user has created a context, a second user may wish to assign him or herself to that context. Moreover, a user may cease to use a context for a period of time, but not want to delete it completely, because the context may be useful in the future. If the user de-assigns the context, he is free to re-assign it later.

Figure 2 shows the user creating a new context, called ‘Search’. Contexts can be hierarchical; in this case the user is specifying that ‘Search’ is a sub-context of ‘BT projects’. In the menu on the left there are a number of additional command options, including the option for assigning contexts to other users. The user can also display the currently assigned contexts, as show in Figure 3. As can be seen, a user can also be de-assigned from a context. In another option, a user can view all the contexts in the workspace.

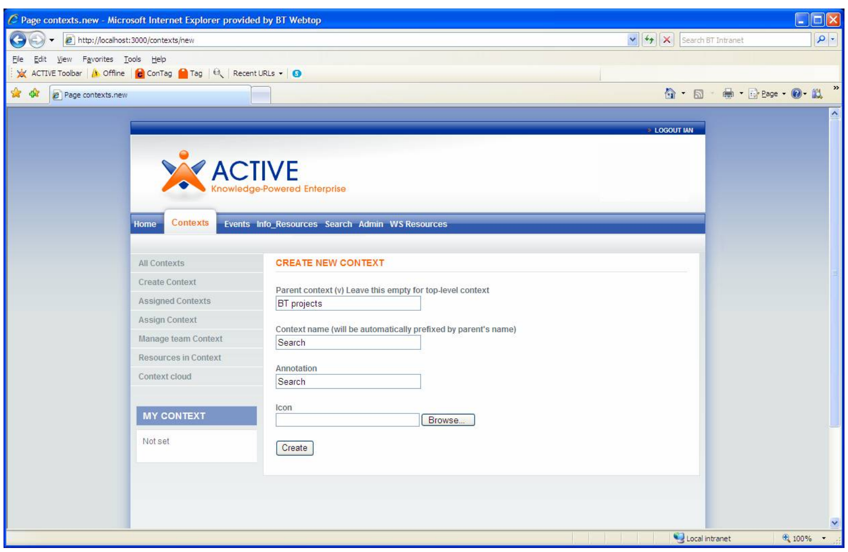

Fig. 2. Creating a context

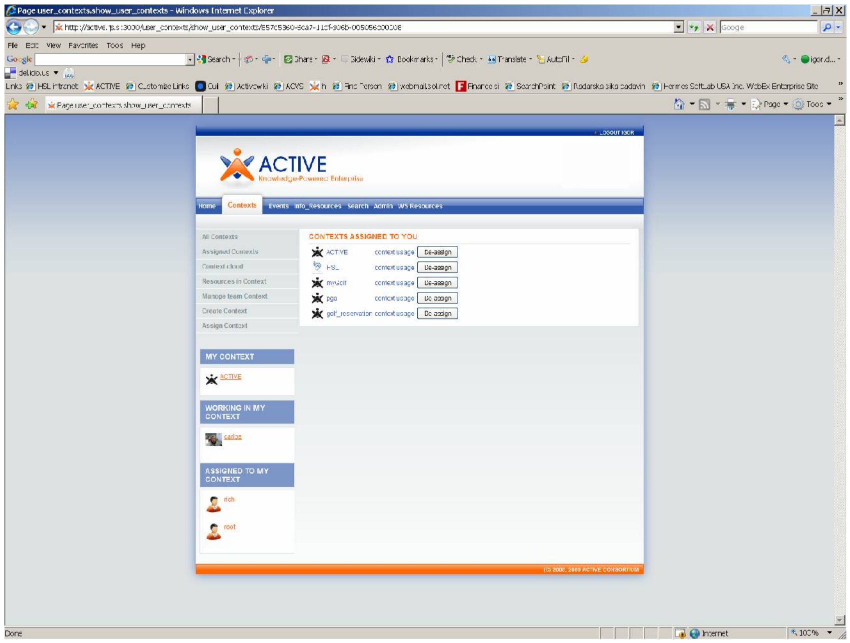

Fig. 3. Displaying assigned contexts

Associating an information object with a context is done via the ‘associate’ feature which is implemented in all the relevant applications. When the user clicks on the associate button, then the opened file, webpage or email is associated with the current context. Figure 4 illustrates this for Microsoft PowerPoint; in the figure the associate button is at the top-right, within the ‘ACTIVE toolbar’. Figure 5 illustrates a number of actions which can be applied to arbitrary files and directories in the extended Windows File Explorer. Files which are inserted into the workspace are marked with a small ACTIVE icon and the file manipulation menu is extended with ACTIVE actions, e.g. context association. Note that in figures 4 and 5, as well as an associate button, there is also a tag button to implement tagging. Clicking on ‘tag’ gives the user the opportunity to create his own tags or to accept system-suggested tags, as mentioned in section 2.1.

Users can also configure their profiles such that, whenever an information object is opened, then an association is automatically made between that object and the current context.

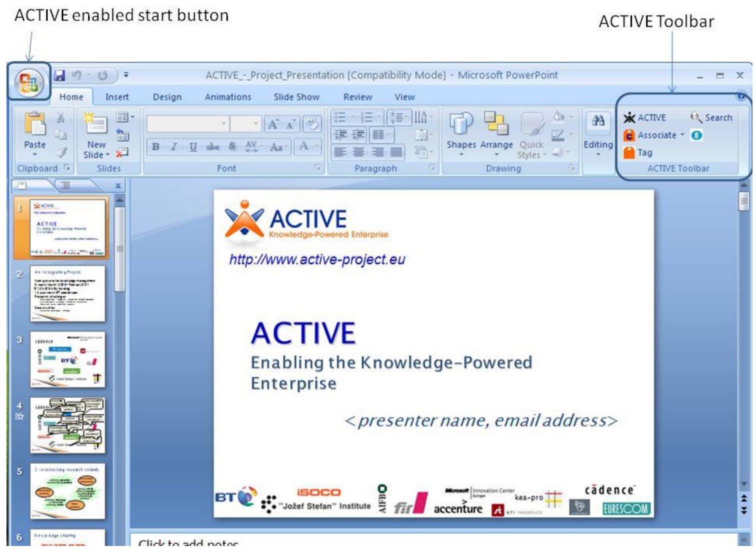

Fig. 4. Illustrating 'associate' in Microsoft PowerPoint

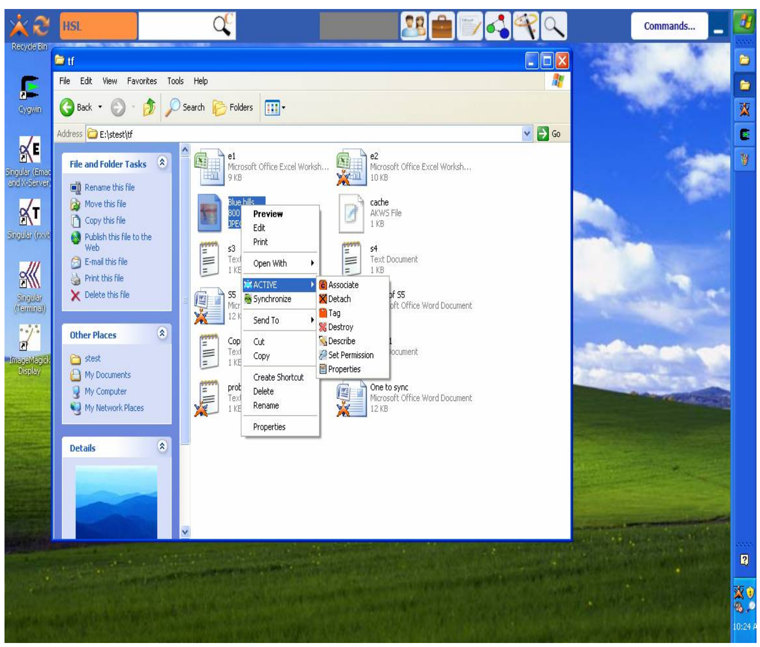

Fig. 5. Illustrating actions on a file in Windows File Explorer

The user switches between contexts by using the ACTIVE taskbar, shown at the top of Figure 6. On the left of the bar the current context (BT) is shown highlighted. Other contexts are shown in the drop-down menu; by clicking on one of these the user switches between contexts. The drop-down menu shows a flat list of contexts. As already noted in section 3.2, it is also possible for a user to create hierarchical contexts, so that one context is a sub-context of another. The figure illustrates a number of other taskbar functions. For example, the icon displaying two faces is used to show all team members assigned to the user's current context, whilst the briefcase icon is used to display all information objects associated with that context. In addition, the Context Visualizer, as described in section 3.4, can be invoked from the ACTIVE taskbar.

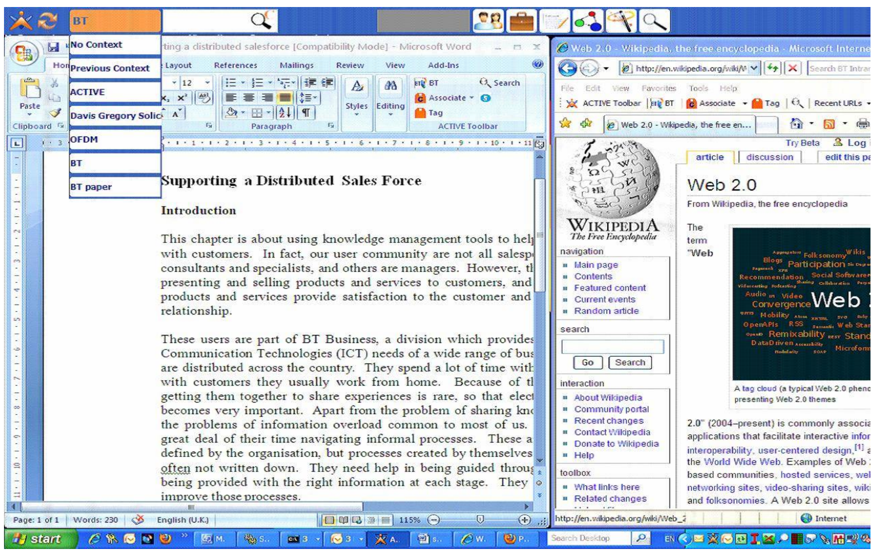

Fig. 6. The ACTIVE taskbar, showing the current context and alternative contexts

In summary, our top-down approach to context management assumes that users will create a set of contexts appropriate to their work; that they will associate information objects and informal processes with those contexts; that they will assign themselves not just to contexts they have created themselves, but to contexts created by others, thus forming shared contexts; and that they will explicitly set their current context in the taskbar. We believe that some users will be happy to follow this approach and will find it beneficial. The extent to which this hypothesis is true is a matter for validation, see section 4 below. We also take the view that many more users will benefit from a context-based approach to information management if there is a degree of automation, e.g. the automatic provision of suggested contexts. This is the subject of section 3.3.

3.3 Bottom-up – letting machine intelligence help

In ACTIVE we use machine intelligence:

- to detect the current context, i.e. based on the user's activities to determine what his current work focus is;
- to discover new contexts, i.e. to cluster the user's information objects in a way which reflects the contexts in which the user works.

As the user is working, the ACTIVE system analyses the documents, spreadsheets, emails, web-pages etc. which the user accesses and, based on this analysis, estimates what is the user's current working context. When a change in the user's context is detected with a sufficient degree of confidence, then one of two things happens. Depending on the user's preference settings, either the system switches context automatically or the user is asked to approve a switch in context.

Context discovery can be used to augment top-down user created contexts, or it can be used as the sole means of creating contexts. As with context detection, user behaviour and information objects are continually analysed and potential new contexts are identified. When the system has a certain degree of confidence that a new context has been discovered, then the user is alerted. In general, of course, this would happen much less frequently than the detection of a change of context. The user is provided with a list of the information objects associated with this new context, and is free to name the context and adopt it as one of his set of contexts, or to simply reject it.

The approach taken for context discovery is described in another paper in this workshop [5].

3.4 The context visualizer

A key goal of ACTIVE is to help knowledge workers manage and understand the context of their daily collaborative processes. The articulation of context using visualization tools supports this goal by helping understand the relationships within collaborative processes.

Context visualization tools need to explicitly articulate knowledge structures and their relationships for a particular given context. They also need to reduce the inherent complexity, facilitating the understanding of complex relationships within a context. They need to combine multiple modalities, e.g. graphs and text, and to provide filter options.

The ACTIVE Context Visualizer is a graphical visualization tool that explicitly articulates knowledge structures and their relationships for a particular given context. This allows the user to explore and browse the context and, for example, see how it is comprised of resources and how these resources are related and influenced by each other. Moreover, it offers other utility features such as zoom and diverse types of filtering. Figure 7 shows a screenshot of the tool.

The Context Visualizer contains two main parts. On the right, the *Diagram Pane* contains the graphical representation of the context. On the left, the *Information Pane* contains information about the context being viewed and the file, user or task currently selected.

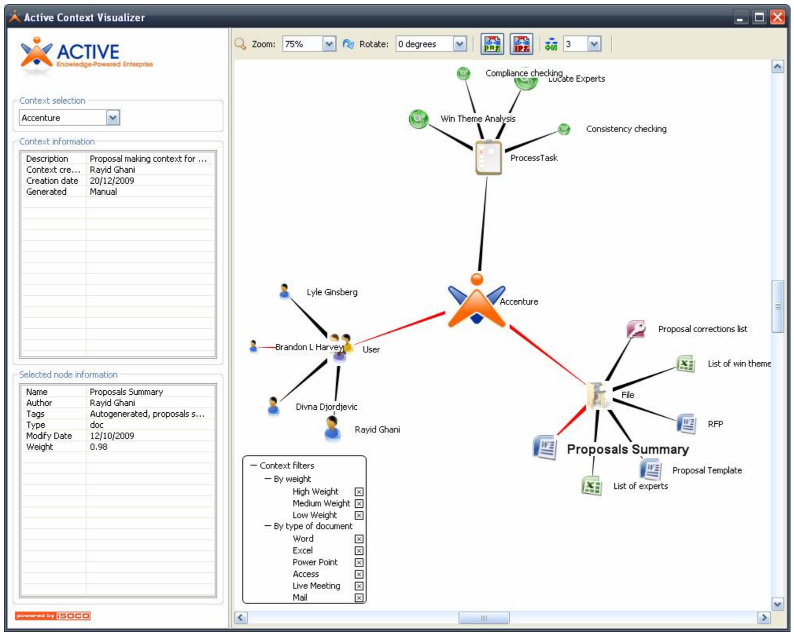

Figure 7 The context visualizer

The Diagram Pane shows the context graph and has three levels: the first level is the actual context, shown in the centre of figure 7; the second level comprises the categories of objects involved in the context, i.e. processes, users and files; the third and final level includes the particular instances of processes, users, and files. Each node is labelled with its name to allow easier exploration, and each node image is resized using a weight attribute depending on its importance within the context so the user can more easily locate the most relevant content. The currently selected file, user or task is shown in bold (‘Proposal Summary’ in the figure). The currently selected object is also shown interconnected, in red, with other resources which have a related use. Thus, in the figure, ‘Proposal Summary’ is interconnected with ‘Brandon L Harvey’ by red lines, implying the file has been created or accessed by Brandon L Harvey.

The Diagram Pane also contains the Filter Pane, which is used to filter files by type or nodes by their weight. In contexts with several hundred nodes, the filter helps the user to focus on the most important information by filtering by weight. As we expect the real-life contexts to have a large number of files, it will also be relevant to be able to filter these by their type, furthermore reducing the complexity of the visualized context and facilitating easier browsing and location of interesting information.

The Information Pane, on the left, is comprised of a combo box “Context selection” to select the context to browse and two information grids: “Context information” and “Selected node information”. The former displays information about the context being viewed. The latter has detailed information on the file, user or task currently selected in the context graph.

The components of the Context Visualizer architecture are:

- The *Context Mapper* receives the incoming context from other AKWS services and maps it to the internally used Context Representation Model.
- The *Context Representation Model* is the underlying model used to represent the context notion in ACTIVE [6]. It is based on PIMO [7] (Personal Information Model Ontology) which uses RDF and NRL (Nepomuk Representational Language) to express personal information resources
- The *Context Presentation Engine* defines a graph-based model for representing contexts based on JPowerGraph [8]. Such an engine enables flexible configuration.

4 Validation in the BT case study

The BT case study is validating a number of the tools being developed in ACTIVE; including the use of context to combat information overload and share knowledge. The trial is taking place within BT Business, a division of BT which works with small and medium-sized businesses across the U.K. The trialists themselves are ICT specialists, technical consultants, solutions consultants and sales managers. Whilst they represent a variety of roles, they are all customer-facing and our expectation was that their particular contexts would relate to the customers with whom they are dealing.

During the first year of the project an exercise was undertaken to understand their problems in handling information and how ACTIVE could help. This took a number of forms. Some potential trialists were interviewed. For a few we undertook job-shadowing, i.e. we observed the user working. To save travelling time and costs, and also to minimize disturbance to the potential trialists, this was done remotely. Microsoft LiveMeeting was used to view the subject's screen, and the voice channel of an instant messaging service was used to monitor conversations.

In parallel with this exercise we were developing the key ACTIVE technologies and understanding how those technologies could be applied. This enabled us to create a number of 'mock-ups' of ACTIVE functionality which were used at the beginning of the second year of ACTIVE to get more detailed feedback from potential users.

Feedback received during this phase suggested that the relevancy of context to a person's work is dependent on their role. For some participants the ability to switch between different contexts was particularly relevant. This was true, for example, of the ICT specialists who were regularly changing their focus of work. For others, in particular where work on a particular activity tended to be less fragmented, the context features of ACTIVE were not considered as relevant.

The exercise identified a number of requirements, some of which related to the theme of context:

- The effect of context switching on the behaviour of applications needs to be capable of being overwritten, i.e. one can always revert to the original form of information presentation. This had already been considered. For example, as explained in section 3.1, when using Word, the user can look at the most recently opened files in the particular context, or simply the most recently opened files in the usual way.

- An information item should be able to be associated with more than one context. This, also, is possible in the implemented system.
- Where users work in many contexts, there is a need for automatic context discovery and detection, as described in section 3.3.
- On context switch, applications which were previously open in the new context but had subsequently been closed, should be reopened. Conversely, applications which were previously closed and subsequently opened, should be closed. This feature can be implemented from the technical perspective, but it would require substantial development efforts. As noted in section 5, the ACTIVE approach is information-centric. This requirements corresponds to an alternative, activity-centric, approach.

As already explained, context is also seen as a way of enabling and encouraging knowledge sharing. When asked about knowledge sharing, the issue of trust was raised by participants. The majority of participants stated that they would be most likely to share information with members of their immediate team. Conversely, they would be most likely to make use of information which originated within their immediate team. In practice, an ACTIVE knowledge workspace could be limited to a small team, or extended to a department or even the whole organization, according to the needs of the business environment.

At the time of writing the ACTIVE system is being deployed to members of the BT Business trial community. During the closing months of the project an extensive validation exercise will take place, starting with around 20 users and ramping up to potentially 100. The overall goal will be to understand, and hence through feedback to improve, the users' perception of the system. All trialists are volunteers and are aware of the nature of the trial. Steps are taken to ensure confidentiality. In particular, the text of documents is not automatically stored on the ACTIVE server; for the purposes of context detection and discovery, only the 'bag of words' representation of the documents is required. Furthermore, the bag of words is tokenised such that each word is represented by a token whose meaning is held in a separate file. Moreover, users are free to specify that *https* web pages should be excluded from the analysis. Users will be questioned about their experiences with the system after a period of time using it. In addition, more immediate feedback on ACTIVE functionality is being obtained. Use of ACTIVE features will be monitored, and periodically the user will be requested to provide feedback relating to particular system interactions. For example, when the system detects a potential context switch, the user could from time to time be queried as to the appropriateness of the suggested switch.

5 Related work

There has been a great deal of research in recent years in the area of context. However, a glance at the proceedings of a conference devoted to the subject of context, e.g. [9] will illustrate the breadth of topics covered. Discussions of context include location-aware services, physical context and, closer to the interests of this paper, the use of context to resolve ambiguities. For a brief overview of these topics, and a list of some key recent papers, see [10]. The NeOn project (http://www.neon-project.org) is also looking at an aspect of context. The goal of NeOn is to improve "the capability to handle multiple networked ontologies that exist in a particular context, are created collaboratively, and might be highly dynamic and constantly evolving". NeOn recognizes that information is processed differently by different people and applications according to their context. Domain ontologies do not always fully specify knowledge because that knowledge is implicit in the context in which the ontologies are used. If the ontology is to be used in a different context, then this knowledge does need to be made explicit. This requires that the ontology be contextualized, i.e. that the relations be found with other ontologies which express this context [11]. More closely related to ACTIVE is work within the APOSDLE project (http://www.aposdle.tugraz.at/). APOSDLE provides support for e-learning within the "the user's immediate work environment and context". To achieve this, and to support user collaboration, it needs an understanding of user context. For more information on the APOSDLE use of context see, e.g., [12] and [13]. Another approach to context, also addressing the need for context switching, is described in [14]. Here, the user is able to switch between tasks and reinstate a group of windows associated with each task. This approach is activity-centric, in that it reinstates the user's total environment. Our approach is information-centric; more limited in that it simply reinstates the information prioritization associated with a context; more ambitious in the way it handles information objects and in the use of machine intelligence. Finally, parallels can be seen between our concept of context and the web application Google Wave (http://wave.google.com/about.html). In ACTIVE, context can contain a range of different kinds of information objects and supports the sharing of those objects. For Google Wave, each wave can be seen as analogous to a context, also comprising different kinds of information objects.

6 Conclusions and next steps

We have described a system which permits both the top-down creation and management of user contexts, and also the bottom-up discovery and detection of such contexts. Our system extends the standard Microsoft applications so that individuals can use their familiar applications whilst benefitting from the enhanced functionality. We believe that, whilst the top-down approach is on its own beneficial to some users, others will find real benefit in combining this with a bottom-up approach based on the use of machine intelligence. This has been confirmed in discussions with potential trialists.

As we have already indicated in section 4, during the final six months of the project, there will be an extensive validation exercise in the BT case study. This will include a detailed analysis of users' reactions to our contextual approach to knowledge management. Similar validation exercises will be undertaken in the other two case studies. In the Accenture case study, context is also being used to enhance knowledge worker productivity. The same underlying ACTIVE technology is being used as in the BT case, but implemented through a different toolset. Comparison of user reaction in the BT and Accenture case will further enhance our understanding of how context can best be used to help knowledge workers achieve their goals.

Acknowledgements

The work described in this paper has been funded as part of the IST-2007-215040 EU project, ACTIVE.

References

1. Driver, E., Moore, C., and Markham R., *Context is king in the new world of work*, Forrester Research (2008)
2. Krötzsch, M, Vrandecic, D., and Völkel, M., *Semantic MediaWiki*, Proceedings of ISWC2006, LNCS4273, edited Cruz et al., pp. 935-942, Springer-Verlag (2006)
3. Mitchell T M: 'Machine Learning', The McGraw-Hill Companies Inc (1997).
4. Hill, C., Yates, R., Jones, C., and Kogan, S., *Beyond predictable workflows: enhancing productivity in an artful business process*, IBM Systems Journal, vol. 45, no.4, pp. 663-682 (2006)
5. Stajner, T., Mladenic, D., Groblenik, M. Exploring contexts and actions in knowledge processes, Workshop on Context, Information and Ontologies (CIAO2010) held in conjunction with EKAW 2010 (2010)
6. Mladenić, D., Getoor, L., Ferlež, J., Grobelnik, M., D2.3.1 Context definition and representation within enterprise environments, ACTIVE project (IST-2007-215040) http://www.active-project.eu/fileadmin/public_documents/D231_context_definition_and_representation.pdf (2009)
7. PIMO: Personal Information Model. http://sourceforge.net/apps/trac/oscaf/wiki/PIMO
8. JPowerGraph. http://sourceforge.net/projects/jpowergraph /
9. Context 07, Sixth International and Interdisciplinary Conference on Modeling and Using Context Roskilde University http://context-07.ruc.dk/ (2007)
10. Mehra, P., Context-Aware Computing: Beyond Search and Location-Based Services, computing now, IEEE, August 2010, http://www.computer.org/portal/web/computingnow/archive/august2010 (2010)
11. Le Duc, C., Matching ontologies for context: The NeOn Alignment plug-in, NeOn deliverable D3.3.2, ftp://ftp.inrialpes.fr/pub/exmo/reports/neon-332.pdf (2006)
12. Musielak M., Hambach S., Christl C. APOSDLE Contextualized Cooperation. In: ACM SIGCHI: ACM Conference on Computer Supported Cooperative Work Electronic Proceedings: CSCW 08 [CD-ROM]. New York: ACM Press (2008)
13. Lokaiczyk, R.; Godehardt, E., Faatz, A., Goertz, M., Kienle A., Wessner, M., Ulbrich, A. Exploiting Context Information for Identification of Relevant Experts in Collaborative Workplace-Embedded E-Learning Environments, EC-TEL 2007, Crete, Greece (2007)
14. Robertson, G., Smith, G., Meyers, B., Baudisch, P., Czerwinski, M., Horvitz, E., Robbins, D., and Tan, D. Explorations in Task Management on the Desktop, in Beyond the Desktop Metaphor, ed. Kaptelinin V. and Czerwinski, M., p.p. 101-138, the MIT Press (2007)