

# ACTIVE – Enabling the Knowledge-Powered Enterprise

## Semantic Technology for Knowledge Worker Productivity

Paul Warren and Ian Thurlow  
Centre for Information & Security  
Systems Research  
Group CTO, BT  
Adastral Park, Martlesham Heath,  
Ipswich, Suffolk, UK, IP5 3RE  
[paul.w.warren@bt.com](mailto:paul.w.warren@bt.com),  
[ian.thurlow@bt.com](mailto:ian.thurlow@bt.com)

Rayid Ghani and Katharina Probst  
Accenture European Technology Park  
Les Bouillides - Les Genêts, 449, Route  
des Crêtes BP99  
F-06902 Sophia Antipolis, France  
[rayid.ghani@accenture.com](mailto:rayid.ghani@accenture.com),  
[katharina.probst@gmail.com](mailto:katharina.probst@gmail.com)

Eyck Jentzsch and Vadim Ermolayev  
Cadence Design Systems GmbH  
Mozartstr. 2  
85622 Feldkirchen, Germany  
[jentzsch@cadence.com](mailto:jentzsch@cadence.com),  
[vadim@ermolayev.com](mailto:vadim@ermolayev.com)

**Abstract:** ACTIVE, a three-year EU integrating project which began in March 2008, is using semantic technology to address three particular requirements of knowledge workers: the need to share information easily and effectively; the need to give priority to information which is relevant to the current knowledge worker task; and the need to reuse and share informal knowledge processes. ACTIVE technology is relevant to knowledge workers in all sectors. Specifically the results of ACTIVE are being validated in the consultancy; telecommunications and engineering sectors.

**Keywords -** *semantic technology; ontologies; knowledge sharing; knowledge processes; task context.*

### I. INTRODUCTION

#### A. The challenges of knowledge work

This paper describes a project, the goal of which is to use semantic technology to materially improve the productivity of knowledge workers in all sectors of the economy.

The project focuses on three particular challenges which we believe are significant for knowledge workers:

- Sharing knowledge easily and effectively, and without overwhelming with information<sup>1</sup>.
- Obtaining the information needed for their current task, and screening out information which, although it may be important, is not relevant to the task in hand.
- Reusing and sharing the informal knowledge processes which characterize so much of their work.

We believe that semantic technology can make a contribution to each of these. In particular, we are using ontologies to describe information, task context and knowledge processes.

---

<sup>1</sup> In this paper we will refer on occasions to *knowledge* and on occasions to *information*. There has been much debate over the difference in meaning between these two words. However, our goals are practical, and we do not find such debates helpful to these goals. In our paper, ‘knowledge’ is used in phrases where this usage is common practice, e.g. *knowledge worker*, *knowledge management*, *knowledge description*, and ‘information’ is used in less specific contexts. Generally we use whichever word seems the more natural English usage.

Supporting this we are using machine learning techniques to learn these ontological descriptions.

#### B. Project background

The work described here forms part of ACTIVE, an EU integrating project, which started in March 2008 and runs until February 2011. The project has twelve partners: BT, the project coordinator which also supplies a case study; the Universities of Karlsruhe and Innsbruck, the Jozef Stefan Institute, Microsoft Innovation Centre, and iSOCO, all of whom chiefly provide the research input; Accenture and Cadence Design Systems, who both provide case studies; FIR and kea-pro who lead the organizational and user validation activities respectively; Hermes SoftLab who undertake systems integration; and Eurescom who provide the project management. The project budget is €1.9 Million, of which €8.2 Million comes from the European Commission as part of the FP7 Programme. More detail about the project and project partners can be found on the project web-site (<http://www.active-project.eu>).

### II. KNOWLEDGE WORK AND KNOWLEDGE WORKERS

The term *knowledge worker* is generally attributed to Peter Drucker [1]. Since Drucker’s original use of the term, much has been written about knowledge workers, who they are, what they do, and how their number is growing e.g. [2]. For our purposes we accept as a given that work for an increasing proportion of people in the world economy involves the use of mental effort rather than, or for some as well as, physical labour. To significantly increase economic productivity it is necessary, therefore, to increase the productivity of this knowledge-based and knowledge-driven work.

Whilst knowledge workers appear often to be working alone, e.g. when writing a report, we know from our own experience that they are critically dependent on colleagues. The documents they create are often multi-authored and in any case they frequently need the knowledge of others to assist

them in achieving their own goals. Therefore improved *knowledge sharing* is a key objective of our work.

We also know that these knowledge workers are dependent on a variety of software tools on computers or other electronic devices. The shipment figures for Microsoft software stand as evidence of this. In fact, the tools knowledge workers use can generally be classified into generic which are used by all knowledge workers, e.g. word-processing, email; and domain-specific, e.g. engineering design tools used by specialist engineers. ACTIVE is considering both sorts of tools.

A significant role of these tools is to deliver information to their users. However, through these tools knowledge workers are frequently overwhelmed by information but still lack information important to their current task. We believe that the concept of *task context* is central to improving this situation. Knowledge workers need easy access to all available information which is relevant to their current task. On the other hand, they need to be screened from information which is not currently relevant. The availability of a mass of information irrelevant to the current task makes it harder to find what is relevant. Constant interruptions which are not relevant to the current task lead to frequent task-switching; such task-switching can undermine productivity. The management of information, to ensure that what is currently relevant is given prominence, is therefore vital to increasing productivity. Of course, we do not mean that other information, e.g. incoming emails, is simply discarded; it may be relevant later. Above all, the user himself or herself must always be in control.

All knowledge workers engage in processes. Some of these processes are formal business processes. Typically they are defined by the organisation; often it is not permissible to deviate from them. Languages for such processes now exist, e.g. WS-BPEL<sup>2</sup> for process execution and BPMN (Business Process Modelling Notation)<sup>3</sup> for process description. However, we observe that knowledge workers create their own informal processes, which we term *knowledge processes*. Examples are the processes which an organisational employee uses to obtain information on a customer from a variety of sources; or more generally any processes to obtain information, e.g. where an experienced individual uses his or her expertise to determine the most appropriate information sources; or the processes of arranging the flow of technological operations in engineering design based on the specific character of the design artifact and the skill profiles of involved designers. The processes are frequently not written down, or if they are then only very informally. This hinders their reuse even by their creators, and certainly hinders sharing between colleagues. We need, therefore to provide assistance to knowledge workers in creating, reusing, sharing and also improving on these knowledge processes.

---

<sup>2</sup> WS-BPEL (Web Services Business Process Execution Language) is an OASIS standard, see: <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-v2.0.pdf>

<sup>3</sup> BPMN is an OMG standard, see: <http://www.bpmn.org/>

### III. SEMANTIC TECHNOLOGY FOR KNOWLEDGE TOOLS

Sections I and II introduced and described three significant challenges for knowledge workers. This section describes how semantic technology can be used to create tools to help respond to those challenges.

#### A. Knowledge sharing

Effective knowledge sharing has been widely studied, both from a technological and an organisational viewpoint. The challenge is to ensure that knowledge within an organisation is available to all who need it. For some global organisations this means that knowledge created on one continent needs to be available to colleagues on another continent. Sometimes creators and users of the knowledge will not know each other, or even be aware of each other.

The innovation which ACTIVE brings to this problem is to combine two parallel approaches to knowledge representation: formal and informal. On the one hand, over the last decade or so, considerable progress has been made in developing formal techniques for describing and reasoning about knowledge. These developments are referred to broadly as 'semantic technologies' and are characterised by the use of ontologies to represent knowledge<sup>4</sup>. At the same time, and starting in the hobbyist and consumer world, informal tags have been used to describe documents and media objects on the web. Such tags are said to constitute folksonomies. Flickr (<http://www.flickr.com>), the web-site for sharing photos, is an example of this. Recently organisations have begun to use such techniques for their knowledge sharing. Ontologies and folksonomies offer different advantages and disadvantages. The former enable rich knowledge descriptions and machine reasoning over those descriptions; however the development and maintenance costs, e.g. the cost of creating the associated metadata, is relatively high. The latter are easy to use but lack a descriptive richness and hence do not permit reasoning. The challenge for ACTIVE is to combine these two approaches, to obtain as far as possible the benefits of both and minimise the disadvantages. For example, from observation of how tags are created and used we aim to learn additional structure, not defined explicitly by users, and hence to define and refine ontologies.

#### B. Knowledge delivery in context

The exploitation of context has already been the subject of significant research<sup>5</sup>. Context, though, means different things

---

<sup>4</sup> Much of the development of semantic technologies, and particularly their standardization, has been under the auspices of the W3C, see: <http://www.w3.org/2001/sw/>

<sup>5</sup> E.g. see the *Sixth International and Interdisciplinary Conference on Modelling and Using Context, Context 07*, <http://context-07.ruc.dk/>

to different researchers. To some it means primarily the environment in which the person is working, e.g. in or away from the office, and the tools he or she is using, e.g. laptop or PDA. To others, context is concerned with the specific task an individual is performing. Our interest is with this latter meaning, and to make this clear we use the phrase 'task context'. The challenge here is to learn task context from the user's behaviour, e.g. web-browsing, email usage; to describe task context semantically; and to prioritise information delivery according to that context.

### C. Knowledge processes

To date, the majority of research relating to processes has had to do with formal business processes. There has been some discussion of knowledge processes, often under other names e.g. [3]. Our challenge is to enable such processes to be created and edited by users, and also to create process descriptions automatically for users. The latter means (automatically) learning from a user, or users, what behaviour constitutes a commonly used process. As with task context, we are describing these knowledge processes semantically.

### D. Three interlocking themes

We do not see knowledge sharing, task context and knowledge processes as being independent themes. They all three interact closely with each other. Knowledge sharing needs to take account of the particular current task context, and of the current process being negotiated by the user. Both task context and current process influence what shared knowledge can be presented usefully to the user. Conversely, both knowledge processes and task contexts are themselves shareable. Certainly, a big advantage of describing knowledge processes is being able to share them with co-workers. Similarly, context, i.e. the knowledge that a set of information objects are all related and relate to a particular task, is also something which can be valuably shared. Finally, the way a user negotiates a process can be a clue to the current task context; and knowing the task context will assist in guiding the user through a process.

The technology also has to come together at the user interface. For this we have the ACTIVE Knowledge Workspace (AKWS). This does not replace any existing user applications; we expect that users will continue to use their current applications. The AKWS sits alongside those applications and influences how they present information to the user.

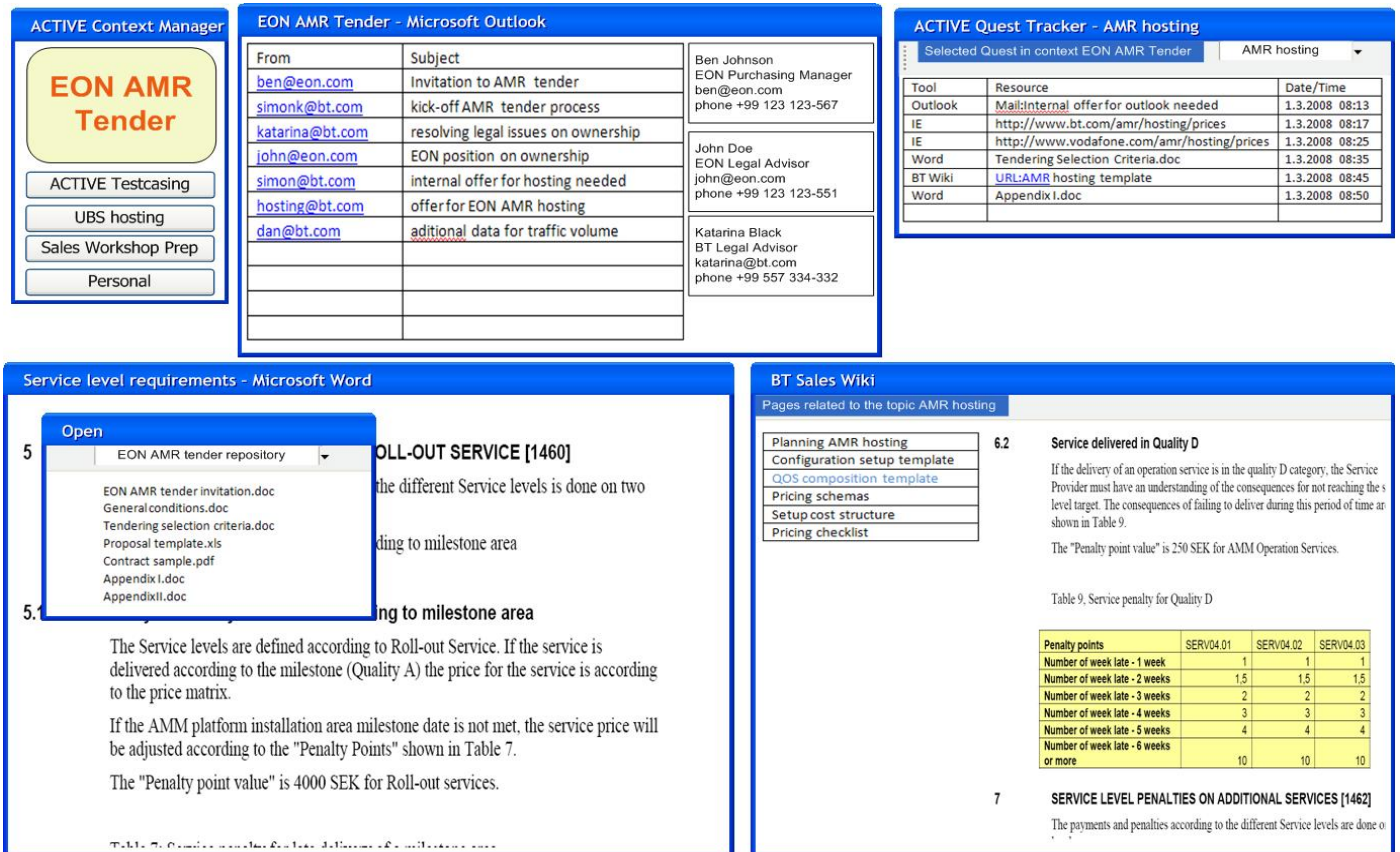


Figure 1 ACTIVE Knowledge Workspace user interface

Figure 1 shows a conceptual design of the AKWS user interface. The top left box is the ACTIVE context manager. This shows that 'EON AMR Tender' is the currently active context. Other contexts available to the user are 'ACTIVE Testcasing', 'UBS hosting', 'Sales Workshop Prep', and 'Personal'. The user would be able to switch context by clicking on one of these; or the system might determine from the user's activities that a context switch has taken place. Within the Outlook application, top centre, the emails either all relate to the current context, or are ordered according to their relevance to that context. The same principal applies for the other applications shown in the Figure.

#### IV. APPLYING THE TECHNOLOGY

Our technology is applicable to knowledge workers in all sectors. However, we are initially applying it in three sectors: consultancy, telecommunications and engineering. These three case studies all present particular challenges and together they are enabling us to thoroughly validate the technology, both from the technical, user and organizational viewpoints.

The three case studies are in many ways complementary. The consultancy case study will very much test the scalability of our technology. The telecoms example has similar requirements to that of consultancy, with a particular need to source information from a number of different systems. Finally, the engineering case study has very specific requirements related to a very specific engineering situation. This case study will validate that our technology can be applied to specialized knowledge worker roles.

##### A. Consultancy

Our consultancy case study is based within Accenture (<http://www.accenture.com>), a global outsourcing and consulting company. Accenture's business derives largely from large-scale client projects in the areas of technology, business process outsourcing, and business consulting. While Accenture's portfolio is extremely broad, a few notable examples of specific types of client engagements include systems integration, application renewal, application outsourcing (e.g., for supply chain applications), business process outsourcing (e.g., for help desks), business strategy consulting, customer relationship management, and marketing. Accenture has clients in a wide variety of industries: for example, Accenture services clients from retail to government to the financial and automotive industries. In effect, every major industry is represented in Accenture's client base. Accenture currently has about 180,000 employees around the world. As a consequence, as noted above, this case study will significantly test the scalability of our technology.

The Accenture case study supports three particular user activities:

- non-task specific knowledge inquiry;
- proposal writing;
- creation of marketing material.

##### 1) *Non-task specific knowledge inquiry*

Most enterprise knowledge management systems, including Accenture's, are set up for the "prototypical" user with no specific task in mind. This is necessitated by the fact that Accenture is a large and complex company, which results in its employees needing information for various different tasks. For instance, they may want to find information on previous Accenture projects, they may want to get an overview of a specific technology, or they may be interested in learning about a particular group within Accenture, such as Accenture Technology Labs. Also, they may need information within the context of various enterprise tasks; two of which are described in the following sub-sections. Because of the variety of tasks, as with any search system, it is very important to provide a "general-purpose" knowledge management system with no specific user or task in mind, general enough that it can be used in many different contexts. Of course, even in striving for maximum generality, we still want to direct our effort towards the typical Accenture user, rather than, for instance, the typical internet user.

##### 2) *Proposal writing*

As most consulting companies, Accenture wins contracts from clients by responding to Requests for Proposals (RFPs), which requires writing a proposal. Such proposals include a discussion of the current problem, how Accenture plans to address it, why the client should choose to hire Accenture rather than its competitors, the experience Accenture has in the given space, etc. The current practice of proposal writing at Accenture involves the use of standard processes on the one hand, but also a variety of ad-hoc mechanisms on the other. Accenture has a set of standard processes called Accenture Selling Methods. Those methods are collections of guidelines for common selling tasks, such as proposal writing. The methods guide the employee through the task by giving instructions about what information to gather, what documents to create, and so forth. Every employee is strongly encouraged to use this methodology. but it is also supplemented by other practices: calling contacts that may have information on similar projects, using search to locate similar proposals, using instant messenger or email to gather additional information from colleagues and friends, and so forth. This is in line with the goals of ACTIVE that merge the top-down formal business processes with bottom-up informal processes that individuals in enterprises tend to use. In ACTIVE we will be investigating a number of research directions which we believe can make the proposal writing process more efficient by supplying users with information that is particularly useful in this context. For instance, we can allow users to find experts, which in turn will help them find relevant information.

##### 3) *Creation of marketing material*

Accenture Marketing, a division of Accenture, frequently designs marketing information for specific clients. For instance, a variety of marketing material may accompany a proposal that is submitted to a client. Such marketing material will commonly include information about previous projects that Accenture has done for the same client, credentials

information, as well as general information about Accenture and its successes. We plan to improve the ability of people to create these marketing materials by helping them find better information faster. Oftentimes, information about previous projects cannot simply be obtained from the central Accenture data repository. This is because many documents created within the context of a client project are client-proprietary and may not be shared with all Accenture personnel, and so information does not always reach those that should get access to it. The second challenge is the duplication of work. Within marketing, many client-specific artifacts are created that could be reused in other projects. This, however, is sometimes difficult, simply because people often create similar materials without each others' knowledge. While related to the first area for improvement, the two are not exactly the same, and should be addressed in a different way.

### *B. Telecommunications*

BT Business, a business unit of BT Retail<sup>6</sup>, is tasked to supply integrated Information Technology (IT) and communications services and solutions to UK-based small to medium enterprises (SME)<sup>7</sup>. It is estimated that the UK's SMEs total spend on their IT and telecoms needs is around £18bn a year. Accordingly, there are significant opportunities for BT to grow its revenues in this sector. However, unless BT can become more responsive and get better and quicker at turning sales opportunities into contracted business, significant revenue generation opportunities may be missed.

At a senior management level, there is a strong awareness that maximizing collective knowledge and skills across BT Business, through, for example, more effective knowledge sharing and reuse, has the potential to reduce the length of the sales cycle significantly<sup>8</sup>. The BT case study will therefore focus on the knowledge transfer and information needs of the front line sales people that operate in the 'Corporate Sales' and 'ICT & Solutions' units of BT Business, that is the technical consultants, the solution consultants, and the sales specialists working on current sales opportunities.

Front-line sales people tend to work on multiple projects and need to be agile in their response to various work demands. In general, there is an increasing demand on the use of their time. As a result, they need to be able to switch quickly between tasks and have all information of relevance to their current task easily at hand; they should not have to search around for

such information as they change from one working context to another. Currently, within BT, much of the information and documentation needed to support a sales opportunity is held in multiple systems, and is not easy to access quickly. By taking into account the context in which people work, ACTIVE technology will give access to relevant information and knowledge more easily, for example it will be used to prioritise the presentation of sales, product and design documentation, and relevant emails. ACTIVE technology will also identify and put people in contact with other experts in the relevant area.

Customer relationship building is also a key area of focus. BT's sales specialists and sales consultants are encouraged to build an informed knowledge of their customers, for example to understand their customers' principal areas of business, their operations, their goals and, where possible, their business plans. Currently, the information that aids this understanding is held in a number of different systems and information sources, and takes time to compile. In order to help the specialist or consultant enhance their customer knowledge, ACTIVE technology will be used to collate customer-specific information according to a person's work-context.

There is considerable potential for people operating in various roles in BT Business 'Corporate Sales' and 'ICT & Solutions' to re-use existing solutions, or at least elements of those solutions, in support of new sales opportunities. The key problem here is that people may not be aware of prior relevant work, and so it is quite possible that time is expended unnecessarily designing solutions that are similar, at least in part, to those that have already been worked on. A better awareness of current solutions to specific business problems and the business domains in which those solutions are being applied should enable common solution patterns to be identified, and enable reuse of previously deployed solutions. This has the potential to reduce the sales cycle time considerably.

Finally, there is a wealth of tacit knowledge within BT which could be used to support new sales opportunities and which is not currently being used to best effect across BT Business. ACTIVE technology will therefore be used to convert tacit and currently unshared knowledge, the so called 'hidden intelligence' of BT, into transferable, interoperable and actionable knowledge that will support more effective collaboration across the various units of BT Business involved in a sales opportunity.

### *C. Engineering*

Cadence Design Systems GmbH (<http://www.cadence.com>) is an engineering design services provider in the domain of microelectronics and integrated circuits. Cadence's case study deals with eliciting tacit design process knowledge from the everyday practice of design engineers in this engineering design domain. The objective of the case study is to develop and validate a system for acquiring and managing explicit and

<sup>6</sup> For information about the BT Group, see <http://www.btplc.com>; for information about BT Retail in particular, see <http://www.btplc.com/Thegroup/Companyprofile/Groupbusinesses/BTRetail>

<sup>7</sup> SMEs are characterised by their diversity, for example an SME could be anything from a start-up or 'micro-business' with 1-10 employees through to a substantial 'medium-sized business' having between 250 and 500 employees.

<sup>8</sup> In the context of this paper, the sales cycle is the period from initial customer engagement to contract sign-off.

finely grained knowledge about electronics design (in terms of methodologies, processes, activities, actors, best practices, software tools etc.) Though design technology in the field is well defined, the processes still remain very stochastic. Some important reasons for process dynamics are: (i) design goals may change in the course of a process; (ii) involved designers may have different skills, capabilities and tool preferences; (iii) even the order of some phases (activities) can not be deterministically pre-set for any possible type of an artifact under design. A design process therefore develops dynamically. At every step, a decision needs to be made concerning which of the admissible continuation design paths is more productive. The challenge is that the knowledge to support such decisions is tacit, subjective and can not be fully elicited and made explicit by interviewing designers. In particular, designers often employ their intuition and experience when reasoning about facilitation dependencies among design activities situated in design environments.

Our approach in ACTIVE is to record the performance of engineering design processes in the everyday practice of experienced designers. Several concurrently and autonomously executed design processes will be tagged by the ACTIVE software components injected in the Cadence Workbench software. The components produce process tags in the form of “ticks”. These ticks are stored in the knowledge base. Many design engineers, distributed over several development teams and several engineering design processes, will use the ACTIVated Cadence Workbench in their work. Therefore, we may consider such a usage as a collaborative engineering design process tagging. Process mining components use this “ticks” knowledgebase to devise the sequences of atomic activities and compound tasks executed in the tagged processes. Process execution paths are thereby recorded in the knowledge base. At each point of decision making a designer has to choose the best productive continuation path for his current process using activity or task patterns prescribed by the technology or in-house policy. Several admissible alternatives may appear to be possible. For those, advice is solicited from the ACTIVE Workspace software. The ACTIVE Workspace suggests the most highly ranked continuation path appropriate for the context of the current process using performance comparisons of other processes in similar contexts, already recorded in the past.

This case study draws on previous work by Cadence in which background knowledge was developed to formalize loosely defined, dynamic, structurally ramified, time bound, stochastic processes of engineering design in microelectronics. In this work both a modelling framework and a suite of ontologies were developed. The core and extension ontologies of the suite [5, 4] are used for representing knowledge about dynamic engineering design processes and design systems as process environments comprising actors, resources and tools. The upper-level ontology [6] serves as a semantic bridge to foundational theories representing common sense. The modeling approach is based on the understanding that an engineering design process is a process of knowledge transformation which passes through several states. Each state

is the state of affairs in which a particular representation of a design artifact or several representations are added after being elaborated by a design activity leading to this state. Evidently, the overall goal of a design process is reaching the target state of affairs in which all the representations are elaborated in a way to meet the requirements. Each time the continuation of the process is decided by choosing an activity from the set of admissible alternatives for that state. Engineering design processes are situated in and factually executed by the design system comprising designers, resources, tools, normative regulations.

## V. NEXT STEPS

The ACTIVE project is currently at an early stage; in parallel with initial technical activity by the technology providers in the project, the case studies have been undertaking a more detailed understanding of user requirements. The development philosophy is cyclical. That is to say, at the end of each year case study prototypes will be developed to obtain user feedback. The objective of the feedback phase will be to validate the technology technically, and also from the user and organization perspective, i.e. by the benefit it brings the user and the user’s organization. The validation phases will occupy two to three months at the end of each of years 1 and 2. In the final six months of the project there will be a fuller validation phase. The overall development philosophy is illustrated in Figure 2.

The next step therefore is to bring the current technology developments together to create initial case study prototypes which will be presented to the case study communities late in 2008 and early in 2009. Feedback from this exercise will then influence developments during the second year of ACTIVE.

## ACKNOWLEDGMENT

The ACTIVE project (IST-2007-215040) is part-funded by the European Commission. The authors would like to thank Igor Dolinsek, of Hermes SoftLab for the conceptual design of the AKWS user interface shown in Figure 1

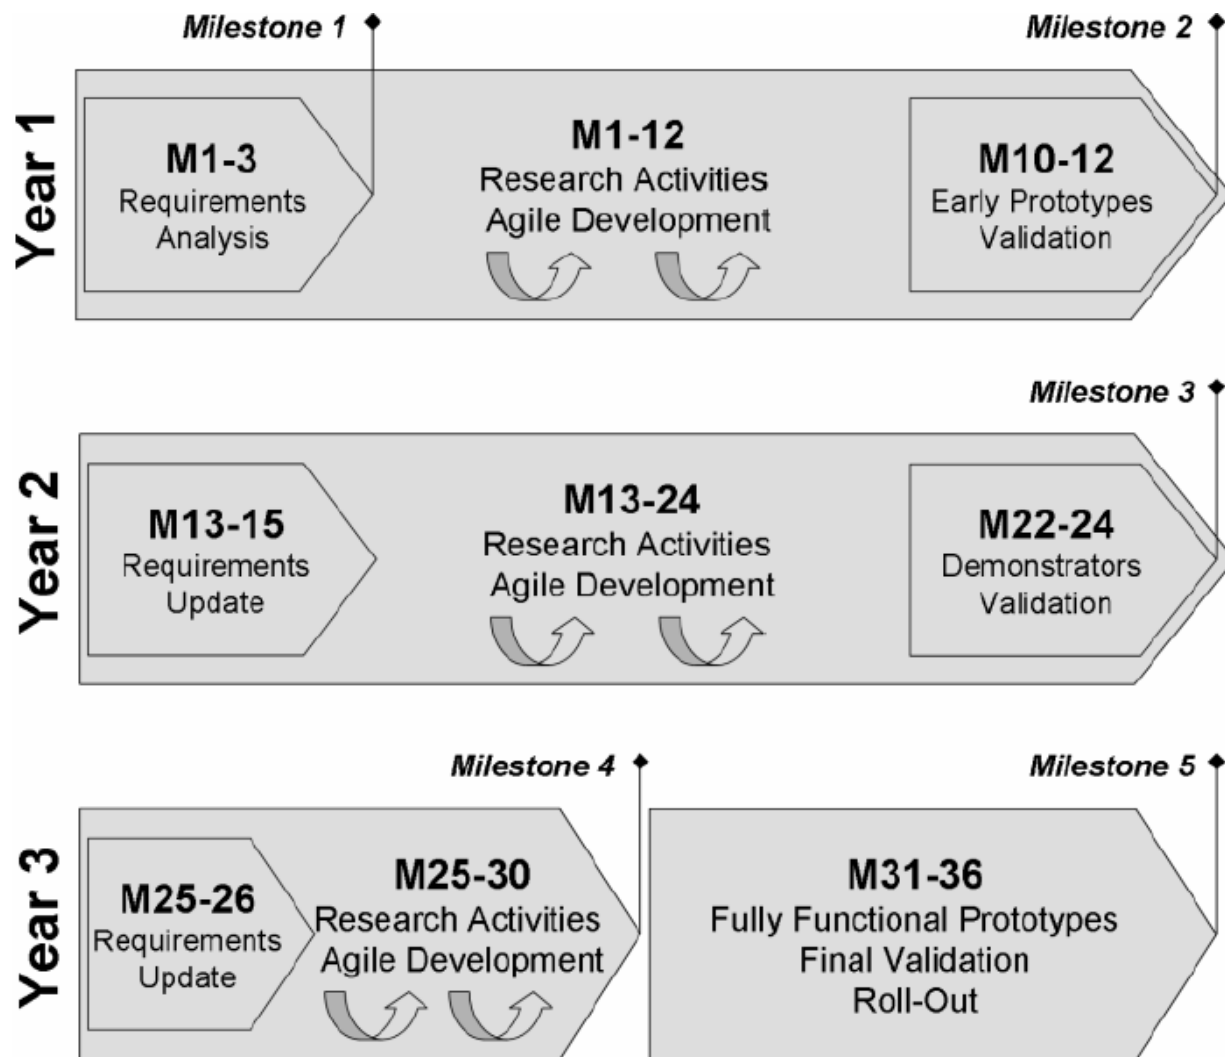


Figure 2 ACTIVE development methodology

#### REFERENCES

- [1] Drucker, P.F. 1969 Management's New Role, Harvard Business Review pp. 49-54
- [2] Davenport, T.H. Thinking for a Living, Harvard Business School Press, Boston Massachusetts
- [3] Hill, C., Yates, R., Jones, C., and Kogan, S.L. 2006 Beyond predictable workflows: Enhancing productivity in an artful business process, IBM Systems Journal, Vol. 45, No. 4, pp. 663-682
- [4] Ermolayev, V., Keberle, N., Matzke, W.-E.: An Ontology of Environments, Events, and Happenings. In: Proc 31st IEEE Annual International Computer Software and Applications Conference ([COMPSAC 2008](#)), Turku, Finland, Jul. 28 - Aug. 1, 2008, 539-546
- [5] Ermolayev, V., Jentzsch, E., Karsayev, O., Keberle, N., Matzke, W.-E., Samoylov, V., Sohnius, R.: An Agent-Oriented Model of a Dynamic Engineering Design Process. In: Kolp, M., Bresciani, P., Henderson-Sellers, B., and Winikoff, M. (Eds.): Agent-Oriented Information Systems III. Revised Selected Papers, 168-183, 2006, LNCS Vol. 3529
- [6] Ermolayev, V., Keberle, N., Matzke, W.-E.: An Upper-Level Ontological Model for Engineering Design Performance Domain. In: Q. Li et al. (Eds.): ER 2008, LNCS 5231, pp. 98-113, 2008.