Journal of Universal Computer Science, vol. 8, no. 1 (2002), 85-96
submitted: 3/12/01, accepted: 24/1/02, appeared: 28/1/02

On a New Powerful Model for Knowledge Management and its Applications[1]

Hermann Maurer
Graz University of Technology, Graz/Austria
and Know Center Graz/Austria
hmaurer@iicm.edu

Klaus Tochtermann
Know Center Graz/Austria
ktochter@know-center.at

Abstract: In this paper we present the Maurer - Tochtermann Model for Knowledge Management (KM) and present strong evidence that this model has powerful ramification. First, it shows clearly that KM is not just "old wine in new bottles" but an important and new area of research and applications; second, it shows clearly where KM differs from classical distributed information systems or data bases; third, it is shown to embrace a number of pragmatic problems that have often been considered the heart of KM; and fourth, it gives a clear indication of the areas that will be of increasing importance for KM in the future. We claim that the model can and should be the basis of future efforts in IT-oriented KM.

Keywords: Knowledge Management, Information Systems, Data Bases, Document Management, Intelligent Agents

Categories: H.1, H.2, H.4

1 Introduction

After having been a very scientific discipline for many years, Knowledge Management (KM) has become focal point of much research, applications and commercial interest since about 1998. However, there is much discrepancy on what KM really is: attempted definitions have sometimes stressed organisational components behind KM with little or no emphasis on information technology, others see KM mainly as a way to measure the value of the "human component" in an organisation, i.e. see knowledge assessment as central issue, and a third school of thought sees information technology as the central aspect of KM. Even those who emphasize IT end up with a range of attempted definitions. Thus, there is no uniform agreement on what really constitutes KM, as can e.g. be seen from the different approaches taken in publications such as [Woods and Sheina 98], [Studer et al. 99], [Karagiannis and Telesko 01], [Davenport and Prusak 98], [Sivan 99], and [Ives et al. 98]. This of course implies that there is no universally accepted definition. Worse, most definitions, if at all given, are rather wishy-washy. It is the claim of this paper

[1] A preliminary version of this paper was prepared in German for a keynote address at the Bielefeld Conference, Feb.5- 6, 2002

that another approach is more illuminating than most previous attempts: based on what are often considered the main issues of KM the Maurer - Tochtermann Model for KM is presented: it is shown to clarify many open issues; it demonstrates impressively in which way KM goes beyond traditional approaches to information management in information systems or data bases; and it makes it clear what (substantial) agendas still have to be resolved as KM is getting more and more powerful.

Lest we mislead readers of this paper we want to clarify our position concerning KM: we think that in actual applications organisational aspects of KM play a large role: they will thus not be ignored in our model, yet we do not delve into any details in this paper; we feel that knowledge assessment (KA) is a topic outside KM: results of KA may influence KM and conversely. However, we do not treat in our model KA as integral component of KM. Thus, the Maurer - Tochtermann Model (MT-Model for short) of KM is, on purpose, IT centric.

One further point is worth mentioning: we believe terms such as “KM” and “KM-Systems” are hard to define in detail, since even the term “knowledge” is very elusive. The distinction made between data, information and knowledge is often fairly artificial, fuzzy or both. After all, when we give you the *information* that the Knowledge Management Center (Know-Center) in Graz was conceived in 1998, you now *know* (although “2000” is really a simple *data*- item) that ideas on KM have been around for quite some time: this seems to indicate that a possible definition of “knowledge” could be “information in context” or “structured information”. Clearly, this is not terribly satisfactory either: we all *know* that “knowledge” is more than this. However, this paragraph and the play with words shows that since “knowledge” is already hard to define, how can we expect to really nail down things like “KM” or “KM-Systems”? It also shows that as first approximation equating knowledge with linked and structured information is at least a good crutch. Hence we will keep this idea in the background of our minds throughout.

The main part of this paper is structured as follows: in the next Section 2 we present pragmatic situations showing what KM is all about; in Section 3 we present the Maurer - Tochtermann Model (MT-Model), followed by core techniques of KM in Section 4: those techniques shed further light on the MT-Model. Section 5 gives a brief summary, followed by a number of references in Section 6.

2 Pragmatic Starting Points for KM

The frustrated statement of some managers: “ If our employees only knew what our employees know, we would be a perfect company” illuminates one of the central aspects of KM: a group of persons always knows more than any single individual, and even those individuals having similar knowledge may use and view their knowledge in rather different ways. The challenge that derives from above statement is clear: how is it possible to get at least part of the knowledge residing in the brains of people into some kind of networked computer system, subject to two major constraints: first, to get the knowledge out of the brains should not create more than at most a very modest effort for the persons involved; second, the knowledge should be made available to others “actively” when they need it, without requiring an explicit request for the knowledge: this latter constraint is clearly important, since persons will often not know that such knowledge exists, and hence never search for it.

Suppose for a moment that some organisation is capable of achieving above kind of KM: the benefits would clearly be enormous. It would:

- avoid duplication of work
- support collaboration between persons
- avoid loss of knowledge if some person becomes unavailable
- ease the training of new employees
- let people learn from other persons' successes or failures
- etc.

It will be shown later that techniques to achieve at least part of the above aims do already exist, and are actually built into modern KM systems such as Hyperwave [Maurer 96], [Hyperwave], [Hyperwave 99]. We will call this kind of KM "KM for Organisations" in what follows. Note also that in many cases much reduced scenarios can already be very useful: for instance, it may not be necessary or feasible that all employees of a company know everything, yet an employee working on some topic x should at least be able to find out whether some other employee has already experience with x. This kind of KM is clearly much less ambitious, has been realized in a number of contexts, and solutions are often referred to as defining "Knowledge Domains" [Helic et al. 01] or "Yellow Pages".

Conversely, it is conceivable that KM goes far beyond the boundaries of a company or organisation, but applies to much of our society: combined with ubiquitous computing it is not totally unrealistic that much of mankind's knowledge is available to everyone at some stage in the future. Putting it differently, the knowledge outside each individual's brain is of course much larger than in the brain of any individual: this knowledge outside may one day become a veritable extension of the human brain using sophisticated future techniques of KM. Why and how this might work is e.g. discussed in [Maurer 01].

There is another pragmatic approach to KM that is based also on a certain amount of frustration, but in a different environment: anyone who has ever had to deal with large amounts of data coming from many heterogeneous sources will have sighed more than once: "If we were able to somehow automatically classify and associate incoming information with existing material we would finally have solved the problem of information archival." We will call this problem "KM for Archives".

Let us have a brief look at two examples to explain the situation:

In the Journal of Universal Computer Science, J.UCS, [JUCS], [Krottmaier and Maurer 01] papers are not just classified according to the ACM system and hence allow complex queries, but also provide for "links into the future". This term refers to the following fact: if a contribution A was written e.g. in 1955 and a new paper B in, say, 2002, refers to A, then in any digital library [C.ACM 98] B will clearly have a link to A, i.e. a "link into the past". In J.UCS, however, a link from A to B is added, thus providing a "link into the future", specifically from 1955 to 2002. Such new types of links (that will clearly ease finding the most recent paper on a certain topic) can be generated in a digital library fairly easily if literature references follow a well defined format. However, even in the absence of a standard for references they can be generated, and generated beyond the boundaries of one digital library using tools such as the Citation Index. Techniques like “similarity recognition” or “recognition of connections” as will be discussed later will provide still more powerful ways to add “links to the future”.

Another example is the electronic version of the largest German encyclopaedia, the Multimedia Brockhaus Premium ['Brockhaus 01]. For each contribution a “knowledge net” showing related contributions (and displayed in a graphical fashion) is generated automatically. This is done using the already mentioned “recognition of similarities”, and is supported by fine grained metadata [Duval 01]: each entry is associated with one or more categories thus avoiding cross links between contributions that do not at all belong together. Indeed the “knowledge net” will be generated by 2003 based on a complete ontology [Uschold and Gruninger 96] of the German language, making cross-linking still more precise.

As will become clear in the next section, “KM for Archives” can be considered as a subset of “KM for Organisations”. For this reason, we will concentrate only on the latter in what follows.

3 The Maurer-Tochtermann Model (MT-Model) for KM for Organisation

The Maurer - Tochtermann Model (MT-Model) for KM, as shown in Figure 1, has been introduced in similar fashion as “communication model” by the authors before [Tochtermann and Maurer 2000]. Work with it has now matured to a point that it is worth- while to present the final version.

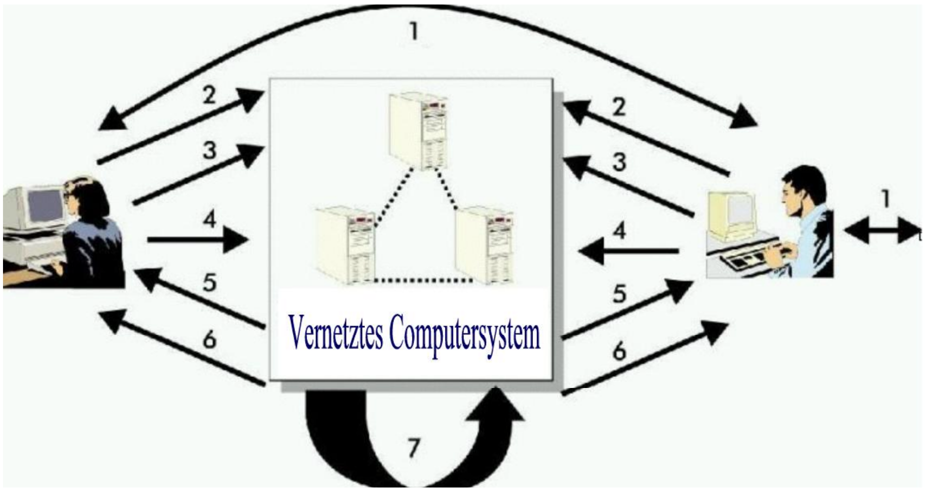

Figure 1: The MT- Model for KM

Figure 1 shows a group of persons who exchange knowledge with each other. A large amount of this exchange uses a networked computer system (and hence enables asynchronous use). The arrows labelled 1 through 7 have each a special significance, and indeed represent the core ideas of the model.

We start with a rough discussion of the meaning of the arrows and go into more detail, afterwards.

Arrow 1 indicates that persons can communicate directly (in coffee- breaks, at an office party, through a telephone call, …):

It is this arrow 1 that subsumes all organisational aspects of KM, and there are indeed many of them. In this paper and this model the fact that we leave a "slot open" for the rich material on organisational aspects of KM will have to suffice.

Arrows 2 to 4 symbolize the various ways in which information (knowledge) can enter the KM-System. In contrast, arrows 5 and 6 indicate that there are two very distinct ways how knowledge can pass from the system to users. To be more specific, arrow 2 stands for the explicit input of information into the KM-System, much as data is usually entered into an information system. Arrow 3, however, symbolizes the implicit input of information into the KM-System: information and knowledge is entered into the system as by-product of activities users would be doing, anyway: here, new knowledge is created without burdening the user, a very important aspect. Arrow 4 indicates that a KM- System can also "systemically" create new knowledge by observing users. Arrow 5 symbolizes the traditional query as used in classical information systems or data bases: users formulate such queries in some way (by using a query language, by filling out a form, by clicking on a succession of links,…) and obtain chunks of information from the KM-System: such chunks can be small pieces of data, or large coherent documents like manuals, books or even courseware. Arrow 6 is more unusual: it indicates that the system can generate and offer knowledge without being explicitly asked by the user. Arrow 7 symbolizes the fact that a good KM-System is able to generate new knowledge based on existing one.

Figure 1 shows very clearly the difference between classical information systems (data bases) and KM-Systems: if we omit arrows 3, 4, 6 and 7 in Figure 1 we have a classical information system! Thus, KM-Systems go beyond traditional systems if the actions indicated by arrows 3, 4, 6 and 7 can indeed be implemented. We will prove that this is indeed true in the next section using concrete examples.

Figure 1, a model for KM- Systems for organisations is also valid for "KM for Archives" if we just ignore arrows 1 and 3 in Figure 1. Thus, our earlier claim that it suffices to study "KM for Organisations" is justified: indeed, this is not very surprising: after all, KM for any organisation will involve a substantial body of material that could well be called a (specific) digital library or digital archive.

4 Some Current Techniques in KM

In this section we indicate how the "critical" arrows 3, 4, 6 and 7 that distinguish a KM- System form classical information systems can be realized.

Arrow 3 in Figure 1 symbolizes the implicit input of information, i.e. the generation of new knowledge as by-product of actions that would occur, anyway. The list of such actions is quite lengthy, and a few examples will have to suffice. There are simple actions like sending an announcement of some event to a group of persons by email. Performing such an action in conjunction with a KM-System will send the announcement also to the system, into a folder "Upcoming events", open to the public and sorted by date. When an event is over (it has its end-date as default expiration date as "metadata") the announcement is shifted into a folder such as "Past events of year xxxx", and at year's end is moved again to a new folder, e.g. "Events of the last ten years", a list that may well come in handy for a yearly or a ten-year report! What has been said for the announcement of events holds, of course, for all information that is put on a Web-server of the organisation, from telephone- directories, to the structure of the organisation at issue, to the tasks of various subgroups, etc. Note that no sizeable organisation can live without ISO 9000 certification anymore, today. To remain certified, extensive documentation about each project, persons, resources and tools involved, the time-lines, milestones, documentations, minutes of meetings, etc. etc. have to be gathered, anyway: all this must become part of the KM-System. Such information, properly interlinked (see below) presents valuable knowledge about the organisation and projects carried out. It also provides much insight in general, e.g. by allowing to determine why past projects worked out well or ran into difficulties. The work- flow of an organisation is also available in electronic form today in most organisation and should be integrated into the KM-System. The same is true of yearly reports, of lists of products with description and pricing information (various types of information accessible only with appropriate authorisation), manuals and other internal reports. Most important, existing information systems and data bases have to be integrated, leading to so- called knowledge portals [Hyperwave 99]. Email should also be administered in the KM-System centrally, clearly with suitable authorization techniques, rather than being handled by each user separately. Why it is important to bring all this information together will become clear in what follows: only if we have a substantial body of information does it make sense that KM- Systems create automatic linkages between pieces of information and classify documents according to potentially a multitude of views. This generates structured information and hence knowledge according to what was said in Section 1.

Arrow 4, symbolizing systemic generation of knowledge deduced from observing behavioural patterns is currently probably the weakest part of all KM-Systems. The basic idea is that inputs, coming from specific sources (data bases or employees) allow the derivation of general rules and procedures that can be made available in similar situations. Such rules are often intertwined with actions symbolized by arrow 5 (explicit queries): a KM-System will e.g. realize that certain search paths are used by some persons over and over again, and hence might provide abbreviations, generate bookmarks automatically, or note that when information x is retrieved, often y is also of interest. As a consequence, a user retrieving x is automatically offered (in the sense of arrow 6) the item y. Observe that many current software packages, including Winword which the authors are using right now, help automatically in many situations (sometimes to the chagrin of authors who wish they could temporarily turn off this function): this approach corresponds to arrow 6, and to arrow 4, if the system is actually trying to learn from user behaviour.

To be able to better explain arrows 6 and 7 it is convenient to mention a few of the techniques that are essential components of a KM-System.

One of the more important concepts in this connection is the notion of "active document" introduced by the first author some time ago. The idea behind this concept is this: whenever a user sees a document on the screen, the user can ask an arbitrary question (e.g. by typing it, or maybe even by using a microphone!) and the *system* provides the answer immediately.

Of course it is impossible that any KM-System can answer any question formulated. However, there are two approaches that deliver, from a pragmatic point of view, results that are fairly close to what has been specified. First, it is sometimes possible to convert a question into a data base query and have the data base answer it. We will not go deeper into this interesting area of "natural language data base queries" here but rather point out the second less well-known technique: if a question is asked more than once for a particular document it can be answered the first time by an expert, possibly asynchronously. The question and answer are stored in the KM-System. If a semantically equivalent question is asked later, the KM-System has "only" to determine this equivalence to present the answer. This approach has proven particularly useful if a document remains "stable" over a reasonably long period and is used by many persons (this is e.g. true of some manuals, general information and often courseware!).

The problem with this approach is clearly how to recognize that two pieces of text x and y are "semantically equivalent", i.e. in our case represent the same question. There are many approaches to this problem.

One is just to compare the frequency of words in the two documents at hand: this technique can be improved by excluding trivial words (like "the" or "and"), by using stemming to reduce nouns to first case singular and verbs to their infinitive, by using synonym- dictionaries, or by even using semantic nets [Meersman et al. 99] and ontologies. If it looks as if the question x now asked means the same as a former question y, the user is presented with question y and a text such as "Is this what you mean?" If the answer of the user is affirmative, the system presents the answer previously given to y. If the answer is negative, and another previously asked question z also looks similar, the process is repeated with z, and so forth, until alternatives run out. At that point the user gets feedback of the type: "This is a very good question. You will receive the answer by email as soon as possible". In this case, an expert will answer the question (potentially some time later), and of course this new question x and the answer are now also entered into the data base for future use. Putting it differently, this approach employs the intelligence of the user to make the final decision whether two questions are indeed identical in meaning or not. The method can be further refined by using approaches from artificial intelligence such as case based reasoning, but none of those have had a wider impact, so far.

Of course it would be more elegant if one were able to *prove* that two pieces of text x and y are semantically equivalent: however, with current day and foreseeable techniques this is not possible in general. It is feasible only if the area of discourse and the syntax of how the text is written are severely restricted [Heinrich and Maurer 2000].

However, there is one other very simple approach for providing for of active documents. It could be called "local FAQ's". More specifically, if a user has a question concerning a document currently on the screen of the computer the user marks the sentence, formula, graph or whatever at issue and types a question. This question is sent by email to an expert (including the document at issue to make it easy for the expert to answer), and the expert sends the answer back: during "office hours" immediately, else at some later stage. The main point is, however, that the area highlighted by the person asking a question will now be preceded by a special icon that indicates to later readers "someone has asked questions here and expert answers are available." Thus, a user coming later and having a query concerning the highlighted material will first click at the icon mentioned to see the question- answer dialogues that have taken place. Only if the question of interest is not among the ones posed before, does this new user also type a question, thereby increasing the number of questions and answers attached to this particular area of the document. Note that such “local FAQ lists” will never get long for two reasons: first, there won’t be that many questions concerning a small fragment of a document; second, if many are asked, the system hopefully reports this to the author since such a large number of questions indicates that something is at miss with the explanations given. Putting it differently, this technique is an easy solution for active document and also provides excellent feedback for authors. Note that this method will not work if documents are very much time dependent (hence each question- answer dialogue may have an expiration date attached to it.). Observe further that this approach is particularly useful if very large numbers of persons are going to read the same document. In one of the first applications of this technique a company with 150.000 employees provided extensive information and learning resources for everyone. The local FAQ’s generated “stabilized” very rapidly: after the first 600 persons had read the material only 0.03% of the other employees posed new questions. As a result, experts were made available initially around the clock to answer queries with little delay until the first 600 employees had worked through the material. Of the remaining other 149.400 employees (more than 99.6% of all employees) only 45 asked new questions, justifying “time delayed” answers in those few cases.

Above discussion should make it clear that no good KM-System can do without active documents. Hyperwave [Hyperwave] was the first system providing this feature, and has continued to improve it.

In connection with active documents we have already encountered the problem of determining whether two pieces of text or documents x and y are similar. A number of techniques to test documents for similarity are available, most based on an extension of the idea of checking for important identical words in both documents as mentioned above. Such methods allow KM-Systems to automatically classify documents (in the sense of arrow 7 in Figure 1), but also to actively notify users (in the sense of arrow 6 in Figure 1) about suspected similarities. This gives rise to many practical applications, as the four concrete examples that follow will show.

Example 1: Let us consider a large distributed company working on a myriad of development projects. If a new project is started in location A a document outlining the project as required by ISO 9000 standard will be prepared. This is automatically translated by the KM-System from the native language into (passable) English. This English version is compared with all other project descriptions in the company. If a similarity is discovered with a project carried out at location B, both A and B (and often also a supervisory agency) are notified by the system that some duplication of work might be about to happen. Of course the KM-System may be wrong in the sense that the similarities are irrelevant. Still, if only a small number of project duplications are avoided the gain is significant.

Example 2: In a company with sizeable research departments an employee A enters a new paper into the company’s digital library. Almost instantaneously A obtains the information that there are already two similar documents in the library, authored by B and C. At the same time B and C obtain information that a contribution similar to what they have written before has been entered into the digital library by A. Like in Example 1, this kind of approach minimizes the danger of duplication and fosters collaboration. Ideally, the search for similar documents is not done just in the company's digital library but also in other digital libraries accessible via the Internet, such as e.g. J.UCS [JUCS], [Krottmaier and Maurer 01].

Example 3: The KM-System checks all "non-private" emails (or all emails accessible to some group of persons) to discover similarities and notifies persons accordingly. This very important application requires fairly sophisticated deliberations concerning privacy and authorization!

Example 4: In a discussion forum some topic is started. The KM-System realizes that this topic has been discussed extensively before. It avoids an entirely new and repetitive discussion by pointing out the contributions in the forum made earlier.

Summarizing, the use of similarities of documents is one of the most powerful tools available in today's KM-Systems. Similarities and connections (see below) can often be shown graphically in a very intuitive way. Such a representation is often called a "knowledge- net" and is used extensively at various levels of granularity in e.g. the electronic encyclopaedia Brockhaus. Figure 2 shows such a knowledge- net at the level of course granularity as is generated automatically for each entry as mentioned earlier. In the example the knowledge net is shown for "Raumsonde" (i.e. "spaceprobe"). It could be refined (more entries generated) by clicking at the button "ERWEITERN".

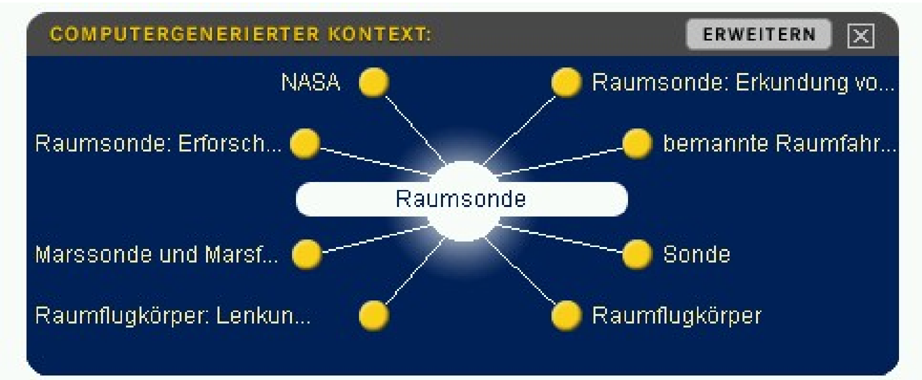

Figure 2: A small knowledge -net

A more complex variant of similarity recognition is what is sometimes called "connection recognition". Techniques for this are not generic enough at this time to be useable in general, hence they have to be customized for specific situations. Let us explain the basic idea by again using an example.

Suppose we have a very large collection of documents, e.g. all publicly accessible documents on WWW servers and data bases, including electronic newspapers, reports by news agencies, etc. The problem to be solved is to find out persons who are likely to have been in contact with some other person X.

If the KM-System finds somewhere "X stayed in Nassau October 15, 2000", this fact is entered into the "recognition data base" as e.g. (X, 15/10/2000, Nassau). When analysing the entry "Maurer made vacation on North Eluthera October 10 - 20, 2000" the KM- System does the following: (a) it recognizes "Maurer" as name of a person putting (Maurer, 10-20/10/2000, North Eluthera) into the recognition database; (b) it compares all entries in that data base with the currently added one. Since October 10-20, 2000 overlaps with the date October 15, 2000 and since "Nassau" is recognized as capital of the Bahamas and North Eluthera as one of the islands of the Bahamas, there is a chance that X and Maurer have met. As a consequence something like (X, Maurer, 1) is entered into the "X- connection database" indicating that so far there is one indicator that X and Maurer have some connection. If the KM- System finds "X has met the person Z" and "Maurer and Z went to school together" in potentially completely different documents at some later stage, the triple (X, Maurer, 1) in the X-connection database is replaced by (X, Maurer, 2). If the third component of this triple, the "counter of indicators", reaches some threshold, i.e. 100, i.e. the triple turns into (X, Maurer, 100) the system sends out an alarm: it is now very likely, that X and Maurer have some connection: this alarm is a typical action corresponding to arrow 6 in Figure 1. The computation of the triples mentioned corresponds to arrow 7.

It should be clear that the establishment of connections even in the example described is fairly complex. However, such techniques have proved invaluable in the past in the case of e.g. tracking down criminal activities.

There are a number of much simpler applications, e.g. in connection with e-Commerce. A typical example is to use the shopping habits of two persons A and B to find out that they have similar habits concerning books. If A buys a book from a new author, and shortly thereafter another book from the same author, the system would guess that A likes this author: this causes the system to point out the author at issue also to B.

Much research in this area of connection recognition is still necessary and going on. However, it is clear by now that connection recognition will be one of the major tools that must be supported or supportable by any good KM-System.

5 Summary

The Maurer - Tochtermann Model (MT-Model) that has been presented in this paper shows very clearly where KM-Systems differ from traditional systems and shows where progress has been made, and further progress is essential. The functions described that a KM-System has to satisfy, like active documents, similarity recognition, knowledge- nets, etc. are not science fiction features, but available in up-to-date KM-Systems such as Hyperwave [Maurer 96], [Hyperwave].

References

[Brockhaus 01] Der Brockhaus Multimedial 2002 Premium; DVD, Brockhaus Verlag, Mannheim (2001).

[C.ACM 98] Communication of the ACM: Special Issues on Digital Libraries: 38, 4 (1995); 41, 4 (1998); 44, 5 (2001).

[Davenport and Prusak 98] Davenport, T., Prusak, L.: Working Knowledge: How Organizations Manage What They Know; Harvard Business School Press, Boston (1998)

[Duval 01] Duval, E.: Metadata Standards: What, Who & Why; (eds. Tochtermann, K., Maurer, H.) Proceedings of Int. Conference on Knowledge Management I-KNOW 01, Springer Pub. (2001), 137-147. See also: http://www.jucs.org/jucs_7_7/meta_data_what_who.

[Heinrich and Maurer 2000] Heinrich, E., Maurer, H.: Active Documents: Concept, Implementation and Applications; Journal of Universal Computer Science 6, 12 (2000), 1197-1202. See also: http://www.jucs.org/jucs_6_12/active_documents_concept_implementation.

[Helic et al. 01] Helic, D., Maurer, H., Scerbakov, N. : Knowledge Domains: A Global Structuring Mechanism for Learning Resources in WBT Systems; Proceedings of WEBNET 2001, AACE, Charlottesville, USA (2001), 509-514.

[Hyperwave 99] Hyperwave: Hyperwave Information Portal - White paper; www.hyperwave.de/publish/downloads/Portalwhitepaper.pdf (1999).

[Hyperwave] www.Hyperwave.de

[Ives et al. 98] Ives, W., Torrey, B., Gordon, C.: Knowledge Management: An emerging discipline with a long history. Journal of Knowledge Management 1, 4 (1998), 269-274.

[JUCS] www.jucs.org

[Karagiannis and Telesko 01] Karagiannis, D., Telesko, R.:Wissensmanagement – Konzepte der Künstlichen Intelligenz und des Softcomputing; Lehrbücher Wirtschaftsinformatik, Oldenbourg Verlag (2001).

[Krottmaier and Maurer 01] Krottmaier, H., Maurer, H.: Transclusions in the 21st Century; Journal of Universal Computer Science 7, 12 (2001), 1125-1136. See also: http://www.jucs.org/jucs_7_12/transclusions_in_the_21st.

[Maurer 96] Maurer, H. (Ed.), HyperWave: The Next Generation Web Solution; Addison-Wesley Longman, London (1996).

[Maurer 01] Maurer, H.: Die (Informatik) Welt in 100 Jahren; Informatik Spektrum, Springer Verlag, (April 2001), 65-70.

[Meersman et al. 99] Meersman, R., Tari, Z., Stevens, S. (Eds.): Database Semantics; Kluwer Academic Publishers, USA (1999).

[Sivan 99] Sivan, Y.: The PIE of Knowledge Infrastructure: To Manage Knowledge We Need Key Building Blocks; WebNet Journal 1,1 (1999), 15-17.

[Studer et al. 99] Studer, R., Abecker, A., Decker, S.: Informatik-Methoden für das Wissensmanagement; Angewandte Informatik und Formale Beschreibungsverfahren, Teubner-Texte zur Informatik, Vol. 29 (1999).

[Tochtermann and Maurer 2000] Tochtermann, K. & H. Maurer: Knowledge Management and Environmental Informatics. Journal of Computer Science 6, 5 (2000), 517-536. See also: www.jucs.org/jucs_6_5/knowledge_management_and_environment.

[Uschold and Gruninger 96] Uschold, M., Gruninger, M.: Ontologies: principles, methods and applications; The Knowledge Engineering Review 11,2 (1996), 93-136.

[Woods and Sheina 98] Woods, E., Sheina, M.: Knowledge Management - Applications, Markets and Technologies; Ovum Report (1998).