The Times They Are A-Changin' — The Corporate History Analyzer

Jürgen Angele[b], Hans-Peter Schnurr[a,b], Steffen Staab[a,b], Rudi Studer[a,b]
[a]Institute AIFB, University of Karlsruhe, 76128 Karlsruhe
http://www.aifb.uni-karlsruhe.de/WBS
[b]Ontoprise Gmbh, Haid-und-Neu Straße 7, 76131 Karlsruhe, Germany
http://www.ontoprise.com

Come gather 'round people
Wherever you roam
And admit that the waters
Around you have grown
And accept it that soon
You'll be drenched to the bone.
If your time to you
Is worth savin'
Then you better start swimmin'
Or you'll sink like a stone
For the times they are a-changin'.

Bob Dylan, 1963

Abstract

Up-to-date corporate information research, i.e. the active tracking and management of knowledge relevant to one's business, is a major task for knowledge-intensive companies. While the correct analysis of market situations and competitors are critical requirements for success, the failure to provide adequate knowledge about one's business environment may incur large losses. In order to support the tracking of external business knowledge and its analysis, we have built CHAR, the Corporate History AnalyseR, a new type of system for gathering and analysing temporal data about companies. CHAR integrates facts about organizational structures and strategic activities into a common knowledge base. Building on an ontology with an elaborate temporal structure, knowledge pieces are amalgamated in order to derive a plentitude of summarizing views onto current and previous states of affairs.

1 Introduction

Global markets changing with an ever faster pace are one of the key topics in business nowadays. Keeping track of what changes in the markets and recognizing what moves one's competitors is vital to any enterprise that wants to succeed on a global scale. While this has been a truism for management at the topmost strategic level for a long time, the very same necessity is currently springing up on all management levels down to the expert knowledge worker who manages his own skills.

The trouble is that these management levels and professionals do have a hard time gathering information, analyzing it, and performing their operational work, such as researching and developing new products or streamlining their production lines. This means that the evaluation of what happens in one's own business may often be neglected.

Here comes corporate research into play. The task of the corporate researchers is the tracking of relevant knowledge in the outside business setting and the communication of important knowledge to stakeholders within the company. The general difficulty, then, is the delivery of knowledge to the right persons at the right time and at the right level of abstraction. People do neither want to be flooded with facts that they do not need nor do they want to be swamped with nitty gritty details, but they rather need a summarizing view onto the single pieces of knowledge.

With these aims in mind, we have developed CHAR, the *Corporate History Analyzer*, a system that allows for collective gathering of information and for intelligent access to vast amounts of small pieces of knowledge at the appropriate level of abstraction. Thus, CHAR constitutes an new type of system that is interesting from the application as well as from the technical point of view. Seen from the side of the users, i.e. the corporate analysts and their customers, it surpasses existing systems regarding the range of support it offers on the querying and on the provisioning side. Regarding the technical background, it incorporates a range of ontology-based knowledge management tools that are interesting in themselves and which are combined to support a multitude of views based on an ontology with an elaborate temporal structure.

In the rest of this paper we will first elaborate the business case that the history analyzer may be used for (cf. Section 2). Then we describe the views onto the knowledge base that support corporate researchers and their co-workers and that allow the former (or possibly both groups) to contribute to the knowledge base. In a final step we describe how the overall system has been conceived and implemented (Section 4). Before we conclude, we give a short survey of related work and refer to some business cases that are based on the same premises as our study, but which work on a basis much less sophisticated than CHAR.

2 A Case Study — Restructuring in the Chemical Industry

Consider the case you are key account manager of ColorBasic Inc, a supplier of basic material for the color industry, responsible for several key accounts in the European color Industry. The other day you read in a newspaper article, that one of your biggest customers, the German Bigcolor AG is in acquisition talks with their french competitor Grandcouleur S.A.

So what? In a first happy reaction about this development, you discuss with your production manager the ability of your company to increase the production of basic material in a short time period. You plan to convince your key account at Bigcolor AG to enhance your supply agreement also for their new subsidiary immediately after aquisition. Nice analysis! Your production manager confirms a possible increase of production and prepares a quick switch of production line. Two month later, you read the news about the acquisition of Grandcouleur S.A. and immediately, you meet your key account at Bigcolor AG to get his agreement for higher sales volume of basic material with very flexible delivery conditions. Nice strategy!

Your key account is impressed about your quick reaction and your fast and flexible supply conditions, but tells you, that Bigcolor AG acquired Grandcouleur S.A. due to the fact that they have a subsidiary in the basic material industry, GrandBasic Inc., your biggest competitor. Therefore, Bigcolor AG made a vertical integration and bought very interesting supply conditions via the acquistion of Grandcouleur S.A. Thus, unfortunately, instead of higher sales volume you wont have any sales in near future with Bigcolor AG. Bad News!

Now, it is easy to recognize, that your analysis of the situation was wrong. You had not enough information about the acquisition to create a better strategy. Later on, you will hear a lot of colleagues telling you stories about the well-known relationship between Grandcouleur S.A. and GrandBasic Inc.

Even worse, but how should you handle such situations? There is an ever faster pace of change in global markets. There is no day without news about mergers, acquisitions and a lot of other strategic moves of players in every industry. One of the most dynamic markets in the past years has been the chemical industry. Here, big global players permanently restructure their organization, try to acquire small innovative technology leaders or fight about entry strategies in new markets. This very dynamic situation leads to the problem, that no market participant has a permanent clear picture about the industry.

Nowadays, the market analysts, consultants and inhouse market research departments try to track the activities of their industry with traditional methods. Newspaper articles, online databases and annual company reports were analysed and competitors web pages were thoroughly tracked. The results are presented to the management, published in reports and distributed. There exist several problems in this research process:

- Information archives are document-based. For a collective gathering of facts this view is too coarse. By this way one may introduce an abundance of redundancies on the one hand and one may account for only few of many aspects that occur in a particular document (e.g. only the information given in the headlines), thus missing the crucial pieces of information again.
- The next problem stems from the fact that typically document management systems rely almost exclusively on information retrieval techniques, which are rather unreliable.
- Even if convential annotation is used for these documents, it is most often too cumbersome to describe the relevant implications, too. E.g., if one company sells one of its units, the implication is that this unit may no longer be the sole supplier for the overall company as it was before.
- These implications may only be made transparent if background knowledge is used, but document management systems are typically restricted to thesauri — the most advanced ones employ topic maps.
- Finally, different executives need different views onto the same basic pieces of information. What is relevant to the one employee, e.g., a competitor's expansion in his regional market, may be irrelevant from a more global point of view.

Thus, what we wanted to have was a system that, *(i)*, supports the collective gathering of information on the level of facts rather than documents, *(ii)*, integrates the gathering task smoothly into the common research process, *(iii)*, allows to intelligently combine facts or give a high-level view, e.g., a summarization or a statement about a general trend, about a whole set of facts, *(iv)*, checks new facts against the available background knowledge, *(v)*, allows a multiple view access to the knowledge via a single entry portal, and, *(vi)*, allows to route derived facts back into the common workplace environment.

Additionally, in our specific business process, we want our system to improve the quality and speed of the knowledge gathering and distribution process in a companys research department. Therefore we have to assure an intuitive graphical user interface, tailored to the specific user needs.

Being given this system, the key account managers of ColorBasic Inc. could have really profited from the work done in his company's research department. He would have seen the strategic implications following from the acquisition of Grandcouleur S.A. by Bigcolor AG. The system would have used the fact that GrandBasic Inc. is a subsidiary of Grandcouleur S.A. and, thus, a probable supplier for Bigcolor AG in order to derive a vertical integration as a possible objective for the acquisition. Hence, the key account manager could have done the right things.

3 CHAR — The system

In the previous section we have seen some pitfalls that may appear while tracking one's markets and competitors. From these problems we have derived the requirements that a system for corporate history tracking should meet:

1. It should provide multiple views onto the same knowledge base for different time frames, for different regional foci, for varying intra-organisational structures and for different strategic questions, to name but a few.

2. It should allow many people to contribute factual knowledge in a way that is embedded into their common work process and that is organized around a semantic background.

In this section we elaborate on these two requirements in Subsections 3.1 and 3.2. Thereby, we only summarize the second part as it is composed of parts that we have already described elsewhere ([7, 12, 13]).

The reader may note here that the mechanisms described here rely on a common conceptualization of the domain, an ontology. This and some further technical assumptions will be explained in more detail in Section 4.

3.1 Querying Knowledge

The query interface of the Corporate History Analyzer has been developed in order to deal with organizational and strategic questions that depend on spatio-temporal constraints. It renders views that may be seen on a common web browser. Actually, they look just like common web sites with some choices for selection. These selections are also controlled by the knowledge of the backend system, e.g. one can only select companies that are known to exist in the knowledge repository.

Figure 1 depicts the main views that are offered by CHAR, *viz.* Activities, Organization, Know How, Strategic Questions and General Query Possibilities (indicated by "Search"). While the latter may be asked by means such as known from standard tools for OLAP/business intelligence, we consider to help the user with the corporate history analyzer through views that directly reflect his interests on the organizational and strategic level.

3.1.1 Answer organizational questions

The first major category of queries relevant to the corporate history is about organizational structures and the activities that change organizational structures. For instance, the view of "acquisitions of M.A. Hanna" returns all its purchases (cf. Fig. 2), and corresponding views are offered for Sellings, Mergers, Restructurings and Management Changes (Figure 1).

What is interesting to note at this point is that it is rather difficult to get a clear picture of what is really happening with M.A. Hanna. It is difficult and time-consuming for the human analyst to detect some trend in lists of single facts. Observations become much easier, when different types of facts may be related and contrasted. For instance, Figure 3 depicts two snapshots of M.A. Hanna's organigram that are automatically derived from single activities, like acquisitions and restructurings, and that give the analyst a neat picture of how formerly isolated purchases that M.A. Hanna made before 1994 were more tightly integrated in the company in 1997 (e.g., "Compounding Technology" having been reorganized into the Business Area "Plastic Compounding").

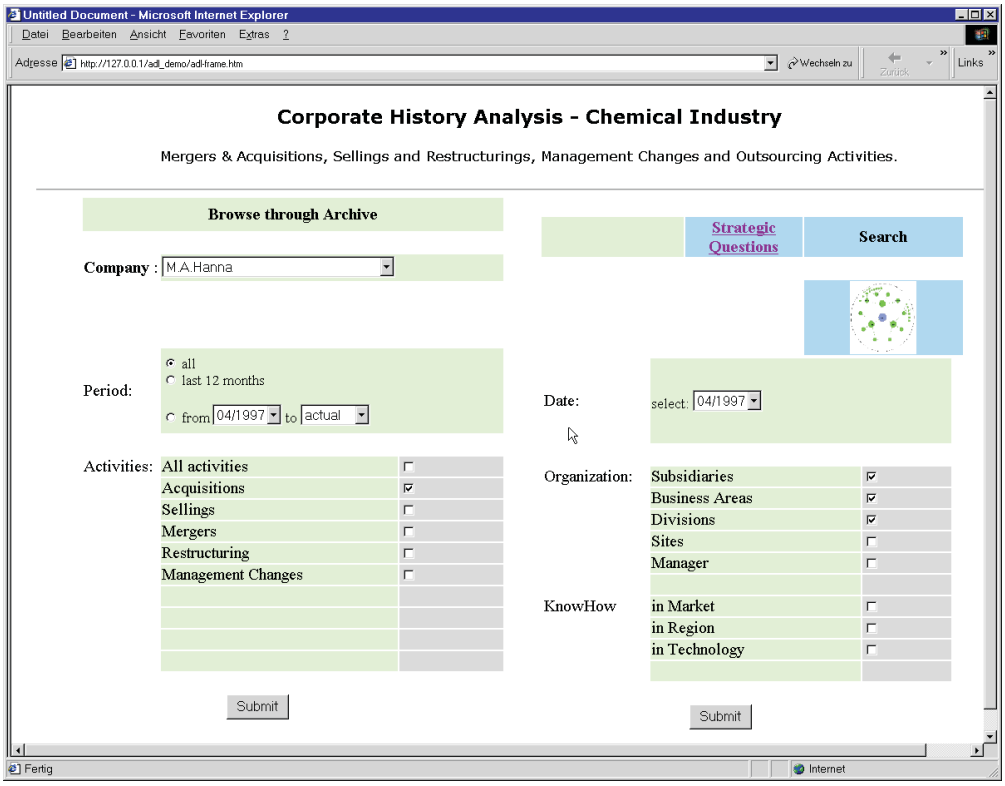

Figure 1: The entry point to CHAR

It should have become clear that this type of comprehensive support is beyond the scope of straightforward database solutions, but rather requires extensive programming or the type of mechanisms that we rely on for CHAR — described in detail in Section 4.

3.1.2 Answer strategic questions

This far, we have described how CHAR provides sophisticated support that is based on “hard facts” and that does not aim at *interpretation* of single facts. With its support for strategic questions, CHAR tries to give answers to questions about business competitors which cannot be answered definitely, but which rely on some conjectures. For instance, the purchase of a company from abroad may lead to a gain of market share in that area, and thus to a regional diversification. Thus, the first question in in Figure 4 is a prototype strategic question. The overall list that we currently support reads as:

1. Which Activity of X leads to operation in region Y?
2. Is there any regional expansion of X due to any acquisition since T?
3. Is there any regional expansion of X due to the acquisition of Y?
4. What activity of X modified their Business Area Y?
5. What experience gains X from acquisition of Y?

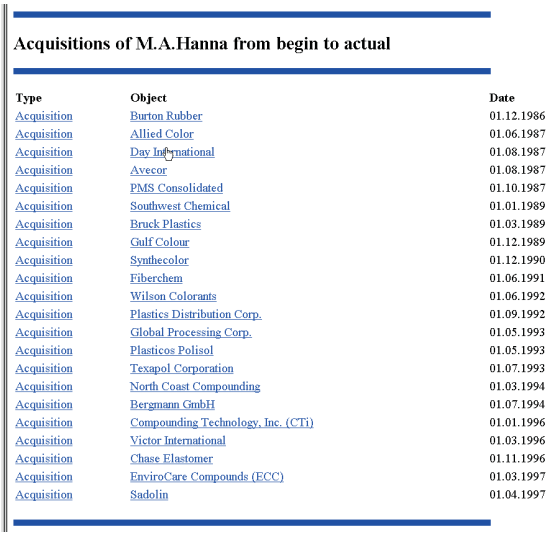

Acquisitions of M.A.Hanna from begin to actual

Type	Object	Date
Acquisition	Burton Rubber	01.12.1986
Acquisition	Allied Color	01.06.1987
Acquisition	Day International	01.08.1987
Acquisition	Avecor	01.08.1987
Acquisition	PMS Consolidated	01.10.1987
Acquisition	Southwest Chemical	01.01.1989
Acquisition	Bruck Plastics	01.03.1989
Acquisition	Gulf Colour	01.12.1989
Acquisition	Synthecolor	01.12.1990
Acquisition	Fiberchem	01.06.1991
Acquisition	Wilson Colorants	01.06.1992
Acquisition	Plastics Distribution Corp.	01.09.1992
Acquisition	Global Processing Corp.	01.05.1993
Acquisition	Plasticos Polisol	01.05.1993
Acquisition	Texapol Corporation	01.07.1993
Acquisition	North Coast Compounding	01.03.1994
Acquisition	Bergmann GmbH	01.07.1994
Acquisition	Compounding Technology, Inc. (CTi)	01.01.1996
Acquisition	Victor International	01.03.1996
Acquisition	Chase Elastomer	01.11.1996
Acquisition	EnviroCare Compounds (ECC)	01.03.1997
Acquisition	Sadolin	01.04.1997

Figure 2: Acquisitions of M.A. Hanna

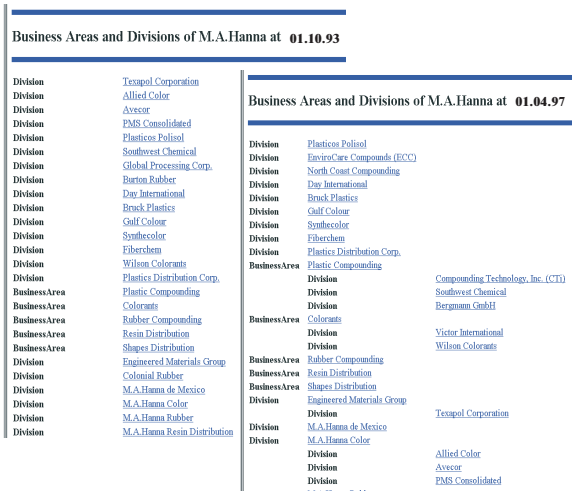

Business Areas and Divisions of M.A.Hanna at 01.10.93

Division	Texapol Corporation
Division	Allied Color
Division	Avecor
Division	PMS Consolidated
Division	Plasticos Polisol
Division	Southwest Chemical
Division	Global Processing Corp.
Division	Burton Rubber
Division	Day International
Division	Bruck Plastics
Division	Gulf Colour
Division	Synthecolor
Division	Fiberchem
Division	Wilson Colorants
Division	Plastics Distribution Corp.
BusinessArea	Plastic Compounding
BusinessArea	Colorants
BusinessArea	Rubber Compounding
BusinessArea	Resin Distribution
BusinessArea	Shapes Distribution
Division	Engineered Materials Group
Division	Colonial Rubber
Division	M.A.Hanna de Mexico
Division	M.A.Hanna Color
Division	M.A.Hanna Rubber
Division	M.A.Hanna Resin Distribution

Business Areas and Divisions of M.A.Hanna at 01.04.97

Division	Plasticos Polisol	
Division	EnviroCare Compounds (ECC)	
Division	North Coast Compounding	
Division	Day International	
Division	Bruck Plastics	
Division	Gulf Colour	
Division	Synthecolor	
Division	Fiberchem	
Division	Plastics Distribution Corp.	
BusinessArea	Plastic Compounding	
	Division	Compounding Technology, Inc. (CTi)
	Division	Southwest Chemical
	Division	Bergmann GmbH
BusinessArea	Colorants	
	Division	Victor International
	Division	Wilson Colorants
BusinessArea	Rubber Compounding	
BusinessArea	Resin Distribution	
BusinessArea	Shapes Distribution	
Division	Engineered Materials Group	
	Division	Texapol Corporation
Division	M.A.Hanna de Mexico	
Division	M.A.Hanna Color	
	Division	Allied Color
	Division	Avecor
	Division	PMS Consolidated

Figure 3: Before and after restructurings

Answers for some of these questions are presented in the lower part of Figure 4. For example, Hanna's acquisition of Bergmann in 1994 led to a new operation in Western Europe — exploiting the background knowledge that Bergmann's German base lies in Western Europe (lower left of Figure 4). Question 2 triggers a search for any kind of regional expansion within some time frame, which may be answered positively with the case of EnviroCare Compounds (middle), Finally, Hanna's business area of Plastic Compounding was affected by restructurings such as the one of the company part that formerly was Bergmann.

By now it should be obvious that for this type of questions one needs comprehensive reasoning support, such as rules about geographic relationships (Japan is in Asia, Italy is in Southern Europe, etc.), rules about categorization of know how, as well as rules about strategic indications (purchase may be indicative for regional expansion).

3.2 Providing Knowledge

The process of providing new facts into the knowledge warehouse should be as easy and as smoothly integrated into the common working tasks as possible. For this reason we offer various modes of contributing knowledge.

First, one may enter information through a *form-based interface* that is extended by a canonical and a hyperbolic tree view onto the concept hierarchy (cf. [13]).

Second, when the information that is to be provided is produced during the writing of documentations or reports one may use a template approach in order to generate knowledge by writing these papers (cf. [12]).

Third, one may use wrapper mechanisms in order to provide data from tables and lists on the web. For instance, the approach described in [2] uses the same basic representation (viz. F-Logic, cf. next section) as its semantic backbone and may, thus, directly feed into our knowledge base.

Fourth — and most important for CHAR, one may use our annotation tool in order to add metadata to data given in documents (cf. [7]). A snapshot of the annotation tool noting some action about M.A. Hanna is shown in Figure 5: The user reads or works with documents using a text or spreadsheet processing tool or an internet browser. When he detects some relevant change being described in

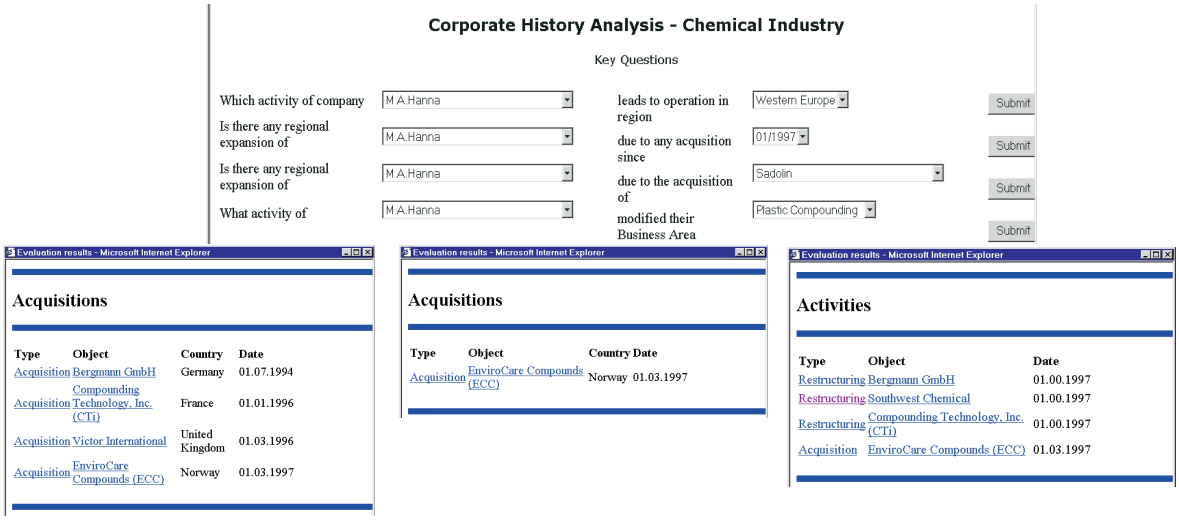

Figure 4: A view onto strategic questions

the document and this change might become relevant[1], he highlights the word or phrase and uses the annotation tool to select the type of the phrase (e.g. “Hanna M.A.” is a company) and its relation to other material (e.g. Hanna sells Shapes Distribution Business to GE Plastics on May 11, 2000). The document, these facts and metadata about the annotater, the time of annotation, etc. is all stored in the back-end database. Currently this is a purely manual process. In the future we plan to use information retrieval methods to recognize phrases (or words) that had been seen and annotated before (e.g., if there is a second document about “Hanna M.A.”) and information extraction methods that propose types and relations between objects for annotation.

4 Developing CHAR

Starting from the principle problem (Section 1), from scenarios such as seen in Section 2, and from use cases such as described in the previous section, the question was how to bring the required conceptual structures and reasoning capabilities into action. To answer this question, we here describe how we have applied our general methodology for developing ontology-based systems to corporate history analysis (Subsection 4.1) including the ontology engineering part (Subsection 4.2) as well as the interface and integration issues (Subsection 4.3).

4.1 The Development Process

We distinguish different phases in the development process that are illustrated in Figure 6. For the main part this model is a sequential one. Nevertheless, at each stage there is an evaluation as to whether and as to how easily further development may proceed with the design decisions that have been accomplished before. The results feed back into the results of earlier stages of development. The main stages of the development process and their key characteristics are given in the following:

[1]Our first intuition was that every major organizational change of one’s competitors or clients should be noted by the company’s IT research department. Over time this rough estimation will have to be specified more precisely — also depending on the type of change and its financial implication.

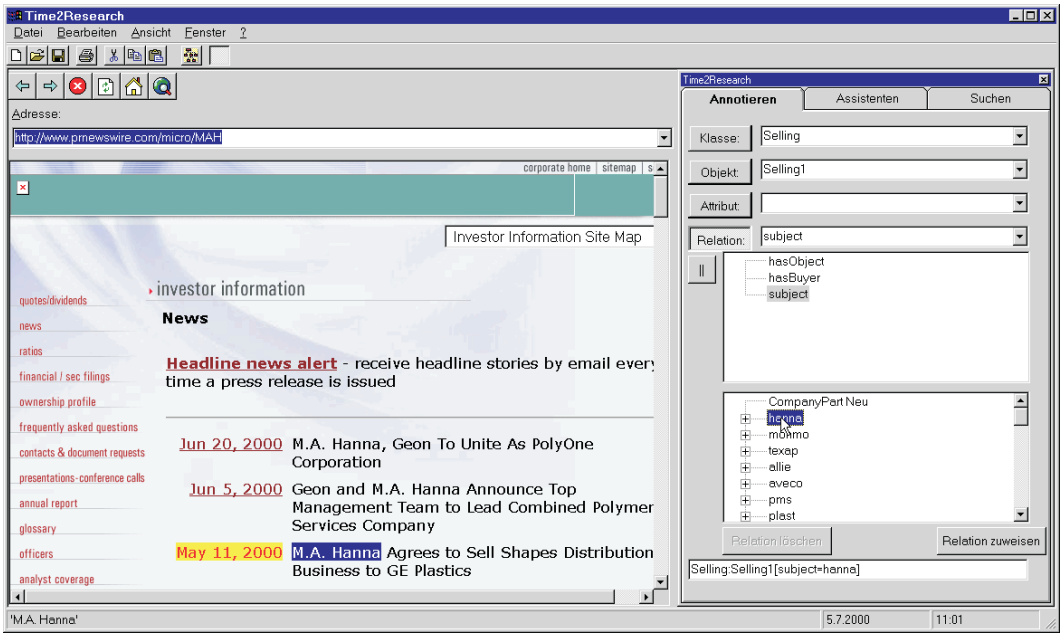

Figure 5: The process of annotating documents

- User requirements are collected in the requirements elicitation phase. Results of this very first phase constitute the input for the design of the web site and for preliminary HTML pages and influence the formal domain model embodied in the ontology.

- In parallel to the development of the structure and layout of the web site an ontology engineering process is started. An ontology defines a common language and schema used to structure, find and exchange knowledge about the domain. For these purposes interesting concepts together with attributes and relations between concepts are identified and represented. Thereafter, rules are developed. Rules describe more abstract dynamic relations between objects which are used to derive new objects, attributes or relations that are not given explicitely. Developing the ontology containing static parts (concepts, attributes, roles) and dynamic parts (rules) is the central issue in developing such an application and is described for our case at hand in more detail in Section 4.2.

- In the query development step the views and queries described in one of the earlier phases are formalized. At first, their functionality is tested independently from the web site design in evaluating those formalized queries using the ontology and test facts.

- Finally, web pages are populated, *i.e.* the queries and views developed during website design, and formalized and tested during query formalization are integrated into query interface.

- After this core development processes information is provided to the knowledge base by the tools mentioned in Section 3.2, *i.e.* by annotating documents, by directly providing facts using the fact editor etc. This process also runs in parallel to the use of the system.

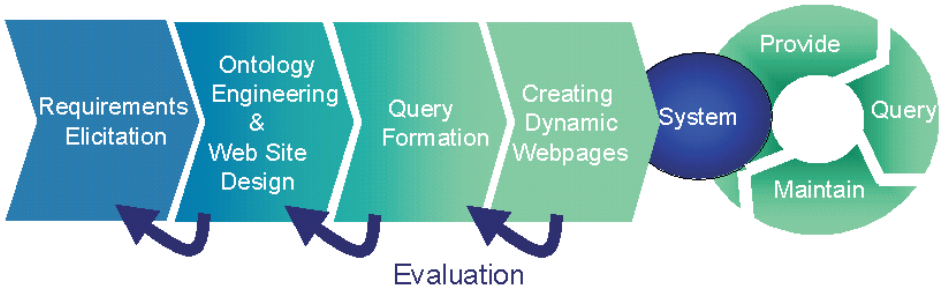

Figure 6: Developing ontology and web-based systems

4.2 Ontology engineering

The ontology for this application consists of (i) concepts defining and structuring important terms, (ii) their attributes specifying properties and relations, and (iii) rules allowing inferences and the generation of new knowledge. The language we use to represent this ontology is F-Logic [9], which provides adequate modeling primitives integrated into a logical framework. We use our tool OntoEdit [14] to edit this static part of the ontology graphically.

One speciality in this domain is that time issues have to be modeled. There are many objects which have a life time. For instance a company has a start time when it is founded and there is an end time when this company is acquired by another company or is merged with another company. So durations had to be modeled. Likewise, an `Acquisition` is an `Event` which occurs at a specific time, which leads to the modeling of `Timepoints`. The following example shows the modelling of `Acquisitions` as `Activitys` which itself are events in F-Logic. The attribute `occursAt` represents the time point. `Activity` has the role `subject` which represents the company executing the activity. `Acquisition` has the properties `price` and `share` and the role `hasObject`, which represents the `CompanyPart` the activity applies to.

```
Event [occursAt => Timepoint].
Activity::Event [subject => CompanyPart].
Acquisition::Activity [
    hasObject =>> CompanyPart;
    price =>> NUMBER;
    share =>> NUMBER ].
```

One of the key characteristics of this application was that the user provides facts about new actions like acquisitions and mergers to the systems and the system derives the consequences. For this purpose rules had to be modelled for all possible activities which incur such consequences, for instance *"If a company is renamed, it remains the same company but from this timepoint on this company has the new name"*, *"If two companies are merged, a new company with a new name is created, the old companies disappear, they are from now on subsidiaries of the new company"*, *"If a company is acquired by another company with share 100% then the first one disappears at this time point and then belongs to the second one"*, or *"If a division is outsourced it becomes part of a new company or it becomes an company on its own"*.

All these examples show that the modelling of time is a central issue in this application. Modelling of time has non-monotonic effects. For instance when a company is founded it potentially lives forever. But if it is later acquired by another company its life time ends.

In addition to the modelling of the effect of actions possible strategical consequences had to be modelled by rules. For example, *"A company which has no experience in a certain technology and which acquires another company which has invented this technology may have the strategy to become technology leader in this area"* or *A company which has no base in a certain region and which acquires another company which operates in this region may have the strategy to expand into this region"*.

4.3 Glueing all parts together

During website design the presentation strategy and actual layouts are developed. In the query formation phase, queries are developed which deliver the information. Instead of expressing those queries directly in F-Logic, it is more convenient to use our Query Builder tool. In addition to the graphical definition of queries it provides output templates for the answers. These templates may then be used as starting points for the development of the final presentation of the answers to the user. There are also templates for selection lists, radio boxes or check lists, which may be used as initial settings of HTML form fields.

In the case that a user has to click to a fixed query without providing further information such a query may be integrated directly into the web page as a hypertext link by copy and paste from the Query Builder tool into the HTML editor.

In the other case that a user has to fill out a form which then generates a query, this query has to be created out of the contents of the form fields by a Javascript program. E.g. the query that returns the names N of all companies C at time t (t is between start and end time of the company) may be expressed in F-Logic by:

```
FORALL N,C, ST,ET <- C:company[name->>N] AND starts(C,ST) AND
ends(C,ET) AND ST <= t < ET.
```

The user has to provide time t in a form field. The Javascript program has to read the time t out of the form field, has to generate the query as a string, has to add the url of the web page the answers are fed into and send it to the Ontobroker system.

5 Related Work

In the business community various approaches are used to analyze the development of markets and the bahviour of competitors. Portfolio analysis [10] is one of the most popular methods that are used. CHAR is well-suited to support portfolio analysis since this method depends on an up-to-date and comprehensive collection of information about the behaviour of competitors, a process that is considerably improved by CHAR. Furthermore, the rules that are available as part of the background knowledge assist in deriving (strategic) conclusions about the behaviour of these competitors.

Corporate research itself is also an important business area. A large collection of companies offer such research services. Typically, the provided information is stored in databases that can be accessed via different indices. E.g. The Dialog Corporation offers such databases addressing among others mergers and acquisitions [6]. However, these databases do not provide any linking between the different M&A actions nor do they derive additional information from these facts. Reuters Business Briefing [1] is offering a broad collection of services including also historical information about companies. These services can be integrated into an intranet environment thus providing access to these services at the workplace of the knowledge worker. Nevertheless, these services do not offer derived

knowledge that is based on the explicitly specified facts. In the same way Hopenstedt [8] is offering a large collection of company information either in online databases or as CDs. Compared to these commercial services the CHAR system is unique with respect to its integration and interlinking of information and the inference capabilities for deriving new knowledge from the given facts.

Considering the technical part, there exist, e.g., numerous publications about handling temporal information. For instance, the Later system for manipulating and querying a temporal knowledge base is described by Brusoni et al. [4]. Later offers operations to assert and delete facts that specify qualitative and quantitative temporal relations. However, temporal reasoning that includes the handling of non-monotonicity is not supported by Later.

In [3] a general outline of an ontology-based approach to knowlege management is given. There, three subtasks are described: ontological engineering, characterizing the information sources in terms of the ontology, and intelligent knowledge retrieval. The CHAR system shows how this general approach may be instantiated to an intelligent KM solution and how these three subtasks are carried out by powerful tools.

An intelligent search facility for documents is described in [15]. They use the so-called Q-Technology for classifying documents and offering intelligent retrieval mechanisms. However, the Q-calculus used in this approach is far less powerful than Frame-logic as used within the CHAR system. Furthermore, the CHAR system delivers knowledge pieces as answers to posed queries, whereas the system outlined in [15] delivers documents as query results.

6 Conclusion

Nowadays, markets are changing faster and faster. Therefore, the need to analyze the changes that occur in a market increases dramatically. Conventional, document-based approaches do not meet the needs anymore: different pieces of information have to be linked to each other and related to the background knowledge about general market behaviour.

CHAR (Corporate History Analyzer) is a knowledge managment system supporting the research process within a company. It is an intelligent assistant-type of system for keeping track of mergers and acquisitions in a specific market and for providing advice about strategic implications that can be drawn from these M&A activities. More specifically, CHAR provides means for *(i)* a decentralized information provisioning process - by using a template approach and offering an annotation tool, *(ii)* collecting information from these different sources and integrating them in the knowledge base, *(iii)* relating these pieces of information to the available background knowledge, and *(iv)* deriving consequences that follow from these acquired facts and the background knowledge - by applying (non-monotonic) reasoning techniques. Based on these techniques CHAR is offering integrated results to various kind of queries including summarizations or trend analyses. These functions are built on top of the collection of methods and tools as offered by the Ontobroker system [5] and the Semantic Community Web Portal framework [13].

Future work will aim at providing better support for extracting relevant pieces of information from various sources: a combination of linguistic and ontology-based methods will provide semi-automatic means for information extraction. A first step towards this direction is described in [7]. Another line of further research will be the integration of the CHAR system into the Smart Task Support approach [12] thus providing assistance for researchers at their workplaces.

References

[1] Dow Jones Reuters Business Interactive LLC. Trading as Factiva. Factiva, a Dow Jones Reuters company. http://www.factiva.com/rbb/companies.asp?node=rbb-link4, Observed at July 3, 2000.

[2] Mathias Bauer and Dietmar Dengler. Trias: Trainable information assistants for cooperative problem solving. In *Proceedings of Agents-1999*, pages 260–267. ACM Press, 1999.

[3] V. R. Benjamins, D. Fensel, and A. Gomez. Knowledge management through ontologies. In Reimer [11].

[4] V. Brusoni, L. Console, P. Terenziani, and B. Pernici. Later: Managing temporal information efficiently. *IEEE Expert*, pages 56–63, July/August 1997.

[5] S. Decker, M. Erdmann, D. Fensel, and R. Studer. Ontobroker: Ontology Based Access to Distributed and Semi-Structured Information. In R. Meersman et al., editors, *Database Semantics: Semantic Issues in Multimedia Systems*, pages 351–369. Kluwer Academic Publisher, 1999.

[6] DIALOG. Smart tools for company intelligence from the dialog corporation. http://library.dialog.com/smart_tools/compint/index.html, Observed at July 3, 2000.

[7] Michael Erdmann, Alexander Maedche, Hans-Peter Schnurr, and Steffen Staab. From manual to semi-automatic semantic annotation: About ontology-based text annotation tools. In P. Buitelaar and K. Hasida, editors, *Proceedings of the COLING 2000 Workshop on Semantic Annotation and Intelligent Content. Luxembourg, August 5-6, 2000*, 2000.

[8] Hoppenstedt. Hoppenstedt Verlagsprogramm Firmendatenbank. http://www.hoppenstedt.de/fi_hfdb1.asp, Observed at July 3, 2000.

[9] Michael Kifer, Georg Lausen, and James Wu. Logical foundations of object-oriented and frame-based languages. *Journal of the ACM*, 42, 1995.

[10] M. E. Porter. *Wettbewerbsstrategie (Competitive Strategy)*. Campus Verlag, Frankfurt, 1992.

[11] U. Reimer, editor. *PAKM 98 Practical Aspects of Knowledge Management — Proceedings of the Second International Conference*, 1998.

[12] H.-P. Schnurr and S. Staab. A proactive inferencing agent for desk support. In *Proceedings of the AAAI Symposium on Bringing Knowledge to Business Processes*, Stanford, CA, USA, 2000. AAAI Technical Report, Menlo Park.

[13] S. Staab, J. Angele, S. Decker, M. Erdmann, A. Hotho, A. Maedche, R. Studer, and Y. Sure. Semantic Community Web Portals. In *Proceedings of the 9th World Wide Web Conference (WWW-9), Amsterdam, Netherlands*, 2000.

[14] S. Staab and M. Maedche. Axioms are objects, too — Ontology engineering beyond the modeling of concepts and relations. In *Proceedings of the ECAI 2000 Workshop on Ontologies and Problem-solving Methods. Berlin, August 2000*, 2000.

[15] K. Stanoevska-Slabeva, A. Hombrecher, S. Handschuh, and B. Schmid. Efficient information retrieval: Tools for knowledge management. In Reimer [11].