Towards Explanations for CBR-based Applications

Marco Spinelli, Martin Schaaf
University of Hildesheim,
Department for Mathematic and Applied Computer Science
Data & Knowledge Management Group
P.O.Box 101363 D-31113 Hildesheim, Germany
{spinelli, schaaf}@dwm.uni-hildesheim.de

Abstract

Case-based Reasoning (CBR) is a mature technology for building knowledge-based systems. Unlike with reasoning approaches making use of deductive inference, CBR-based applications are capable to produce useful results even if no answer matches the query exactly. Result sets presented to users are ordered by means of similarity and utility. However, for knowledge intensive domains we discovered that results sets enriched by calculated similarity values for particular answers are not sufficient for experts. Such users have a demand for additional information and explanations making the proposed results more transparent. By presenting additional explanations to them, their confidence in the result set increases and possible deficiencies, e.g. in the weight model, can be revealed and corrected. In this position paper we investigate explanation approaches for CBR from the user level perspective. Besides identifying potential uses cases, we sketch techniques for creating different kinds of explanations and relate them to already existing approaches from other areas of CBR research, e.g. conversational CBR.

1 Introduction

CBR-based applications can become highly complex with regards to the knowledge stored in the various knowledge containers. It is gained during a work intensive and difficult elicitation process, usually hidden from the user and compacted into a single similarity value used for ranking the retrieval results [Bergmann 2002]. For highly knowledge intensive tasks executed by domain experts, we observed a growing demand for additional explanations that make the retrieval process and the proposed results more transparent. Such explanations cannot only reveal the knowledge behind the CBR application but also provide additional assistance in interpreting the result set. The reasons are manifold: Weights encoded in the similarity model may not reflect the users preferences or the underlying case base may not have been consolidated.

In this paper, we discuss our approach of providing the user with additional explanations to result sets of CBR applications. With this approach, users can have a more detailed view on the results and decide, which is best suited for their purposes.

In the following section we will briefly introduce the prototypical CBR application where we encountered the demand for explanations. We will then analyze the problem and possible solutions in more detail. Before we conclude our discussion, we will provide an example that illustrates the current state of work.

2 About the Context of this Work

The problems tackled in this paper occur within the context of the IPQ[1] Project (Intellectual Property Qualification) that aims at supporting developers of microelectronic circuits in their search for design components to be reused. For such components the term *Intellectual Property* (IP) [Lewis, 1997] has been assigned within the *Electronic Design Automation* community. The application of CBR to the selection of IPs has been published for example in [Schaaf *et al*, 2002a] and [Schaaf *et al*, 2002b]. Within the scope of this paper, it is sufficient to mention that IPs are characterized by a set of formal attributes typically consisting of a value from a well-defined type and an associated metric. The set of types includes, beside primitive types like real or integer values, taxonomies and intervals. For the retrieval of IPs the user specifies a query by providing a subset of attributes reflecting his/her current design situation. According to these attributes the system tries to find suitable IPs by assessing the similarity between query and IPs stored as cases in a case-base. The set of retrieved IPs is used as input for subsequent steps towards the final decision about the IP to be integrated into a microelectronic design.

Our CBR application for IP Retrieval is based on the structural CBR approach [Bergmann 2002] that makes use of a default similarity model containing local similarity measures for each IP attribute type of the characterization as well as global similarity measures facilitated by aggregation functions for higher-level categories.

Furthermore, there exists an appropriate weight model reflecting the importance of particular IP attributes and

[1] IPQ Project (12/2000-11/2003). Partners: AMD, Fraunhofer Institute for Integrated Circuits, FZIKarlsruhe, InfineonTechnologies, Siemens, Sciworx, Empolis, ThomsonMultiMedia, TU Chemniz, University of Hildesheim, University of Kaiserslautern, and University of Padeborn. See http://www.ip-qualifikation.de/

categories as well. Finally, it is assumed that the local similarity measures are fixed, while the global similarity calculation can be adjusted by user specific weights. These weights will be stored together with other information about the IP user in a profile that builds the context of a particular IP Retrieval query.

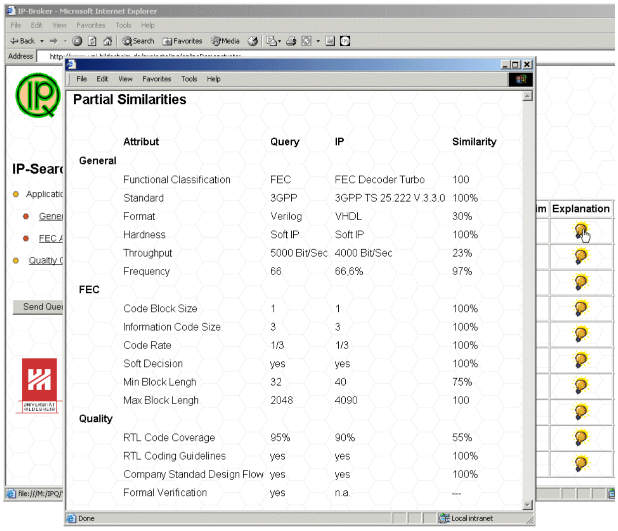

Figure 2: User Interface for IP Selection

Figure 2 shows the interface presented to users of the IP Retrieval service. To be more specific, it illustrates the simplest type of explanation possible where the particular calculated similarity values are shown in an additional window. This kind of explanation support was implemented by the READee prototype [Oehler *et al,* 1999], a predecessor of the CBR-based application developed for our project. Of course, the examination of the intermediate results is a cumbersome task and we will discuss more sophisticated approaches later in the paper. The IP Retrieval system is available online at the University of Hildesheim and will be publicly accessible if we receive the necessary approval from our project partners. The ideas discussed here constitute the necessary preliminary work for future extensions of the prototypical system.

3 Use Cases for Explanations within the IPQ Context

Within the context of the IPQ Project we identified a variety of use cases together with our industrial partners. Without going into details of the IP qualification process, we can distinguish between three different use cases on the conceptual level.

3.1 Report Generation

Generating a report for each proposed IP with respect to the user query is a fundamental demand for explanation support. It includes the determination of the most relevant attributes leading to the proposed result by calculating the absolute relevance of each attribute. Furthermore, users can select short text explanations for each attribute involved in the similarity calculation to be included into the report. Such explanation may be *design rationales* of the domain expert or *best practices* associated to attributes specified by other users.

While this kind of explanation support mostly operates on the IP characterization and the similarity model, it also includes text patterns connected to concepts defined in the vocabulary of the CBR application.

3.2 Increasing User Confidence

As mentioned, the selection of IP is a highly knowledge intensive task. Before an IP is considered as candidate for the subsequent entry check, a task that consumes significant time and money, the designer needs more confidence in the results proposed by the IP retrieval system. This is especially true if the coverage of attributes specified in the query and in the case is low. Explanations providing appropriate visualizations of the result set, e.g. by determining and graphically rendering the attributes with the highest impact on the assessment, aid the designer in getting a quick overview.

3.3 Determine Deficiencies

Another purpose of explanations is the identification of deficiencies in the result set. Again, a low coverage of attributes in the query and the cases is the starting point for appropriate analysis techniques that can lead to the following:

- Reengineering of the default similarity model by the knowledge engineer.
- Supporting the IP User in specification of his/her user specific adjustments that have impact on the retrieval assessment.

The distinction between the purpose of increasing the users confidence and determining deficiencies results is justified because it leads to contrasting strategies for the explanation support. An optimistic strategy basically assumes a high precision and recall, and thereby a high utility of the proposed result set. A pessimistic strategy assumes to result to be poorly suited. The decision, which strategy to be applied, depends on the quality of the case-base, the maturity level of the similarity model, and the particular query. Consequently, a part of this work is to identify criteria that enable the intelligent selection of analysis approaches from a toolbox.

4 An Overview of Explanation Approaches

When deploying a system based on structural CBR within a knowledge intensive domain, a frequent task for users is to interpret the proposed results according to the query. This is true because of the complex knowledge model that lies behind a structural CBR application. Furthermore, it is often difficult to make any assumptions about the quality of the case base itself.

4.1 Simple "Colored Explanations"

A first step towards explanations is to give an overview, which attributes specified by the user, have the highest impact on the similarity and to separate them from the attributes with lower impact. The realization of such an explanation support is not difficult with the CBR retrieval engine orenge from empolis [empolis 2001] that allows dynamic and conditional rendering of the result set. Accompanied by additional text patterns facilitating natural language explanations even such an easy a straightforward approach is extremely useful for generating reports.

4.2 Explanations as Experiences from the Past

A more sophisticated technique for explanations is to make use of a second CBR system where known explanations to queries and result sets are stored and retrieved in parallel. Beside the knowledge already present behind the original application, such an "embedded" explanation component stores additional "explanation cases". The downside of this approach is the additional maintenance effort caused by the second CBR-based application. However, both are closely related, which makes multi-level case base approaches feasible, as they have been proposed for instance in [Nick *et al*, 2002]. Although the method presented there focuses on the integration of cases at different levels of maturity, it is also a good starting point for decreasing the maintenance effort for integrated IP and explanation case bases.

4.3 "Data Mining" in Result Sets

A third alternative to generate explanations encompasses various techniques for either knowledge-based or statistical data analysis in order to detect interesting relationships or other regularities. For IP Retrieval, we are currently investigating a technique that utilizes CBR, again. Here, we take the best case of the result set as prototype for a new query and compare the original result set and the one corresponding to the new query. A difference exceeding a certain threshold indicates potential deficiencies of the query or the CBR application itself that should be reported to the user. For instance, let us imagine a user specifying a value for a particular attribute that greatly differs from the average value assigned to the cases of the case base. Under such circumstances, specification of the attribute does not contribute anything to the similarity assessment because all cases have nearly the same low similarity with respect to this attribute. Using the best case for a new query could reveal this situation and probably end up with a totally different result set, as we will illustrate in the next section when presenting an example in more detail. The determination of attributes that with high impact on the overall similarity assessment is also tackled by current research for conversational CBR methods [Schmitt, 2003]. Here important attributes are determined step by step during a dialog process with the user depending on the customer's previous answers and remaining potential cases of the case base.
Another approach that is statistic based is a graphical visualization of the result set. The method to visualize case bases for easier management from [McArdle and Wilson 2003] can, of course, also be adapted to visualize only the retrieved cases. The underling MDS[2]-Algorithm transforms mutli-dimensional data points in a two-dimensional space. The generated distance between two points in 2D is preserved from the former n-dimensions. The user obtains a better survey how similar the retrieved cases are to the query and to each other. This provides more insights into the similarity assessment than the usual single dimensioned similarity attribute.

5 Example Application Scenario for Explanations

As mentioned in the previous section, a promising approach to explanation is to apply various techniques analyzing the result set and detect interesting regularities or relationships. In the following we will elaborate on an example in more detail that will illustrate the approach of reusing the best matching result from the result set for a new query. For demonstration purposes we limited the number of attributes from the original application to a small amount.

Attribute	Range	Unit
Power Consumption (PC)	[0..10 *(Max_{PC})*]	W
Frequency (F)	[0.. 100(Max_F)]	Hz
Market Segment (MST)	$\in$ MST	

Figure 3: IP Attributes

As shown in Figure 3, an IP is described by three typed attributes. The *Power Consumption* is a real value from 0 to 10 and describes the fraction of the overall power consumption of the resulting microelectronic circuit, where the particular IP is integrated. The *Frequency* attribute describes the maximum possible frequency where the resulting circuit is allowed to operate. The last attribute, *Market Segment*, refers to intended application scenarios of the resulting circuit. The value is taken from a taxonomy depicted in Figure 4.

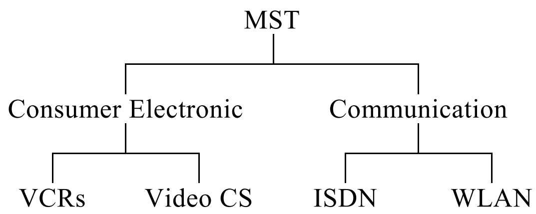

Figure 4: MST: Excerpt from the Market Segment Taxonomy

We will assume the following local similarity measures, which have been simplified, again, for illustration purposes.
If an IP from the case base has equal or lower power consumption than specified by in the query the local similarity is 1 otherwise as described in (1).

$$sim_{PC}(q_{PC}, ip_{PC}) = \begin{cases} 1 \quad if \quad q_{PC} \geq ip_{PC} \\ 1 - \left| \dfrac{q_{PC} - ip_{PC}}{Max_{PC}} \right| \quad else \end{cases} \tag{1}$$

Equation 1: Local similarity for power consumption

The local similarity measure for the IP frequency is described likewise with the difference that a higher or

[2] Multi-Dimensional Scaling

equal frequency, as specified in the query, leads to a similarity of 1.

$$sim_F(q_F, ip_F) = \begin{cases} 1 & if \quad q_F \leq ip_{FC} \\ 1 - \left| \dfrac{q_F - ip_F}{Max_F} \right| & else \end{cases} \quad (2)$$

Equation 2: Local similarity for frequency

For the market segment attribute, the similarity measure is defined by the length of the path from one node to another within the market segment taxonomy.

The maximum path length in the example taxonomy of Figure 4 is 4, e.g. from WLAN to VCRs node, leading to a similarity of 0,0 as defined be the similarity measure (3)

$$sim_{MST}(q_{MST}, ip_{MST}) = 1 - \frac{pathlength}{\max pathlength} \quad (3)$$

Equation 3: Local similarity for market segment

Finally, we use a simple weighted sum as aggregation function for assessing the global similarity (4)

$$\begin{aligned} agg(q, ip) &= 0{,}2 \cdot sim_{PC}(q_{PC}, ip_{PC}) \\ &+ 0{,}2 \cdot sim_F(q_F, ip_F) \qquad (4) \\ &+ 0{,}6 \cdot sim_{MST}(q_{MST}, ip_{MST}) \\ &with \quad w_{PC} + w_F + w_{MST} = 1 \end{aligned}$$

Equation 4: Aggregation function

It is expected that particular attribute values may not have been specified by the user in his/her query or captured in the case description. Hence, we define the following completion rules in order to deal with that situation:

Attribute not available in Query	⇒ local sim = 1,0
Attribute not available in Case	⇒ local sim = 0,5

Table 1: Completion rules

In our example case base we assume the following three cases.

	IP1	IP2	IP3
PC	5 W	6 W	7 W
F	40 MHz	50 MHz	50 MHz
MST	VCRs	ISDN	Communication

Table 2: Case Base

The user may now specify a query like

PC	5 W
F	40 MHz
MST	Not specified

Table 3: Example Query

and receives the following result set together with the individual similarity values:

	IP1	IP2	IP3
Total Similarity	**1,0**	**0,98**	**0,96**
PC (Weight: 0,2)	1,0	0,9	0,8
F (Weight: 0,2)	1,0	1,0	1,0
MST (Weight: 0,6)	1,0	1,0	1,0

Table 4: Example Result Set

Due to the fact that the MST attribute is not specified and the completion rule as defined before, the local similarity for the MST attribute is 1,0 for each case.
If we now reformulate the query by taking the best matching case from the result set, which is IP1, the ranking of the cases changes as follows:

	IP1	IP3	IP2
Total Similarity	**1,0**	**0,51**	**0,38**
PC (Weight: 0,2)	1,0	0,8	0,9
F (Weight: 0,2)	1,0	1,0	1,0
MST (Weight: 0,6)	1,0	0,25	0

Table 5: Result Set after Refinement

Of course, IP1 has a total similarity of 1,0 but IP3 has now a higher similarity than IP2. The reason is that the attribute with the highest impact on the total similarity has not been specified in the original query. When executing the best case as reformulated query, the ranking changes accordingly. The same effect can be observed, if the user specifies a value for the MST that does not closely match values from the case base.
Another important observation in this example is that the total similarities of the cases change significantly. From here we can propose the recommendation to the user to reformulate the query because, obviously, he/she did miss the most important attribute and the result set may be highly ambiguous.

6 Related Work

A very promising approach for explaining retrieval mismatches has been presented in [McSherry 2003]. Originating from the user query, the system generates sub-queries. E.g. a query with four specified attributes Q_{1234} is splitted into the sub-queries Q_1, Q_2, Q_3, Q_4, Q_{12}, Q_{13}, … , Q_{234}, then, for each sub-query it is analyzed if it is not covered by the case library (e.g.). With this information, an explanation for the user is generated informing him that e.g. the attributes corresponding to the query Q_{23} are not covered by any case. By making use of this information the user is able to refine his query accordingly. This approach simulates in a very simple but effective way a salesperson-like behavior.
[McArdle and Wilson 2003] present an approach for supporting the maintenance of large case bases by

means of dynamic visualization. They make use of the *sping based algorithm*, a variant of the already mentioned MDS algorithm.

7 Future Work

In this paper we presented our current research toward explanations for CBR-based applications. We briefly sketched potential application scenarios and identified complementing approaches. The ideas presented here are subject to implementation and integration into the IP retrieval prototype. As the example showed, it is a good starting point for further investigations, especially with large scaled case bases from the real world.

References

[Bergmann, 2002] Ralph Bergmann. *Experience Management – Foundations, Development Methodology, and Internet-Based Applications.* LNAI 2432, Springer Verlag, 2002.

[Lewis, 1997] Lewis, J. Intellectual property (IP) components. *Artisan Components, Inc., [web page]*, http://www.artisan.com(1997) [Accessed 28 Oct 1998].

[empolis 2001] empolis knowledge management GmbH (2001). The orenge Framework - A Platform for Intelligent Industry Solutions Based on XML and Java. *Whitepaper, empolis knowledge management GmbH*, Kaiserslautern.

[McArdle and Wilson, 2003] Gavin P. McArdle, David C. Wilson. Visualising Case-Base Usage. In *Workshop Proceedings of the 5th International Conference on Cased Base Reasoning CBR, Workshop: Long-Lived CBR System, Markus Nick and Mirjam Minor(Eds.) pages 105-124, Dublin*. NTNU Department of Computer and International Science technical report 4/2004.

[McSherry 2003] David McSherry. Explanation of Retrieval Mismatches in Recommender System Dialogues. In *Workshop Proceedings of the 5th International Conference on Cased Base Reasoning, pages 191-199, Dublin*. NTNU Department of Computer and International Science technical report 4/2004.

[McSherry 2003b] David McSherry. Similarity and Compromise. In *Proceedings of the 5th International Conference on Cased Base Reasoning, pages 291-305, Berlin 2003*. LNAI 2689 Springer-Verlag.

[Nick *et al*, 2002] Markus Nick, Klaus-Dieter Althoff, Thomas Avieny, Björn Decker. How Experience Management Can Benefit from Relationship among Different Types of Knowledge. In *Proceedings of the 1st German Workshop on Experinece Management, Mirjam Minor, Steffen Staab (Eds.) pages 95-106, Bonn 2002*. GI-Edition Lecture Notes in Informatics.

[Oehler *et al,* 1999] Peter Oehler, Ivo Vollrath, Peter Conradi, Ralph Bergmann, Thomas Wahlmann. *READee – Decision Support for IP Selection using a Knowledge-Based Approach*. IP'98, 1999.

[Schaaf *et al*, 2003] Martin Schaaf, Andrea Freßmann, Marco Spinelli, Rainer Maximini, and Ralph Bergmann. *A Knowledge Representation Format for Virtual IP Marketplaces*. ICCBR03 Trondheim, Norway 2003.

[Schaaf *et al*, 2002a] Schaaf, M., Maximini, R., Bergmann, R., Tautz, C., Traphöner, R. (2002). *Supporting Electronic Design Reuse by Integrating Quality-Criteria into CBR-Based IP Selection*, Proceedings 6th European Conference on Case Based Reasoning, September 2002.

[Schaaf *et al*, 2002b] Schaaf, M., Visarius, M., Bergmann, R., Maximini, R., Spinelli, M., Lessmann, J., Hardt, W., Ihmor, S., Thronicke, W., Franz, J., Tautz, C., Traphöner, R. (2002). *IPCHL - A Description Language for Semantic IP Characterization*, Forum on Specification & Design Languages, September 2002.

[Schmitt, 2003] Sascha Schmitt. Dialog Tailoring for Similarity-Based Electronic Commerce Systems**.** *PhD. thesis. University of Kaiserslautern, Germany*.