German Workshop on Experience Management (GWEM 2002)

Experience Management for Electronic Design Reuse through Quality-Oriented IP Selection

Ralph Bergmann[1], Rainer Maximini[1], Martin Schaaf[2]

[1]University of Hildesheim, Germany
[2]University of Kaiserslautern, Germany
[1]*{bergmann,r_maximi}@dwm.uni-hildesheim.de*
[2]*schaaf@informatik.uni-kl.de*

Abstract: The growing complexity of today's electronic designs requires reusing existing design components, called Intellectual Properties (IPs). Experience management approaches can be used to support design reuse, particularly the process of selecting reusable IPs. For the IP selection, quality criteria concerning the IP code and the documentation must be considered in addition to functional requirements of the IP. We analyse IP quality criteria in detail and show different concepts for their integration into the retrieval process.

1 Introduction

The design of electronic circuits is a discipline where two contrasting tendencies can be observed: On the one hand, modern circuit designs get more and more complex and difficult to handle by electronic engineers. On the other hand, global competition requires a continuous reduction of development times. At the same time, the correctness and reliability of the designs should, of course, not suffer from shorter development cycles.

These problems have become so dominant that they cannot be met anymore without extensive utilization of *design reuse*. It is getting vitally important for an electronic engineer to reuse old designs (or parts of them) and not to redesign a new application entirely from scratch. Reusing designs from the past requires that the engineer has enough experience and knowledge about existing designs, in order to be able to find candidates that are suitable for reuse in his specific new situation. The idea of design reuse is not new, but until recently, reusable components in electronic designs were of limited complexity and understandable by application designers. To reflect this growing complexity, the term *intellectual property (IP)* [Le97] has been assigned to those designs and today the term *reuse* means more than just plugging a component into a new environment. Due to this overall increased complexity of IP-based design, there is now a demand for knowledge management approaches that support electronic design processes. One main goal of such an approach is to provide user assistance in the question if an IP is suitable for a new design situation. This is one objective of the current project "IPQ: IP Qualification for Efficient Design Reuse" [1] funded by the

[1] IPQ Project (12/2000 – 11/2003). Partners: AMD, Fraunhofer Institut für Integrierte Schaltungen, FZI Karlsruhe, Infineon Technologies, Siemens, Sciworx, Empolis, Thomson Multi Media, TU Chemnitz, Universität Hildesheim, Universität Kaiserslautern, and Universität Paderborn. See www.ip-qualifikation.de

German Ministry of Education and Research (BMBF) and the related European Medea project "ToolIP: Tools and Methods for IP"[2].
This paper starts with a brief overview of the basic knowledge management considerations for electronic design reuse with IPs. Then, it focuses on one specific objective, which is to consider quality criteria of IP in addition to its functional specification during the IP selection process. Starting from existing approaches to measure IP quality in general, we propose and analyze different ways for integrating IP quality assessment approaches directly into an IP retrieval process that is realized using a case-based reasoning retrieval mechanism.

2 Knowledge Management for IP Selection

An IP is a design object whose major value comes from the skill of its producer [Le97], and a redesign would consume significant time. IP is just a new name for what formerly was called macro/mega cells, cores, or building blocks. The difference is the growing complexity of such cells (10k to 100k gates). Today, specialized IP vendors emerge starting to offer their IPs also in the Internet. In the near future, the IP market is expected to grow significantly. Therefore, the knowledge-based support for the exchange of qualified IP is needed, including general constraints and guidelines, as well as executable specifications for intra- and inter-company exchange. Furthermore, qualified IP must be made accessible via methodologies and tools, which will focus on the efficient implementation of all relevant IP management functions including *creation*, *storage*, *intelligent analysis* and *retrieval*, *validation* and *simulation* of IP. A related IP tailored design flow has to consider methods and tools for supporting IP checks done by IP vendors (called IP compliance checks), or customizable to the IP user needs (called IP entrance checks). The above-mentioned activities for IP management stem from the particular view of the electronics design area. However, from a knowledge management perspective, it becomes obvious that these are just particular instances of the typical processes [Wi97] proposed in the knowledge management literature. Moreover, the knowledge about IPs is of very specific nature since each knowledge item describes a particular design. Hence, from a knowledge management point of view, we deal with *experience management* [Be01a] that can be implemented in part by CBR technology as we have shown in the previous project READee [BV99] [Oe98a] [Oe98b].

2.1 The IP Selection Process

Figure 1 provides a brief overview of the IP selection process, which is part of an overall IP tailored design flow. It is assumed that IPs will be stored in an IP library (or several libraries in the future). During the *retrieval process*, an IP user specifies her/his requirements on the IP s/he is looking for and searches with this specification in the IP library for a reusable IP. The question whether an IP is reusable depends on many different technical criteria, such as the function of the IP, the technological realization

[2] See toolip.fzi.de for partners and further information.

on the chip, the design tool used for the design, and, very important, the quality of the IP, the quality of its design and the quality of the documentation.

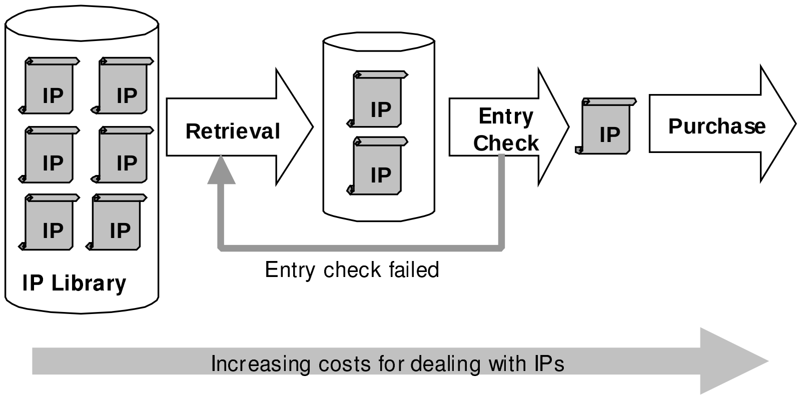

Figure 1: IP Selection Process

Since it is very hard to decide whether an IP is reusable, it cannot be expected that this can already be determined during the retrieval. The selection made during retrieval is therefore not based on the IP (the chip design code) itself but on a description of the IP, which we call *characterization*. Hence, the purpose of the retrieval task is therefore to make a pre-selection of a very small subset of IP candidates. The final decision on whether an IP is really reused is taken in the *entry check process*. The entry check requires getting at least partial access to the chip design code of the IP itself. Then, various compatibility checks and simulations of the IP are made in order to decide whether it can be really reused. Consequently, the entrance-check is very cost intensive regarding the required human and monetary resources. If an IP has passed the entry check, the IP user purchases it from the IP provider and only then the full IP code is disclosed to the IP user. If all selected IPs fail the entrance checked, the retrieval must be repeated or the design problem must be solved without reuse.

2.2 IP Retrieval

Given this design flow, it becomes obvious that the retrieval quality is of very high importance. If IPs are proposed that turn out not to be reusable, the large effort involved in the entry check is wasted. If reusable IPs are available but are not proposed, the opportunity for improving efficiency, quality, and time-to-market due to the reuse of this IP is lost. Hence, in principle, the IP retrieval should (to some degree) anticipate the subsequent entry check. Of course, it cannot replace it due to the large amount of (manual) effort involved, but it can approximate it as good as possible. A database of reusable designs requires a search facility that possesses more intelligence and knowledge than currently available systems can provide in order to achieve a high retrieval quality. In the near future, electronic applications will grow so complex that their designers cannot be expected to be specialists in the reusable components they are going to employ for their project. Hence, retrieval assistants capable of providing real selection support are highly recommended. It is one major goal of the experience management approach to reduce the amount of IPs that will be rejected during the entrance-check and, therefore, to decrease the overall costs for the IP selection process.

2.3 IP Representation

To achieve high precision retrieval it is essential to formalize knowledge about the IP itself as much as reasonable in order to establish a clear semantic that can be utilized by CBR technology. Therefore, we proposed from a retrieval point of view to divide the IP representation into two components, which is common for experience management [Be01a]:

1. The *IP characterization* that describes the IP in a way that allows to assess its quality and reusability in a particular situation.
2. The *IP content* that contains all deliverables of the design itself.

During retrieval only the IP characterization is used to rank the IPs contained in an IP library, not the IP content itself. Parts of the content are first used in the entry-check. The IP characterization describes all facts about the IP that are relevant for deciding whether the IP can be reused in a certain situation. The degree of detail used to characterize an IP determines how accurate its quality and reusability can be assessed, i.e., how accurate the retrieval is. The structure of the IP characterization knowledge (taxonomic and categorical knowledge) leads to an object-oriented (OO) representation [BS98] of the characterization. Such representations are particularly suitable for complex domains where characterizations with different structures occur, which is important when representing IPs.

The attributes of the IP characterization can be structured according to Table 1. The IP characterization is basically divided into two main parts:

- IP Application Attributes define all attributes that are important to decide about the applicability of an IP in a given design situation, and
- IP Quality Criteria characterize the IP and according to its quality.

This structure is compliant with the Virtual Component Attributes (VCA) Standard [VS01] and the OpenMORE Assessment Program for Hard/Soft IP [Sy01], which are documents released by an organization that aims at standardizing all IP related data.

<table>
<tr><th>Format</th><th>Category 0</th><th>Category 1</th><th>Category 2 / Taxonomy</th></tr>
<tr><td rowspan="10">IP Format</td><td>IP Content</td><td colspan="2"></td></tr>
<tr><td rowspan="9">IP Characterization</td><td rowspan="7">IP Application Attributes</td><td>Functional Class (Taxonomy)</td></tr>
<tr><td>Target Market Segments</td></tr>
<tr><td>Provider Claims</td></tr>
<tr><td>Integration Requirements</td></tr>
<tr><td>Reference Environment</td></tr>
<tr><td>IP Instance Attributes</td></tr>
<tr><td>Physical Description</td></tr>
<tr><td rowspan="2">IP Quality Criteria</td><td>Hard IP Criteria</td></tr>
<tr><td>Soft IP criteria</td></tr>
</table>

Table 1: IP Characterization

The advantage of dividing the IP characterization into these two main parts is that retrieval problems can be decomposed into two sub-problems by focusing either only on the application attributes or on the quality attributes. This allows us to decompose the selection process and the related similarity modeling into the two related parts, too, i.e., we will have a similarity measure for the application attributes and a second similarity measure for the quality attributes.

3 Quality-Based IP-Retrieval using the OpenMORE Approach

In the following, we will focus on the essentials of IP quality criteria. In order to give a first impression, we present an existing IP quality assessment approach, namely the *OpenMORE Assessment Program* [Sy01], along with an evaluation from the CBR perspective.

3.1 General Characteristics of IP Quality Criteria

Before we turn to the OpenMORE approach, we fix some assumptions about the nature of quality criteria. This will be done here limited to retrieval aspects. These basic assumptions about IP quality criteria are:

1. To each quality criterion a data type can be assigned. The type defines the range of possible values for quality rating. The type of a particular criterion is strictly ordered, like score values. Here, order means that a higher value is always preferred over a lower value.
2. There exist neither dependencies between different quality criteria nor between application attributes and quality criteria. Hence, quality criteria can be considered as independent variables. E.g. the delivery of test-bench files always results in higher quality no matter which kind of functionality the IP module supports.
3. Quality attributes can be structured into categories giving the user the possibility to specify preferences for a set of thematically connected criteria instead of individual criteria.
4. There exist quality criteria that are not applicable to all types of IPs. For example, depending on the description language VHDL or Verilog different coding guidelines may apply.

This leads to the following general definition of an IP quality criterion:

A quality criterion is an attribute associated with an IP containing a textual description of the quality item that is measured and a result value taken from an associated ordered data type. The type of the score values is a range of natural numbers enhanced by two special values indicating that the criterion is not applicable to a given IP or it has not been assessed. Furthermore, each quality criterion can be associated with a category that allows aggregating thematically connected criteria. The categories can be structured within a hierarchy that will be named the IP quality classification.

3.2 Overview of the *OpenMORE Quality Assessment* Program

In the past large efforts have been already made to identify and classify relevant quality criteria with respect to reusability of IPs. The best-known and probably most complete catalog published is the OpenMORE Assessment Program [Sy01], which will be introduced in this section. It identifies around 150 different quality criteria for soft IPs and 130 for hard IPs. Beside the catalog of quality criteria the OpenMORE Assessment Program also contains some guidelines to the assessment procedure in order to insure comparability between IPs assessed by different organizations.

RMM2 Section	Description	Type	Assessment	Unweighted Score	Max Score
I	**Macro Design Guidelines**			**0**	**542**
3	***System-Level Design Issues: Rules and Tools***			***0***	***88***
3.2	Design for Timing Closure: Logic Design Issues			0	34
3.2.2	Synchronous vs. Asynchronous Design Style			0	10
3.2.2.1	System is synchronous and register based with latches used only to implement small memories or FIFOs. FIFOs and memories designed so they are synchronous and edge triggered. Exceptions fully documented.	R		0	10
3.2.3	Clocking			0	14
3.2.3.1	Number of clock domains and clock frequencies are well documented. Required clock frequencies and associated phase lock loops (PLL) and external timing requirements (setup/hold and output timing) fully documented.	R		0	10
3.2.3.2	*If two asynchronous clock domains interact, then they meet as a single module which is as small as possible. Use the smallest possible number of clock domains.*	G		*0*	*2*
3.2.3.3	*If a phase locked loop (PLL) is used for on-chip clock generation, then a means of disabling or bypassing the PLL is provided.*	G		*0*	*2*
3.2.4	Reset			0	10
3.2.4.1	The basic reset strategy for the chip is documented, especially 1) synchronous or asynchronous, 2) internal or external power-on reset, 3) more than one reset (hard vs. soft reset), 4) each macro individually resettable, for debug purposes	R		0	10
3.4	Design for Verification: Verification Strategy			0	22
3.4.1	The system level verification strategy must be developed and documented before macro selection or design begins.	R		0	10

Figure 2: OpenMORE Assessment Program

Figure 2 shows a part of the OpenMORE catalog illustrating the classification of IP quality criteria. The IP quality criteria are grouped into the three major categories *Design* (50%), *Verification* (35%), and *Deliverable Guidelines* (15%) that are weighted differently. The other sub-categories are simply for structuring purposes. Each quality criterion is either called a *design rule* (denoted by ”R” in the type column) or a *design guideline* (denoted by ”G” respectively). The only difference between the two is that in case of a rule possible scores are 0, 5, or 10, in case of a guidelines possible scores are 0, 1, or 2. If a criterion is not applicable to an IP its score becomes 0 and the possible maximum score for the associated category is reduced.

3.3 The OpenMORE Assessment Approach from a CBR Retrieval Perspective

The objective of the OpenMORE Assessment Program is to calculate one single value for quality that allows a global ranking of different IP modules. This ranking is global because it is independent from individual requirements that an IP user might have. The assessed scores within each of the top-level categories (which are named ”Macro Design Guidelines”, ”Macro Verification Guidelines”, and ”Deliverable Guidelines”) are summed up considering the fixed weight values 0.5, 0.35, and 0.15 respectively.
From a CBR perspective this approach can be easily transformed into a simple similarity measure that is structured according the given categorization of the quality attributes. Of course, since the OpenMORE Assessment Program only supports a global quality value, the resulting similarity measure is somewhat artificial since the similarity does not depend at all on a query specifying a required quality. Furthermore, an OpenMORE induced similarity model would only consist of three categories, each

one aggregating rules and guidelines. The overall relevance of a quality criterion is then as follows

$$\frac{w_{category}}{10 \cdot n_{rules} + 2 \cdot n_{guidelines}} \cdot \frac{assessedScore}{MaxScore}$$

The term *assessedScore* can be 0, 5, or 10 for a rule and 0,1,2 for a guideline. It can be seen that the relevance of a criterion depends on the top-level category (through w_{category}), on the kind of the quality criterion (rule or guideline), and on the number of other rules n_{rules} and guidelines $n_{guidelines}$ aggregated in that category. The overall influence of a single criterion does not depend on the nature of the quality criterion itself. Furthermore, due to the very ambiguous handling of conditional criteria only applicable to some IPs, the number of rules and guidelines within a category may change depending on the particular IP assessed. Hence, the relevance of a quality criterion changes in an unpredictable way for different IPs. It follows that a naive adaptation of the OpenMORE Assessment approach would not be a good base for efficient automated retrieval. Finally, it is not clear how the relevance of IP quality criteria induced by the OpenMORE Assessment Program approximates the overall utility [Be01b] with respect to reusability in any way.

Due to the generally agreed importance of quality criteria it is absolutely necessary to develop a more sophisticated approach better reflecting the relevance of quality criteria. Furthermore, such an approach must be flexible in the sense that it can be adapted to newly identified criteria, newly developed classification schemes, and new insights about the overall utility of particular quality criteria.

4 Improved Quality-based Retrieval

Despite the fact that a naive implementation of the OpenMORE Assessment Program for retrieval would have several drawbacks, it can be taken as a starting point for a CBR-based retrieval. We can make use of the criteria themselves as well as their structuring in categories. Hence, for the following we take the categorization scheme shown in Figure 3 as base for a new structured similarity model.

4.1 Quality Requirements from Users

An important requirement for IP retrieval is to give the IP user the possibility to define a *quality specification*. The quality specification consists of *quality requirement q* and a *weight model w*. Initially, the quality requirement and the weight model use the categorization scheme shown in Figure 3. The similarity model defining the similarity between a quality requirement *q* and the quality attributes of an IP *c* can be developed corresponding to the categorization scheme by weighting *each attribute* and *each category* according to the weight model. The categorization of the criteria allows the IP user to abstract from basic criteria by focusing on the overall relevance of a category. Furthermore, it is intended to allow the specification of relevance of criteria/categories not only by forcing the IP user to provide weight values directly, which could become a very cumbersome task. Instead, relevance can be also expressed by relating (e.g. sorting or partial ordering) different attributes respective categories to each other. This can be seen as specifying their relative relevance and depends on the query. From this ordering of attributes, a weight model can be determined automatically.

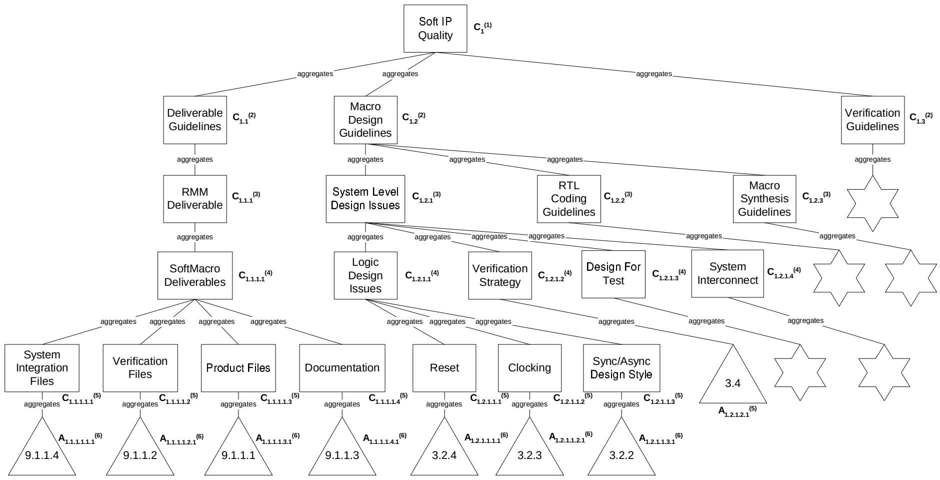

Figure 3: Categorization Schema for Quality Criteria

4.2 Basic Similarity Measure

We apply the traditional local-global principle by, first, defining similarity from a global point of view considering the categories $C_i^{(l)}$ and, second, from a local point of view taking into account the details of the similarity between specific quality criteria [Be01a] [Be01b] [BS98].
As usual, the global similarities are modeled by aggregating local similarity with the aggregation function ϕ, which encodes the knowledge about the importance, relevance and utility. ϕ is defined for each category $C_i^{(l)}$ as

$$\phi_{C_i^{(l)}} : [0,1]^n \rightarrow [0,1]$$

$$\phi_{C_i^{(l)}}(s_1,\ldots,s_{n)}) = \sum_{k=1}^{n} \omega_{C_i^{(l)},k} \cdot s_k$$

Here, *n* is the number of attributes and subcategories aggregated in the category $C_i^{(l)}$, $\omega_{C_i^{(l)},k}$ is the weight of attribute s_k such that $0 \le \omega_{C_i^{(l)},k} \le 1$ and $\sum_{k=1}^{n} \omega_{C_i^{(l)},k} = 1$ holds. The vector $\omega_{C_i^{(l)}}$ is called *category specific weight model* and can be defined for each particular category. Now, assuming *Q* as the set of all possible quality requirements and *C* the set of all possible quality characterizations for each category $C_i^{(l)}$ the similarity function $sim_{C_i^{(l)}}$ can be recursively defined as:

$$sim_{C_i^{(l)}} : Q \times C \rightarrow [0,1]$$

$$sim_{C_i^{(l)}}(q,c) = \phi_{C_i^{(l)}}(sim_{C_i^{(l+1)},1}(q,c),\ldots, sim_{C_i^{(l+1)},m}(q,c), sim_{A_i^{(l+1)},1}(q,c),\ldots, sim_{A_i^{(l+1)},n}(q,c))$$

Here, $C_{i,1}^{(l+1)},...,C_{i,m}^{(l+1)}$ denote the subcategories of $C_i^{(l)}$ while $A_{i,1}^{(l+1)},...,A_{i,n}^{(l+1)}$ are the attributes of $C_i^{(l)}$. The value of $sim_{A_i^{(l+1)}}$ represents the local similarity of two basic quality criteria as defined below. The global similarity of the quality requirement of the IP-User q and the quality part of the IP characterization c will be defined as: $sim(q,c) = sim_{C_1^{(1)}}(q,c)$.

The local similarity measures used here determine how well a certain IP is reusable regardless of a deviation with respect to a single quality criterion. As a general rule we can say that an IP is reusable if it has a higher quality than required by the user. Hence, typical local similarity measures for quality attributes are asymmetric, such as, for example:

$$sim_{A_i^{(l)}} : T_{A_i^{(l)}} \times T_{A_i^{(l)}} \rightarrow [0,1]$$

$$sim_{A_i^{(l)}}(q,c) = \begin{cases} 1 - \dfrac{q-c}{\max_{A_i^{(l)}}} & if \quad q > c \\ 1 & else \end{cases}$$

4.3 User-Defined Categorization Schema

Reflecting the characteristics of IP quality criteria in the similarity model results in a better approximation of utility than it can be achieved with the flat OpenMORE approach. The structure does not only enable IP-users to specify their preferences; it also serves as a base for communication about quality criteria, their characteristics, and their impact on reusability of IPs. Especially the impact on reusability is not very well researched, yet, in the electronics design area. Consequently, the IP categorization scheme must be considered as preliminary being subject of changes according to further developments. Hence, in a subsequent step we want to give the IP user the possibility to modify the categorization scheme, which would allow her/him to define their own categories aggregating basic quality criteria or sub-categories. Therefore, attribute and categories can be moved up and down in the schema.

In order to manage different categorization schemas, it is necessary to store the quality specification together with the individual categorization scheme within a user profile for further evaluation. Collecting this kind of knowledge enables the necessary learning-cycle for improving the IP-quality classification toward a better approximation of utility during usage of the system.

4.4 Combination with Application Attributes

Obviously, any retrieval mechanism for IP must consider both, application criteria and quality criteria together. We now present possibilities of how the application and quality aspects can be combined during retrieval. For the discussion of the integration aspects, both types of selection criteria can now be considered as black boxes, for which the similarities sim_Q and sim_A can be determined.

The first possibility for integration is to apply the simple weighting average of the two. We can assign a weight to each type of criteria. With this approach, IP retrieval returns

similar cases depending on the quality and functionality. The advantage of this technique is that the IP user can specify the importance of quality in relation to the functionality for him/herself. However, the disadvantage of this approach is that a lack of appropriate IP quality can be compensated by a more appropriate application characteristic and vice-versa. There's a chance that IPs with proper application characteristics are not retrieved due to bad quality but other IPs are retrieved with inappropriate application characteristics just because they have a high quality.
For the second possibility called *strict prioritized combination*, we do not combine the similarities to a global similarity. Instead we give the main priority to sim_A, i.e., we first retrieve IPs, which "fit" only by their application characteristics (ignoring the quality). Only in a second step, we take the IP-Quality into account by ordering the retrieved IPs according to sim_Q. The result-set of the first retrieval step is then clustered into different functional clusters. Building these clusters is based on the similarity between the query and the IP as well as on the similarity between the IPs among each other. Conceptual clustering algorithms as developed to large extend in the machine learning literature can be applied to compute such a clustering. As a result of the first step, we now have achieved a set of clusters, each of which contains a set of IPs that are similar to each other with respect to their applicability characteristics. How for each cluster the IPs are sorted according to the quality, i.e. using sim_Q.

5 Architecture of the IP-Retrieval System

The IP retrieval software currently under development will be designed for three different user classes. The *IP Vendor*, the *IP Assessor* responsible for assessing the quality of an IP, and the *IP User*. Throughout this paper we have not distinguished between the *IP Vendor* and *IP Assessor*. Instead, we referred to both by the term *IP Provider*, but for the system architecture it makes sense to distinguish between these user classes because a very likely usage scenario is a web site that acts as a portal for several *IP Vendors*. In this case, the *IP Assessor* has at least the responsibility to maintain the standardized characterization of each IP. Figure 4 shows the necessary steps for releasing an IP. As soon as a standardized IP characterization is available, the IP can be checked into the retrieval system. The development of an XML application for IP description is another task accomplished by the IPQ project.

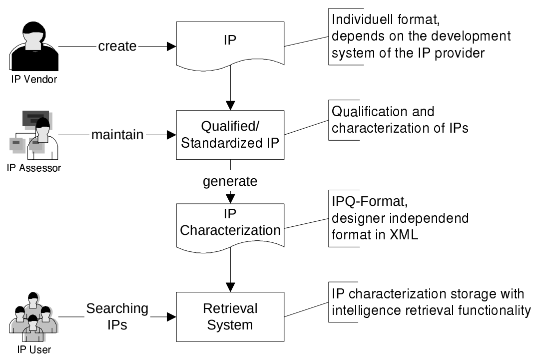

Figure 4: IP Release Process

The IP retrieval system itself integrates several components as shown in Figure 5. One main component is the open retrieval engine ORENGE[3], a modular system providing a variety of functionality beside the essential retrieval engine (e.g. a dialog module, a chat module, a module for text mining, an explanation module and a profile manager). The profile manager is used to store additional information about a user, like personalized similarity measures. This gives an *IP User* the flexibility to specify his own preferences when dealing with quality criteria.

The other components of the IP retrieval system are mainly for interacting with the different users and converting the IP Characterization to a format named OOML. OOML (ORENGE Object Mark-Up Language) is an internal XML format for the ORENGE retrieval engine and is optimized to support the retrieval process.

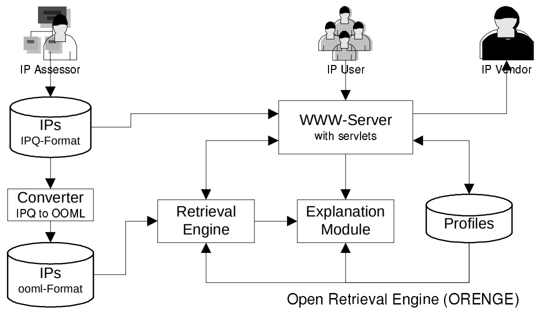

Figure 5: IP Release Process

As depicted in Figure 5, the IP Assessor stores the IP Characterization (e.g. via ftp or http upload) directly in an IP database. The user access is completely handled and controlled by Java Servlets running on a WWW server. In order to give IP Providers important feedback for adapting their products or marketing purposes, the system tracks the user actions and retrieval results.

6 Conclusion

In this paper we developed a concept for the integration of quality criteria into IP retrieval, which is a significant improvement with respect to the current state in IP qualification. The new concepts allow considering detailed quality preferences of IP users; we have also shown how these preferences can be integrated into a CBR-like retrieval. Further, we propose methods for capturing the preferences in such a way that they can be immediately converted into a related graphical user interface for the retrieval engine to be developed. We also propose to establish a quality classification improvement strategy that allows coping with the high dynamics of this field and enables the acquisition of typical quality requirements of IP users. This information is also highly valuable input for IP providers since it contains the knowledge about the demands of IP users concerning IP quality and enables the IP providers to develop IPs toward market needs. Finally, we discussed two approaches for combining quality-

[3] ORENGE stands for Open Retrieval Engine and is the current commercial CBR software product by empolis.

based retrieval with retrieval based on application characteristics. Although most of the presentation in this paper seems very specific to the electronic design domain, we think that particularly the nature of quality criteria and their treatment can be transferred to experience items in other domains where quality plays an important role, too.

References

[Be01a] Bergmann, R (2001). Experience Management: Foundations, Development Methodology, and Internet-Based Applications. Habilitationsschrift, Universität Kaiserslautern.

[Be01b] Bergmann, R., M. M. Richter, S. Schmitt, A. Stahl, & I. Vollrath (2001). Utility-oriented matching: A new research direction for case-based reasoning. In S. Schmitt, I. Vollrath, and U. Reimer (Eds.), 9th German Workshop on Case-based Reasoning.

[BR01] Burkhard, H. D. & Richter, M. M. (2001). On the Notion of Similarity in Case-Based Reasoning and Fuzzy Theory. In: S. Pal, T. S. Dillon, & D. S. Yeung (Eds.) Soft Computing in Case Based Reasoning. Springer.

[BS98] Bermann, R. & A. Stahl (1998). Similarity measures for object-oriented case representation. In B. Smyth & P. Cunningham (Eds.) Advances in Case-Based Reasoning (EWCBR'98), Lecture Notes in Artificial Intelligence, Springer, 25-36.

[BV99] Bergmann, R. & Vollrath, I. (1999). Generalized cases: Representation and steps towards efficient similarity assessment. In W. Burgard, Th. Christaller & A. B. Cremers (Hrsg.) KI-99: Advances in Artificial Intelligence, Lecture Notes in Artificial Intelligence, 1701, Springer, 195-206.

[Le97] Lewis (1997). Lewis, J.: Intellectual Property (IP) Components. Artisan Components, Inc., http://www.artisan.com/ip.html, (1997).

[Oe98a] Oehler, P., Vollrath, I., Conradi, P., Bergmann, R., & Wahlmann, T. (1998). READEE - Decision support for IP selection using a knowledge-based approach. In IP98 Europe Proceedings, Miller Freeman.

[Oe98b] Oehler, P., Vollrath, I., Conradi, P., & Bergmann, R. (1998). Are your READee for IPs?. In A. Kunzmann & R. Seepold (Hrsg.) Second GI/ITG/GMM Workshop "Reuse Techniques for VLSI Design", FZI-Report 3-13-9/98, Universität Kalrsruhe.

[Sy01] Synopsis Inc., Mentor Graphics Corporation (2001). OpenMORE Assesment Program for Hard/Soft IP Version 1.0, http://www.openmore.com.

[VS01] VSI AllianceTM, Virtual Component Transfer Development Working Group (2001). Virtual Component Attributes (VCA) With Formats for Profiling, Selection, and Transfer. Standard Version 2.2 (VCT 2.2.2).

[Wi97] Wiig, K. (1997). Knowledge Management: Where Did It Come From and Where Will It Go?. Expert Systems with Applications (14). Pergamon Press / Elsevier.