

Proceedings of the
First International Conference on
Web Information Systems
Engineering

WISE-2000

Adelstein, 1999
Hong Kong Polytechnic University
D.J. Computer Science, Hong Kong, China
Kathleen Lee, Kyoto University
K.F. Wong Education Foundation
San Francisco, CA



Los Angeles, California
Washington Brussels Tokyo

Proceedings of the First International Conference on Web Information Systems Engineering

Volume I (Main Program)

Hong Kong, China

June 19 – 21, 2000

Editors


Qing Li, Z. Meral Ozsoyoglu, Roland Wagner, Yahiko Kambayashi, Yanchun Zhang

Co-Organized by

City University of Hong Kong
University of Southern Queensland
ACM (Hong Kong Chapter)
Hong Kong Web Society

Co-Sponsored by

AIMtech Centre
Hong Kong Pei Hua Education Foundation
IEEE Computer Society, Hong Kong Section, Computer Chapter
Kambayashi Lab, Kyoto University
K.C. Wong Education Foundation
Sun Microsystems



Los Alamitos, California

Washington = Brussels = Tokyo

TA168
.I518
2000

Copyright © 2000 by The Institute of Electrical and Electronics Engineers, Inc.

All rights reserved

v.1

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 133, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society, or the Institute of Electrical and Electronics Engineers, Inc.

IEEE Computer Society Order Number PR00577

ISBN 0-7695-0577-5

ISBN 0-7695-0578-3 (case)

ISBN 0-7695-0579-1 (microfiche)

Library of Congress Number 00-101162

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1-714-821-8380
Fax: + 1-714-821-4641
E-mail: cs.books@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1-732-981-0060
Fax: + 1-732-981-9667
[http://shop.ieee.org/store/
customer-service@ieee.org](http://shop.ieee.org/store/customer-service@ieee.org)

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81-3-3408-3118
Fax: + 81-3-3408-3553
tokyo.ofc@computer.org

Editorial production by Bob Werner

Cover art production by Alex Torres

Printed in the United States of America by The Printing House




Table of Contents – Volume I

First International Conference on Web Information Systems Engineering — WISE'00

Message from the General Chairs	xi
Message from the Program Chairs	xii
Program Committee	xiii
Conference Organization	xiv

Keynote Address

Streaming Real-Time Audio and Video Data with Transformation-based Error Concealment and Reconstruction	2
<i>B. Wah, D. Lin, X. Su</i>	

Keynote Address

Hyperdatabases	14
<i>H. Schek, K. Böhm, T. Grabs, U. Röhm, H. Schuldt, R. Weber</i>	

Databases and the Web

A Middleware Approach for Combining Heterogeneous Data Sources — Integration of Generic Query and Predefined Function Access	26
<i>K. Hergula, T. Härder</i>	

Efficient Web-based Access to Multiple Geographic Databases through Automatically Generated Wrappers	34
<i>S. Cha, K. Kim, C. Song, Y. Kwon, S. Hwang</i>	

Multicache-based Content Management for Web Caching	42
<i>K. Cheng, Y. Kambayashi</i>	

Shift it to the Server! — Let the Database Server Update Your Websites	50
<i>H. Loeser</i>	

Optimal Placement of Proxies of Replicated Web Servers in the Internet	55
<i>X. Jia, D. Li, X. Hu, H. Huang, D. Du</i>	

XML and Semi-Structured Data

XML Schema Directory: A Data Structure for XML Data Processing	62
<i>E. Kotsakis, K. Böhm</i>	

XML Content Management based on Object-Relational Database Technology	70
<i>B. Surjanto, N. Ritter, H. Loeser</i>	

Interoperability and Semi-Structured Data in an Open Web-based Agent Information System	80
<i>H. Lu, L. Sterling</i>	

Distributed Object-Oriented Schema for XML-based Electronic Catalog Sharing Semantics among Businesses.....	87
<i>K. Kuramitsu, K. Sakamura</i>	
Architecture of a Content Management Server for XML Document Applications	97
<i>T. Arnold-Moore, M. Fuller, A. Kent, R. Sacks-Davis, N. Sharman</i>	
XML based Text TV.....	109
<i>T. Vuorimaa, C. Sancho</i>	
Web Documents and Security/Control Issues	
A Semantic Similarity Approach to Electronic Document Modeling and Integration.....	116
<i>W. Song, D. Cheung, C. Tan</i>	
Global Atlas: Calibrating and Indexing Documents from the Internet in the Cartographic Paradigm	125
<i>F. Lee, S. Bressan, B. Ooi</i>	
Dynamic Aspects of Documents of the Brazilian Web.....	133
<i>N. Fonseca, R. Resende, W. Meira, Jr</i>	
A Role-based Access Control Model for XML Repositories.....	138
<i>H. He, R. Wong</i>	
Logging and Signing Document-Transfers on the WWW — A Trusted Third Party Gateway.....	146
<i>A. Heuer, F. Losemann, C. Meinel</i>	
Using Class Decompilers to Facilitate the Security of Java Applications!	153
<i>V. Tam, R. Gupta</i>	
The Most Efficient Uniquely Decipherable Encoding Schemes.....	159
<i>D. Long, W. Jia</i>	
Transactions and Workflows	
A Web-based Transactional Workflow Monitoring System.....	166
<i>H. Hong, B. Lee, K. Kim, S. Paik</i>	
Web Interface-Driven Cooperative Exception Handling in ADOME Workflow Management System.....	174
<i>D. Chiu, Q. Li, K. Karlapalem</i>	
Timed Workflow: Concept, Model, and Method.....	183
<i>Z. Hai, H. Pung, T. Cheung</i>	
Consistency Preserving in Transaction Processing on the Web.....	190
<i>K. Lee, H. Kim</i>	
A Transaction Management Scheme with Mark in Decision Supporting Environments.....	196
<i>C. Park, S. Park</i>	
Searching and Querying on the Web	
WebSail: From On-line Learning to Web Search	206
<i>Z. Chen, X. Meng, B. Zhu, R. Fowler</i>	

<i>WhatNext: A Prediction System for Web Requests using N-gram Sequence Models</i>	214
<i>Z. Su, Q. Yang, Y. Lu, H. Zhang</i>	
Searching the Web using a Map	222
<i>J. Yates, X. Zhou</i>	
A Semantic Associative Search Method for WWW Information Resources	230
<i>Y. Kiyoki, X. Chen, T. Kitagawa</i>	
Querying Multiple Databases Dynamically on the World Wide Web	238
<i>J. Cardiff, T. Catarci, M. Passeri, G. Santucci</i>	
QUERYAID: A Query Formulation Model for and by Reuse	246
<i>A. Seriai, C. Oussalah</i>	
A Web Site Indexing Process for an Internet Information Retrieval Agent System	254
<i>S. Cazalens, E. Desmontils, C. Jacquin, P. Lamarre</i>	
Modeling and Meta-Data Issues	
Modelling the Webspace of an Intranet	260
<i>R. van Zwol, P. Apers</i>	
Modeling Business Processes for Web-based Information Systems Development	270
<i>J. Dietz, H. Schouten</i>	
Concept Hierarchy based Text Database Categorization in a Metasearch Engine Environment	283
<i>W. Wang, W. Meng, C. Yu</i>	
Model and Research Issues for Refreshing a Very Large Website	291
<i>H. Liu, W. Ng, E. Lim</i>	
Automatic Generation of Ontology based Annotations in XML and Their Use in Retrieval Systems	296
<i>B. Czejdó, J. Dinsmore, C. Hwang, R. Miller, M. Rusinkiewicz</i>	
Information Visualization	
Visualization of Web-based Workspace Structures	302
<i>R. Biuk-Aghai</i>	
Visualization of Linear Time-Oriented Data: A Survey	310
<i>S. Silva, T. Catarci</i>	
Database and Modelling Strategy: A Compliant Way for Display Optimisation	320
<i>N. Farenc, F. Sidler, A. Ferrando, D. Thalmann</i>	
Supporting a Focus+Context Interaction Style for Spatial Databases	328
<i>S. Cicerone, D. Frigioni, L. Tarantino</i>	
Visualization and Control Functions for Distributed Discussion Systems	336
<i>M. Yuriyama, A. Hatanaka, H. Tarumi, Y. Kambayashi</i>	

Visual Support for Text Information Retrieval based on Matrix's Singular Value Decomposition	344
<i>J. Hou, Y. Zhang, J. Cao, W. Lai</i>	
A Visual Representation of Search-Engine Queries and their Results.....	352
<i>R. Grewal, M. Jackson, P. Burden, J. Wallis</i>	
Data Mining, Data Warehousing/OLAP	
Metadata Management Concept for Multidimensional OLAP Data based on Object-Oriented Concepts	358
<i>O. Mangisengi, A. Tjoa, R. Wagner</i>	
Mining Term Association Rules for Automatic Global Query Expansion: Methodology and Preliminary Results.....	366
<i>J. Wei, S. Bressan, B. Ooi</i>	
A Model for Evaluating Materialized View Maintenance Algorithms.....	374
<i>T. Ling, E. Sze</i>	
Handling New Media/Multimedia Information	
A Signature Representation and Indexing Scheme of Color-Spatial Information for Similar Image Retrieval	384
<i>S. Jo, K. Um</i>	
VideoMAP*: A Web-based Architecture for a Spatio-Temporal Video Database Management System	393
<i>S. Chan, Q. Li</i>	
Conflict Control Locking in Distributed Cooperative Graphics Editors.....	401
<i>L. Xue, K. Zhang, C. Sun</i>	
Interactive Graphics for Web based Applications.....	409
<i>M. Sweeney</i>	
Optional and Responsive Locking in Distributed Collaborative Object Graphics Editing Systems.....	414
<i>D. Chen, C. Sun</i>	
Speech Recognition using Artificial Neural Networks	419
<i>C. Lim, S. Woo, A. Loh, R. Osman</i>	
New Services, Applications and Development Techniques	
Advertising Services in the Internet Marketplace	426
<i>J. Yang</i>	
Automatic Generation of Reliable E-Commerce Payment Processes.....	434
<i>H. Schuldt, A. Popovici, H. Schek</i>	
Developing Web-Aware Systems	442
<i>N. Šerbedžija</i>	
A Fast and Flexible Framework of Scripting for Web Application Development: A Preliminary Experience Report.....	450
<i>V. Tam, W. Foo, R. Gupta</i>	

Future Smart-Organizations: A Virtual Tourism Enterprise	456
<i>H. Afsarmanesh, L. Camarinha-Matos</i>	
Design of Data-Intensive Web-based Information Services	462
<i>T. Feyer, O. Kao, K. Schewe, B. Thalheim</i>	
A Static Java Classfile Analyzer	468
<i>M. Jung, D. Kim, K. Park, M. Jin, W. Kim, S. Lee, J. Kang</i>	
Internet EDI Implementation to Support Multiple Document Standards for R&D Project Management.....	473
<i>C. Liang, J. Ma</i>	
Communications and Mobile Computing	
A Whole Correlation Structure of Asymptotically Self-Similar Traffic in Communication Networks	480
<i>M. Li, W. Jia, W. Zhao</i>	
Supporting Online Web-based Teamwork in Offline Mobile Mode Too	486
<i>Y. Yang</i>	
Towards an Integrated Systems Approach for Mobile Traveller Applications.....	491
<i>P. Peinl, B. Mitschang</i>	
Index of Authors	497

XML Content Management based on Object-Relational Database Technology

B. Surjanto, N. Ritter, H. Loeser

University of Kaiserslautern, Department of Computer Science,
P.O. Box 3049, D-67653 Kaiserslautern, Germany
E-mail: {surjanto, ritter, loeser}@informatik.uni-kl.de

Abstract

XML (Extensible Markup Language) is a textual markup language designed for the creation of self-describing documents which contain textual data combined with structural information describing the structure of the textual data. We introduce XCoP (XML Content Repository) as a repository which is based on an object-relational database management system (ORDBMS) and improves content management of XML documents thereby exploiting their structural information. It allows users to reuse and process textual portions of document contents, called fragments, which are flexibly configurable. Moreover, it enables collaborative development of documents and facilitates synchronization of fragment modification and versioning. Thus, XCoP offers comprehensive content management functionality by taking advantage of the availability of structural information.

1. Introduction

XML (Extensible Markup Language) has been recommended by W3C (World Wide Web Consortium) as a new standard document markup language for the Web [5]. It was developed as a textual markup language designed for the creation of self-describing documents, i.e., XML documents contain textual data interspersed with structural information. The structural description may also be used to express specific document semantics in order to facilitate the organization of and search in large sets of documents.

Unfortunately, current products and approaches for document-oriented application domains such as web content management (WCM), document management (DM), document content management (DCM), or digital libraries, concentrate on the textual representation when processing and analyzing documents. They do not take advantage of the availability of structural information and do not adequately and completely support other relevant aspects of content management. In contrast, existing approaches for structure-oriented application domains, e.g., data integration, management of semistructured data, or search engines, prefer

very fine granularity such as markup elements within a document in order to exploit the corresponding structural information and give less attention to operations processing textual document contents. We are convinced that processing support to both, textual document contents and structural information, would be very useful for both kinds of application domains and could lead to substantial improvements when large sets of documents have to be organized, classified and searched.

In this paper, we introduce XCoP (XML Content Repository) which is based on object-relational database (ORDB) technology and manages textual XML document contents together with their structural information. XCoP allows the reuse of document contents by providing so-called *fragments*. The use of our fragment concept reduces many operational problems. It avoids data redundancy and, therefore, update anomalies on replicated document contents. Furthermore, it enables collaborative development of documents, since multiple users may work on different fragments of a document concurrently. Unlike existing tools and approaches, the fragment granularity in XCoP is flexibly configurable by the users. Depending on their needs, users may specify, split, and merge fragments at runtime in suitable units they want to share. At the same time, XCoP extracts the structural information of document contents in order to make query facilities and other repository services more effective. These structural descriptions together with their corresponding relations to document contents are maintained.

With the advent of object-relational database management systems (ORDBMSs) [25], developers can extend the database system by user-defined data types (UDT), user-defined routines (UDR), and user-defined index structures. Furthermore, ORDBMSs allow users to access external data sources via SQL (Structured Query Language), enable (limited) object-oriented data modeling, still keep the advantages of relational database management systems (RDBMSs), and coexist with legacy data which mostly resides in relational databases (RDBs). We develop XCoP as part of an integrated Web database framework, called *iWebDB* [21], which is built by our database research group

using the extensibility features of an ORDBMS. In this paper, we deal with the content management of XML documents in XCoP and do not discuss other value-added repository services [2] provided by XCoP such as version control, configuration management, collaboration, context management as well as workflow control.

The rest of this paper is organized as follows: Section 2 provides a list of requirements for the content management of XML documents. In Section 3 and 4 we describe our approach for managing XML document contents and its implementation on the top of an ORDBMS. Previous related work is then reviewed in Section 5. Finally, Section 6 gives concluding remarks and an outlook to future work.

2. Requirements for the Content Management of XML Documents

We believe that XML is the key technology for the next generation of Web applications. Its simplicity, extensibility, and structuring capability make XML very attractive for various applications domains. In addition to the unidirectional hyperlinks of HTML (Hypertext Markup Language), a supplemental technique of XML, called XLink (XML Linking Language, [13]), provides more sophisticated linking constructs (e.g., extended links). Moreover, the capability of RDF (Resource Description Framework, [22]) to describe metadata (representing semantics) of Web resources makes the Web more appealing for both information providers and consumers. Finally, XML and XSL (Extensible Stylesheet Language, [8]) together offer a solution for a media-independent publishing that is achieved through the separation of the content and presentation format. Content creators can therefore concentrate on producing their document contents and information consumers may choose their preferred views. The potential impact of XML is obvious. Using XML as a single language for specifying the huge amount of information produced and consumed over the Web would provide better interoperability among information providers as well as easier handling of documents by applications. Consequently, we will have to be prepared to deal with a very large number of XML documents and their contents will have to be managed adequately and efficiently. As we think, a corresponding content management system has to fulfill the following requirements:

- **Content Reuse**

Sharing information for reuse must be considered being an essential property of a content management system when many information producers work with a large number of documents containing lots of information. The main benefits of information reuse are obvious: decrease of information replication and avoidance of update anomalies. In order to achieve these benefits, content management systems should not only allow

information reuse at a coarse and fixed level of granularity, e.g., documents, but should be more flexible and allow dynamic determination of the appropriate granule of reuse.

- **Exploitation of Structural Information**

An XML document contains different kinds of structural information, such as markup elements and attributes used, hierarchies of markup elements, or document structure information optionally defined in a corresponding DTD. Also XLink contains this kind of information. Structural information is a kind of metadata that apparently is very useful for improving the search quality on document contents and enabling automatic content consistency checking (see next point).

- **Content Consistency Enforcement**

The physical and logical structure of documents defined in the DTD as well as links between documents need to be handled adequately by enforcing their consistency, e.g., avoiding links to non-existing resources. Corresponding structural information can effectively be used for consistency enforcement purposes.

- **Integrated Management of XML Documents and Other Resources**

XML documents are by nature interconnected with each other and other kinds of resources like various types of image files, a variety of word processing file types, etc. Therefore, functions for an integrated management of XML documents and other resource types are required.

- **Utilization of DB Techniques**

Due to the considerable amount of XML documents anticipated in future applications, approved DB technology, e.g., for secondary storage management, query processing, concurrency control, transaction management, recovery, etc., is likely to be exploited for efficient content management.

- **Dynamic Content Generation**

The delivery of static XML document contents is not sufficient for many application domains. Structured business data stored in DBS are often needed to be combined with XML data dynamically. Furthermore, dynamic generation of XML documents based on existing data is wanted. Thus, functions for accessing external data sources and document contents as well as for processing document templates have to be provided.

- **Repository Functionality Support**

In order to manage several states of XML contents and to support the overall development life cycle, basic repository functions such as version control, configuration management, collaboration, context management, and workflow control are required. This additional functionality completes an XML content management

system and leads to what we call an *XML content repository*.

3. XML Content Management in XCoP

In this section, we describe our approach for managing XML document contents. XCoP manages both the textual representation of XML document contents, called *textual contents*, and the corresponding structural information. Document type definitions (DTDs), documents as well as (logical) document parts belong to textual contents. For that purpose, a conceptual model and a set of operations that support the processing of textual contents and structural information are provided. Everything that is managed by XCoP is conceptually modeled as repository objects. They are further characterized by their attributes (e.g., owner, creation date, last modification date, etc.) and structurally organized into various repository object types such as resources, documents, fragments, elements, etc.

3.1. Managing Textual Contents

Although XML supports content reuse through external entities, it lacks functions for maintaining and ensuring the consistency of document parts belonging together. Indeed, physical structures in XML need to be maintained manually using entity declarations specified in DTDs. Changes in the physical structure of a document require manual updates in the DTD. In contrast, XCoP permits users to logically split document contents into parts without having to change DTDs, maintains these parts and ensures consistency automatically. Moreover, splitting document contents enables collaborative work since users can work on their respective parts independently.

Definition of Fragment

Comparable with the terminology used in the XML Fragment Interchange [18], we introduce the concept *fragment*. A fragment consists of textual contents and may further comprise a set of other fragments (child fragments) but not itself. A fragment is well-formed, if its textual contents match the production content of XML 1.0 and all its child fragments are well-formed. Apparently, the *overall textual contents* of a well-formed fragment resulting from a proper integration of the textual contents of all direct and transitive child fragments is well-formed. A well-formed fragment is called a *root fragment*, if it has exactly one markup element at the top of its element hierarchy. A DTD fragment contains a set of markup declarations which are specified properly as defined in the XML specification and may also be composed of other DTD fragments. A DTD fragment is called a *DTD root fragment*, if it contains text delimited by brackets (“[...]”) as demanded by the production doctypedec1 of XML 1.0. Consequently, we define an XML document being well-formed, if it consists of

exactly one root fragment, and being valid, if it has a DTD (with only one DTD root fragment) and exactly one root fragment whose textual contents meet the DTD.

Based on the fragment concept, we can flexibly determine portions of XML document contents as well as DTDs, so that they can be reused for building new (DTD) fragments. Figure 1 illustrates an example of fragment reuse and the corresponding structuring relationships spanning a so-called *fragment graph*. Fragment graphs are acyclic and directed. Their graphical representation uses rectangles as vertices for fragments and arcs to describe the *aggregation* relationships between parent and child fragments. Arcs are ordered according to the appearance of the child fragments in their parent fragment.




Figure 1. An example of fragment reuse

Determining Fragment Granularity

One of the main strengths of XCoP is that it permits users to specify the fragments they want to reuse in arbitrary size and in a flexible way. We provide two simple techniques, named *explicit fragmentation* and *implicit fragmentation*, for determining (DTD) fragments by the help of XCoP-specific PIs (*processing instructions*), called *markup PIs*, that act like markup tags.

Explicit fragmentation happens by explicitly enclosing those parts of given document contents or fragments, which are supposed to become new fragments, by pairs of markup PIs especially provided for that purpose (`<?XCOP TAG=<EXPL-FRAGMENT>?>` and `<?XCOP TAG=</EXPL-FRAGMENT>?>`).

Implicit fragmentation happens by marking an excerpt of a DTD or DTD fragment. Corresponding markup PIs are `<?XCOP TAG=<IMPL-FRAGMENT>?>` and `<?XCOP TAG=</IMPL-FRAGMENT>?>`. The marked DTD excerpt may contain several element declarations. The meaning is that each element occurrence of an element declaration contained in the marked DTD excerpt is always to be treated as an own fragment.

Both techniques allow users to share and reuse fragments without having to take care of DTD updates. By default the entire contents of a document (without DTD) are treated as a fragment, if there are no fragmentation instructions. Thus, based on their needs users may freely deter-

```
<?XML VERSION="1.0"?>
<!DOCTYPE PRODUCT-CATALOG [
  <?XCOF TAG=<IMPL-FRAGMENT?>
    <!ELEMENT PRODUCT (NAME, (QA-NOTE | MFG-INSTR)*)>
    <?XCOF TAG=</IMPL-FRAGMENT?>
  ]>
  ...
<PRODUCT-CATALOG>
  <PRODUCT PROD-ID="PP-60" TYPE="raw-mat">
    <NAME> Polypropylene 60 </NAME>
  </PRODUCT>
  <PRODUCT PROD-ID="MIX-750" TYPE="sf-good">
    <NAME> FlexiMix </NAME>
    <?XCOF TAG=<EXPL-FRAGMENT?>
      <QA-NOTE>
        For european customers, assure mixing quality
        using <QA-METH>M-16</QA-METH>. Otherwise, use
        <QA-METH>M-30</QA-METH> with standard params.
      </QA-NOTE>
      <?XCOF TAG=</EXPL-FRAGMENT?>
    </PRODUCT>
  </PRODUCT-CATALOG>
```

Conceptual Model for Managing Textual Contents

```

classDiagram
    class RepositoryObject
    class Resource
    class BinaryResource
    class TextualResource
    class DTD
    class Document
    class Fragment
    class DTD_Fragment["DTD Fragment"]

    RepositoryObject <|-- Resource
    Resource <|-- BinaryResource
    Resource <|-- TextualResource
    TextualResource <|-- DTD
    TextualResource <|-- Document
    TextualResource <|-- Fragment
    DTD "0..*" o-- "0..1" Resource
    Document "0..*" o-- "0..*" Resource
    Fragment "1" o-- "0..*" Resource
    DTD_Fragment "1" o-- "0..*" DTD
    DTD_Fragment "1" o-- "0..*" Fragment
    Fragment "0..*" o-- "0..*" Parent
    Parent "0..*" o-- "0..*" Child["Child {ordered}"]
  
```

The diagram illustrates the relationships between various classes in a system. At the top is the **Repository Object** class, which is a generalization of the **Resource** class. The **Resource** class is further specialized into **Binary Resource** and **Textual Resource**. The **Textual Resource** class is specialized into **DTD**, **Document**, and **Fragment**. The **Fragment** class is further specialized into **DTD Fragment**. The diagram also shows associations between these classes, including multiplicity and role names like **Parent** and **Child {ordered}**.

represented by a repository object which can be identified by a URI (Uniform Resource Identifier, [4]). Resources are either binary resources (e.g., images or executable files) or textual resources (e.g., ASCII text files). The latter are refined in DTDs, documents, and fragments. A fragment may consist of several child fragments, which are ordered and may be shared with other parent fragments. Fragments include markup elements and are further specialized in DTD fragments that consist of markup declarations.

Currently, we use the following kinds of structural information in order to support automatic enforcement of document contents consistency and to provide appropriate search facilities:

- ### Conceptual Model for Managing Structural Information

As presented in Figure 4, markup declarations, markup elements, and attributes are also modeled as repository objects. Markup declarations are further classified into element declarations, attribute declarations, and entity declarations.

rations, whereas an element declaration may contain one or more attribute declaration(s) and consist of an ordered set of other element declarations. DTD fragments comprise an ordered and non-empty set of markup declarations. Due to simplicity, we ignore parameter entities and general entities as well as their relationships to entity declarations in the models given in this paper.




Figure 4. Modeling structural information

3.3. Processing Repository Objects

This section gives an overview of the operations provided to support the processing of repository objects, particularly those that are related to textual contents and structural information. Operations are defined specifically for each repository object type. Although our actual modeling approach is object-oriented, we give a more generic description in the following in order to keep our explanations short and simple. We distinguish four categories: manipulation operations on repository objects, manipulation operations on relationships, retrieval operations, and specific operations. Well-formedness respectively validity constraints are checked during operation executions whenever required, especially for such operations that effect content changes.

Manipulation Operations on Repository Objects

This category contains operations for the creation, modification, and deletion of repository objects.

- `createObject (initTextualContents: String): Object`
This operation accepts an input string representing the

initial textual contents of the repository object to be created. After its successful execution, the created repository object is delivered. When creating a document object, the corresponding root fragment and DTD (if any) are identified and created (again by calling `createObject`) and associated with the document object (see `CreateRelationship` operation). Likewise, a create operation for a DTD leads to the creation of the related DTD root fragment and the corresponding relationship. In case of the creation of a (DTD) fragment object, fragmentation markups within the textual contents of the fragment are recognized and the `createObject` operation is recursively performed for each child fragment. Afterwards, the corresponding aggregation relationships are created. Similarly, structural information of a (DTD) fragment, i.e., markup declarations, markup elements and attributes, are extracted and the corresponding objects and relationships are created.

- `updateObject (newTextualContents: String): Boolean`
An update operation requires the new textual contents supposed to replace the existing textual contents of a repository object. Update operations are usually performed by calling deletion and creation operations on subordinated data structures.
- `deleteObject (): Boolean`
This operation deletes a repository object. An object can only be deleted if it does not share (aggregation) subordinated objects with other objects of the same type. Similar to the operations mentioned previously, a delete operation may cause cascaded deletions.

Manipulation Operations on Relationships

The operations of this category are used explicitly to establish and delete relationships between repository objects.

- `createRelationship (partnerObject: Object [, relType: RelationshipType, posExpr: PosExpr]): Boolean`
A relationship to another repository object (e.g., from a document to a DTD, from a parent to a child fragment, from a fragment to an element, etc.) is established using this operation. The optional parameter `posExpr` specifies the relationship position and is needed only in case of an ordered relationship type. For example, a position of a composition relationship between a parent and a child element is given by a non-negative integer number specifying the order position of the child element in its parent element, whereas an `XPointer` (XML Pointer Language, [9]) expression is used to specify a position of an aggregation relationship between a parent and a child fragment. If both participating repository objects are elements, then the relationship type `relType` to be

taken into account can additionally be restricted (composition or reference). In contrast, when creating a relationship between two elements, an element and an attribute, or two markup declarations, changes in the corresponding fragment are performed accordingly.

- `deleteRelationship(partnerObject: Object): Boolean`
This operation just deletes a relationship existing to the partner repository object. The corresponding objects remain in the repository.

Retrieval Operations

Retrieval operations allow us to get the textual contents of a repository object and provide a navigational access mechanism for repository objects based on their relationships to other repository objects.

- `getObjectContents([level: Integer]): String`
By this operation, the textual contents of a repository object are retrieved. If applied to a repository object which is not an attribute object, a non-negative integer number `level` is required representing the number of the hierarchy levels (of the same object type) to be examined. This operation delivers the textual contents which result from integrating the textual contents of all child objects up to the specified level.
- `getObjects(targetObjectType: ObjectType [, relType: RelationshipType, targetObjectRole: ObjectRole, level: Integer]): ObjectList`
This operation delivers a list of objects of type `targetObjectType` that are connected to a repository object. A non-negative integer number `level` representing the number of hierarchy levels to be examined and the role of the target repository object `targetObjectRole` (i.e., child or parent) can optionally be specified if origin and target object types are the same. If both are elements, then the relationship type to be taken into account can additionally be restricted (composition or reference).

Specific Operations

Unlike the operations of the previous three categories, specific operations are exclusively applicable to certain repository object types. In the following, the specific operations on fragments are explained:

- `splitFragment(targetRangeExpr: XPtrExpr): Fragment`
A (textual) excerpt (specified by the XPointer expression `targetRangeExpr`) of a fragment object is extracted to become a new child fragment object.
- `mergeFragment(role: FragmentRole, allowReplication: Boolean): Fragment`
This is the reverse operation to the previous one.

Depending on the role (i.e., child or parent) of a fragment object, its textual contents are merged with its parent or child fragment(s). If shared fragments are involved it depends on the flag `allowReplication`, whether these shared fragments are replicated or the merge is denied.

- `matchFragment(targetExpr: XPtrExpr): Boolean`
This operations checks whether the textual contents of a fragment object contain parts that match the text specified by `targetExpr`.

4. Implementation

In this section, we give an overview of iWebDB, our integrated Web database framework [21], and briefly describe how it can be used to implement XCoP. We want to start with a short introduction to object-relational database (ORDB) technology, our implementations are based on.

4.1. Useful ORDB Technology

Current activities in developing ORDBMSs [25] aim at integrating object-oriented concepts into the relational model. The *extensibility* property of ORDBMS can be considered to be extremely beneficial for efficiently implementing repository functionality, since it allows for:

- adding UDTs consisting of specially designed data structures and corresponding operations (UDRs) to the database schema;
- capturing arbitrary formats within UDTs by exploiting large objects (LOBs) types;
- linking externally stored data with database entries and controlling access to these (externally stored) data by database mechanisms;
- exploiting pre-defined extensions for accessing the database via the web, for managing data of various formats (text, HTML, video, audio, image, etc.);
- dynamically transforming externally stored data into relations, which may be incorporated into SQL processing.

The availability of these features was the reason, why we decided for ORDBMS. In the following, it will be outlined how we used these features.

4.2 iWebDB

iWebDB is implemented based on an ORDBMS and exploits OR extensibility features such as UDT, UDR and external data access. As depicted in Figure 5, iWebDB currently consists of eight modules, five of them being DBS extensions, so-called DataBlades (*Informix*), *Extenders* (*IBM*), or *Catridges* (*Oracle*): *Doc*, *ED* (External Data), *eXtract*, *DG* (Document Generator), and *XCoM* (XML Content Manager). In addition, two client applications, *SM*

(Site Manager) and XCoEx (XML Content Explorer), exist. The module MP (Macro Processor) is a function library that can be used for client-side as well as for server-side applications. XCoM and XCoEx are specific components of XCoP.




Figure 5. Architecture of iWebDB

- **iWebDB/Doc**

All base types and tables for WCM, e. g., types for storing XML, HTML, Postscript, and other text documents as well as images of different formats are provided by iWebDB/Doc. Each document is stored entirely. In addition to document storage, iWebDB/Doc provides functionality for managing, querying, and analyzing the (HTML) contents, e. g., functions to extract the title, headers, hyperlinks, etc.

- **iWebDB/ED**

This module helps to simplify Web site administration by allowing the integration of externally stored files into iWebDB. It is based on the abstract table facility of SQL/MED (Management of External Data, [20]). File system data accessible via Web protocols, such as "file:", "ftp:", and "http:", can be made available in relational tables, the file contents in document formats provided by iWebDB/Doc. Available data can be queried using SQL. Thus, external data can seamlessly be integrated into the system. Users and other iWebDB components can access documents stored in a file system as well as documents managed by the DBS in the same way by using SQL. iWebDB/ED is used by the document generator module to write generated documents to the file system and, in case of document templates provided in the file system, to read them.

- **iWebDB/eXtract**

The search capabilities provided by iWebDB/Doc are enhanced by iWebDB/eXtract. Furthermore, function-

ality is provided for storing extracted data in a special index structure being exploited by a search engine.

- **iWebDB/SM**

This module is the iWebDB tool for site management needed by all information providers. Based on a graphical user interface (GUI), all tasks concerning installation and maintenance of Web documents and directories, user and group administration as well as enable/disable document generation can easily be performed.

- **iWebDB/MP**

This module is a function library designed to be the foundation for building gateway programs for Web-based DB access. It offers an extensible macro processor for embedding special tags into Web documents. Among other things, it provides functionality for the execution of SQL statements, invocation of Java functions, and the composition of documents. The library can be used to realize a fully functional client- or server-side Web database gateway (see Figure 5). In addition, it serves as a basis for iWebDB/DG (see below). The macro processor offers an API to access the context manager, so that namespaces and variables can be created as well as values be read or set. Using this API, applications based on iWebDB/MP can create their own namespace, define variables, and set values.

- **iWebDB/DG**

This module provides a document generator based on iWebDB/MP. Using SQL triggers, static Web pages stored in the Web server file system are automatically refreshed whenever DB data the pages are based on has been modified. Thus, except for Web applications, Web information systems can avoid dynamically generated documents to offer up-to-date information.

4.3. XCoP

After the overview of the iWebDB architecture and most of its components, we want now outline the two XCoP-specific components, XCoEx and XCoM:

- **XCoEx**

This module provides a GUI for manipulating and retrieving repository objects. A text editor for editing textual contents is included. In addition, administration tasks related to repository objects and their mapping to the (web server) file system as well as user and group access privileges are also supported. XCoEx is implemented as a Java application using JDBC (Java Database Connectivity).

- **XCoM**

This module is responsible for the server-side management of repository objects. The implementation of the conceptual models and the corresponding operations on

the top of an ORDBMS is realized in XCoM.

XCoM will be detailed in the following subsections by showing how the conceptual models and operations on repository objects can easily be mapped into an OR data model based on IDS/UDO (Informix Dynamic Server with Universal Data Option, [19]). Current limitations of IDS/UDO related to modeling capabilities are also mentioned.

Mapping XCoP's Conceptual Model to OR Data Model

The repository object types (as described in Section 3) are captured within an ORDB schema as follows (see Figure 6 for examples). Each object type is mapped to a so-called *named row type*, and each subtype relationship to a corresponding *sub-class* relationship type (see keyword UNDER in Figure 6(a)). Each subtype inherits the attributes (i.e., data fields) as well as operational properties (i.e., routines, aggregates, and operators) of its supertype. Based on the defined type hierarchy, a table hierarchy can be specified (also illustrated in Figure 6(a)). The primary key constraint defined in the supertable repository_object is inherited by all subtables.

With IDS/UDO, *routine overloading* can be accomplished as shown in Figure 6(b). Routine overloading allows users to assign a single name (e.g., create_object) to multiple routines having different signatures.

IDS/UDO does not support reference types and provides only limited set operation capability on collection types. Because both techniques can be very useful for the implementation of relationship type semantics (e.g., aggregation), we have to employ referential integrity constraints instead as a workaround (see Figure 6(c)).

Considering that a fragment represents the smallest unit of information specified by users and indicates data locality, it is advantageous to also store it as a unit. Fragments are therefore stored as rows of a single table fragment. Obviously, this storage approach achieves better performance when constructing a document from its fragments as compared to other approaches that compose a document from its elements. Moreover, concurrent work on fragments can be directly controlled by using the row locking mechanism offered by most (O)RDBMS vendors. Aggregation relationships between fragments are mapped to a separate table (fragment_graph). The order_pos column in the fragment_graph table reflects the order position of a child fragment within its parent fragment.

Storing structural information, especially elements and attributes, within relational tables is a delicate matter as reported in [16, 24, 11]. Based on the mapping schemes given in [16], we use the *attribute approach* for storing structural information. Moreover, [16] describes techniques for translating queries posed in an XML query language (e.g., XML-QL [10]) into SQL queries for the alternative mapping schemes. Indeed, search quality on document contents

is improved significantly by exploiting their structural information which is extracted and stored in tables. Furthermore, dependencies between a fragment and its structural information are automatically controlled by triggers. Changes in a fragment will activate a trigger which accordingly maintains the corresponding structural information, and vice versa.

```
CREATE ROW TYPE repository_object_t (
  oid VARCHAR(8) NOT NULL,
  owner VARCHAR(20), cdate DATE);
CREATE ROW TYPE resource_t (
  uri VARCHAR(255), contents BLOB)
  UNDER repository_object_t;
CREATE ROW TYPE textual_resource_t (
  description VARCHAR(255)) UNDER resource_t;
CREATE ROW TYPE document_t (
  content_status VARCHAR(8), dtd_oid VARCHAR(8),
  root_fragment_oid VARCHAR(8))
  UNDER textual_resource_t;
CREATE ROW TYPE dtd_t (
  name VARCHAR(60), root_fragment_oid VARCHAR(8))
  UNDER textual_resource_t;
CREATE ROW TYPE fragment_t (is_root BOOLEAN)
  UNDER textual_resource_t;
CREATE ROW TYPE dtd_fragment_t ()
  UNDER fragment_t;
```

```
CREATE TABLE repository_object
  OF TYPE repository_object_t (PRIMARY KEY (oid));
CREATE TABLE resource
  OF TYPE resource_t UNDER repository_object;
CREATE TABLE textual_resource
  OF TYPE textual_resource_t UNDER resource;
CREATE TABLE document
  OF TYPE document_t UNDER textual_resource;
CREATE TABLE dtd OF TYPE dtd_t
  UNDER textual_resource;
CREATE TABLE fragment
  OF TYPE fragment_t UNDER textual_resource;
CREATE TABLE dtd_fragment
  OF TYPE dtd_fragment_t UNDER fragment;
```

(a) Definition of type hierarchy and typed table hierarchy

```
CREATE FUNCTION create_object (document)
  RETURNING VARCHAR(8);
CREATE FUNCTION create_object (dtd)
  RETURNING VARCHAR(8);
CREATE FUNCTION create_object (fragment)
  RETURNING VARCHAR(8);
```

(b) Routine overloading

```
ALTER TABLE document ADD CONSTRAINT (
  FOREIGN KEY (dtd_oid) REFERENCES dtd (oid),
  FOREIGN KEY (root_fragment_oid)
  REFERENCES fragment (oid));

ALTER TABLE dtd ADD CONSTRAINT (
  FOREIGN KEY (root_fragment_oid)
  REFERENCES dtd_fragment (oid));

CREATE TABLE fragment_graph (
  parent_oid VARCHAR(8), child_oid (8),
  order_pos SMALLINT,
  PRIMARY KEY (parent_oid, child_oid),
  FOREIGN KEY (parent_oid)
  REFERENCES fragment (oid),
  FOREIGN KEY (child_oid)
  REFERENCES fragment (oid));
```

(c) Implementing relationships between repository objects

Figure 6. A sample ORDB schema

Other extensions of iWebDB, particularly MP, DG, and ED, allow dynamic content generation of documents (and fragments) integrating existing legacy databases or external files by using special macro tags and templates. In our

first prototype, we store textual contents as LOBs, since IDS/UDO does not offer mechanisms comparable to the *IBM DataLinks* [7] allowing to handle links to externally stored files. Finally, we again utilize referential integrity constraints and triggers for automatic enforcement of link and reference consistency.

5. Related Work

With XCoP we contribute an integrated solution approach to content management requirements of XML documents as presented in Section 2. Unlike other tools and approaches that are either based on an RDBMS or an object-oriented database management system (OODBMS), the implementation of our approach is based on an ORDBMS. By exploiting the extension capabilities of an ORDBMS many tasks concerning content management can be delegated to the DBMS in an integrative way.

While XCoP supports content reuse extensively, iWeb-DB/Doc [21] only offers basic storage options at document level. An implicit fragmentation technique for structured document contents was introduced in HyperStorM [1], a DB application framework designed for managing SGML (Standard Generalized Markup Language) documents and built on the top of an OODBMS. We have adapted this technique in order to apply it to XML documents in XCoP. Beyond the possibilities provided in HyperStorM, we allow users to modify implicit fragmentation specifications at run-time and additionally offer the explicit fragmentation method. HyperStorM does not exploit structural information of fragments and lacks content management functionality, e.g., content reuse, external data sources integration, etc. Recently, Lore [17], a DBMS specifically designed to handle semistructured data, has been modified in order to manage XML data. Lore focuses on processing fine-grained structural information rather than coarse-grained textual contents. It also lacks sophisticated content management functionality.

Fragmentation of XML documents has been proposed to W3C in XML Fragment Interchange [18]. This proposal only specifies a mechanism for processing parts of an XML document without having to take the entire document into account. Issues concerning fragment reuse and collaboration among concurrently working users are beyond its scope.

There are a number of commercial products for DM as well as DCM, e.g., *DynaBase* [14] of *Inso Corporation*, *Documentum 4i* [12] of *Documentum*, *BladeRunner* [3] of *the e-content company*, and *POET CMS* [23] of *POET Software*. All products use a full-text search engine but offer limited search capabilities w.r.t. structural information. Furthermore these systems support versioning and configuration management, collaborative works, and workflow functionality. Integration with external data sources resid-

ing in file systems is not supported by all products. Most of them rely on OODBMS except for *DynaBase* and *Documentum* that use RDBMS. *DynaBase* and *Documentum 4i* are DM products and thus use files as smallest content units, while others enable content reuse in variable size. Finally, *BladeRunner* and *POET CMS* do not support dynamic content generation.

Recent research approaches [16, 24, 11] propose techniques for storing fine-grained XML data in an RDBMS in order to leverage relational technology. We apply two techniques described in [16] for storing both fine-grained structural information and textual contents in variable granularity, while commercial products such as *eXcelon* [15] of *eXcelon Corp.* and *Tamino* [26] of *Software AG* exclusively support processing of fine-grained XML. This, indeed, may lead to unnecessary fragmentations, complicates the content management of XML documents and, therefore, does not support content reuse adequately.

6. Conclusion and Outlook to Future Work

In this paper, we have presented XCoP, a repository approach designed for the content management of XML documents. XCoP integrates many content management services by utilizing ORDB technology. To the best of our knowledge XCoP is the first approach exploiting ORDB technology for these purposes. Unlike existing products and approaches, processing both textual contents as well as corresponding structural information of XML documents are supported. XCoP facilitates content reuse through fragmentation which obviously avoids content redundancy and enables collaborative work on documents. With our fragmentation methods users can logically decompose an XML document into fragments and dynamically determine the granularity of the units they want to share. Two fragmentation methods have been introduced: implicit and explicit fragmentation. An implicit fragmentation is applied to markup declarations of a DTD and an explicit fragmentation to document contents. Furthermore, conceptual models for the management of textual contents and structural information have been presented and corresponding manipulation operations have been introduced. Finally, implementation techniques of XCoP based on the ORDBMS IDS/UDO with its current limitations have been discussed. The mapping of the conceptual models to an ORDB schema has been exemplary presented and ORDBMS features used, e.g., UDT, UDR, constraints, locking mechanism, and trigger, as well as those that are desirable, i.e., reference types, set operation support on collection types, and links to externally stored files, have been discussed.

Currently, we are embedding structural information contained in RDF and working on version control, configuration management, and collaboration support. In contrast to WebDAV [6], a protocol for distributed authoring and

versioning applying at resource level (i.e., documents), our approach aims at controlling the management of document fragments. Furthermore, context management and workflow control functionality are planned to be accommodated in XCoP.

Acknowledgments

The authors thank T. Härder, M. Flehmig, U. Marder, and H.-P. Steiert for fruitful discussions and helpful comments on earlier versions of this paper.

7. References

- [1] Böhm, K., Aberer, K., Neuhold, E.J., Yang, X., *Structured Document Storage and Refined Declarative and navigational Access Mechanisms in HyperStorM*, The VLDB Journal, Vol. 6, 1997, pp. 296-311.
- [2] Bernstein, P.A., Dayal, U., *An Overview of Repository Technology*, Proc. of 20th VLDB, Santiago, Chile, September 1994, pp. 705-713.
- [3] *BladeRunner 1.5*, the e-content company, 1999.
<http://www.xmlcontent.com/>.
- [4] Berners-Lee, T., Fielding, R., Irvine, U.C., Masinter, L., *Uniform Resource Identifiers (URI) - Generic Syntax*, Network Working Group RFC 2396, IETF, August 1998.
<http://www.ietf.org/rfc/rfc2396.txt>.
- [5] Bray, T., Paoli, J., Sperberg-McQueen, C.M., *Extensible Markup Language (XML) 1.0*, W3C Recommendation, February 1998.
<http://www.w3.org/TR/REC-xml-19980210>.
- [6] Clark, J.R., Kaler, C., *Versioning Extensions to WebDAV*, Internet Draft, IETF Delta-V Working Group, February 2000.
<http://www.webdav.org/deltav/protocol/>.
- [7] Davis, J.R., *DataLinks: Managing External Data with DB2 Universal Database*, IBM Corporation, February 1999.
<http://www-4.ibm.com/software/data/pubs/papers/#datalink>.
- [8] Deach, S., Adler, S., Berglund, A., Caruso, J., Milowski, A., Zilles, S., *Extensible Stylesheet Language (XSL) Specification*, W3C Working Draft, April 1999.
<http://www.w3.org/TR/WD-xsl>.
- [9] DeRose, S., Daniel Jr., R., Maler, E., *XML Pointer Language (XPointer)*, W3C Working Draft, December 1999.
<http://www.w3.org/TR/WD-xptr>.
- [10] Deutsch, A., Fernandez, M., Florescu, D., Levy, A., Suciu, D., *XML-QL: A Query Language for XML*, W3C Note, August 1998. <http://www.w3.org/TR/1998/NOTE-xml-ql-19980819/>.
- [11] Deutsch, A., Fernandez, M., Suciu, D., *Storing Semistructured Data with STORED*, SIGMOD'99, Philadelphia, 1999, pp. 431-442.
- [12] *Documentum 4i*, Documentum, 1999.
<http://www.documentum.com/>.
- [13] DeRose, S., Orchard, D., Trafford, B., *XML Linking Language (XLink)*, W3C Working Draft, December 1999.
<http://www.w3.org/TR/xlink>.
- [14] *DynaBase*, eBusiness Technologies, Inso Corporation, 1999.
<http://www.ebt.com/>.
- [15] *eXcelon*, eXcelon Corporation, 1999.
<http://www.odi.com/>.
- [16] Florescu, D., Kossmann, D., *A Performance Evaluation of Alternative Mapping Schemes for Storing XML Data in a Relational Database*, Technical Report No. 3680, INRIA, France, 1999.
- [17] Goldman, R., McHugh, J., Widom, J., *From Semistructured Data to XML: Migrating the Lore Data Model and Query Language*, Proc. of WebDB'99, Philadelphia, Pennsylvania, June 1999, pp. 25-30.
- [18] Grosso, P., Veillard, D., *XML Fragment Interchange*, W3C Working Draft, June 1999.
<http://www.w3.org/TR/WD-xml-fragment>.
- [19] *Informix Dynamic Server with Universal Data Option 9.1.X*, Informix Software, Inc., 1997.
<http://www.informix.com/>.
- [20] ISO/IEC JTC1/SC32, *Information Technology - Database Language SQL - Part 9: Management of External Data*, ISO, 1999.
- [21] Loeser, H., Ritter, N., *iWebDB - Integrated Web Content Management based on Object-Relational Database Technology*, Proc. of IDEAS'99, Montreal, Canada, August 1999, pp. 92-97.
- [22] Lassila, O., Swick, R.R., *Resource Description Framework (RDF) Model and Syntax Specification*, W3C Recommendation, February 1999.
<http://www.w3.org/TR/REC-rdf-syntax>.
- [23] *POET Content Management Suite 2.0*, POET Software, 1999. <http://www.poet.com/>.
- [24] Shanmugasundaram, J., Tufte, K., He, G., Zhang, C., De Witt, D., Naughton, J., *Relational Databases for Querying XML Documents: Limitations and Opportunities*, Proc. of 25th VLDB, Edinburgh, Scotland, 1999, pp. 302-314.
- [25] Stonebraker, M., Brown, P., Moore, D., *Object-Relational DBMSs - Tracking the Next Great Wave*, 2nd edition, Morgan Kaufmann Publishers, 1998.
- [26] *Tamino*, Software AG, 1999.
<http://www.softwareag.com/>.