
Silvia Hollfelder, Florian Schmidt, Matthias Hemmje, Karl Aberer, Arnd Steinmetz

Transparent Integration of Continuous Media Support into a Multimedia DBMS

December 1997



**GMD – German National Research
Center for Information Technology**

The “Arbeitspapiere der GMD – GMD Technical Reports” primarily comprise preliminary publications, specific partial results and complementary material. In the interest of a subsequent final publication the “Arbeitspapiere/Technical Reports” should not be copied. Critical comments would be appreciated by the authors.

No part of this publication may be reproduced or further processed in any form or by any means without the prior permission of GMD. All rights reserved.

© Copyright



1997

Addresses of the authors:

Silvia Hollfelder, Matthias Hemmje, Karl Aberer, Arnd Steinmetz

GMD-IPSI

Dolivostraße 15

D-64293 Darmstadt

Phone: +49/6151/869-953/ 869-844/ 869-935/869-862

Fax: +49/6151/869-966

Email: {hollfeld, hemmje, aberer, arnd.steinmetz}@darmstadt.gmd.de

WWW: <http://www.darmstadt.gmd.de/{~hollfeld,~hemmje,~aberer, ~steinmetz}>

Address all correspondence to:

Silvia Hollfelder

Abstract*

Multimedia Database Management Systems (MM-DBMS) have to efficiently provide the specific functionalities required by time-dependent multimedia data types. During presentation playout, this requires support for continuous data delivery and media-oriented optimization of presentation quality according to user requirements. In this work, we show how we integrated this type of support transparently with other MM-DBMS functions like media editing and querying. This is achieved by extending an existing data model for continuous media, in particular by means of stream abstractions for editing and presentation, and by extending the system architecture by means of supporting stream processing components, e.g., for intra-media synchronization by adaptation. This approach has been implemented as an extension of existing commercial object-relational DBMS technology.

* This work has partly been funded by the European Union within the framework of the ESPRIT Long Term Research Project HERMES, N. 9141 [<http://www.ced.tuc.gr/hermes/hermes.html>].

1 Introduction

In commercial DBMS there is a clear trend to provide specialized functionality, which has usually been provided as separate software packages and required substantial development efforts for the integration with database applications, as integrated database extensions. Examples are particularly found in the management of nonstandard or multimedia data types, e.g., text retrieval, image handling, statistical functions. This trend is a key factor for the rapid spreading of object-relational database technology among all major database vendors and is the reaction to the requirements of new applications exploiting distributed multimedia technology, like digital video production [S96][SH94], education and training or digital libraries [BAN95]. For static data and media, this approach is relatively straightforward as it does not heavily affect the fundamental processing paradigms used in conventional DBMS. For handling continuous media, like, e.g., video and audio, additional characteristics and requirements related to the presentational aspects of these media need to be considered. Currently, there are two “schools” to that respect (for a nice discussion see also [O96]):

1. Media presentation is not a problem that is to be considered by the DBMS. The DBMS provides a storage framework for the media, including meta data required for search and manipulation. Presentation is supported by external systems, like video servers, or within distributed multimedia applications.
2. The DBMS has to support presentational aspects, since this is a fundamental aspect of accessing the media data managed by the DBMS. Presentational requirements like synchronization are to be considered as a new type of constraints that need to be consistently supported by the DBMS [MS96].

From our perspective there are several arguments in favor of the second school:

1. Presentations of composite multimedia objects require the coordinated scheduling of the available resources based on the characteristics of and relationships between the involved media. Since this information is stored in the database the DBMS itself most efficiently performs the scheduling. For example, when presenting a video with text inserts the DBMS needs to support a buffering scheme for the video that exploits the stream semantics and expected user interactions. In addition, it has to schedule the delivery of the text inserts such that they can be presented timely.
2. In case of concurring requests, the DBMS is the place that can optimize the database usage and usage of available different media, for example by performing quality adaptations [TKCWP97].
3. Different abstractions are used in the different stages of processing of media data. For example, a MPEG video stream is edited at the frame granularity, retrieved at the granularity of scenes or shots, and distributed over a network for presentation at the granularity group-of-pictures. The different data structures representing these abstractions, which are also relevant for presentation purposes, need to be kept consistent during updates.
4. Exploiting application-specific characteristics of multimedia objects for distribution and presentation of the media increases the necessity of integrated presentation support in the DBMS. Examples are quality adaptations for parts of presentations resulting from a retrieval request parametrized by their relevance, buffering strategies taking into account possible future interactions based on application semantics, e.g., favorite hyperlinks, or resource scheduling taking user profiles or usage statistics into account.

In this paper, we show the feasibility of the integration of presentational support for multimedia data with a standard commercial object-relational DBMS (Informix Universal Server IUS). Basic support in the spirit of school 1 is provided by the IUS by means of the Video Foundation

Blade [VFUG97]. We actually take advantage of the mechanisms provided by this database extension to simplify the physical integration of the media server and obtain additional abstraction and querying capabilities. The additional presentational support that we provide addresses the continuous data transport required during the presentation of continuous media streams, client-side buffering mechanisms exploiting stream and interaction semantics, and client-side quality adaptation mechanisms to overcome fluctuations in the network and server resources. Most of the concepts were already developed while we extended the object-oriented DBMS VODAK [V95] to a (proprietary) multimedia DBMS [RKN96]. This paper shows that the concepts for presentational support developed in this earlier proprietary work can be smoothly integrated into standard platforms, taking advantage of the available functionality of these systems. Thus there is no principal difficulty in a modular extension of modern, extensible DBMS with presentational support for multimedia data.

2 General Concepts

In the following, we go into some details of the basic concepts supporting the storage and presentation that have been integrated as multimedia database extensions for presentational support for multimedia objects.

2.1 Continuous Long Fields Data Abstractions

We provide a data type *Continuous Long Field* (CLF) as generic, abstract representation for any kind of continuous media, like audio, video or animations. This datatype supports operations for editing and presentation, like “insert”, “request”, or “delete”. By modelling the generic concept CLF, we combine the concepts of streams and long fields. While streams are only “consumed”, e.g., read by presentation entities, the concept of CLFs allows consumption as well as any access to the data. Structural meta data, like the format and other recording parameters, are stored with every CLF object. Multiple media encoded in a single stream (like, e.g., with MPEG-encoded audio and video) are modelled as separate CLF objects, to enable the individual retrieval and manipulation of each of the media parts. For editing purposes, CLF objects are segmented into a sequence of manipulation units, called Continuous Object Data Units (CODU). In consequence, a CLF is a sequence of CODUs. The granularity of these units is determined individually for each type of media, e.g. Motion-JPEG frames, MPEG-1 frames or audio samples. At this level of granularity, it is possible to manipulate the data stream, e.g., by inserting, deleting or appending CODUs of a Continuous Long Field.

For continuous presentation our approach supports a client-pull mechanisms, i.e., the client continuously requests chunks of data which are provided in our system by the server in a best-effort manner. Since requests at the granularity level of CODUs are not very effective and presentation engines, like, e.g., MPEG players or other video and audio devices, handle larger units of data, a second abstraction is introduced to support presentation. An atomic unit requested during synchronous presentation is called a Continuous Object Presentation Unit (COPU). One COPU may consist of several CODUs as illustrated in Figure 1. By this generic concept, the buffer manager is able to handle any kind of continuous media in the same way by means of generic COPU requests.

The CLF data type that provides the above abstractions has been implemented as part of our database extension for basic continuous media support. It provides, in particular, storage support for the raw media data and associated meta data, like, e.g., content descriptions, physical characteristics or presentational characteristics [HL97]. By means of the explicit representation of meta data entities, we offer the opportunity to support queries on meta data, too. This means it is, e.g., possible to send queries about the content, the structure or technical parameters of a video stream. Currently, the content-oriented meta data has to be managed manually in the database, because we do not support any automatic content-oriented analysis or retrieval mechanisms in the current implementation. More details on the implementational aspects will be provided in Section 3.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.