

Recent Developments for Data Models – Exercise 1

Monday, May 07, 2012 – 15:30 to 17:00 – Room 36-336

1) User-defined Structured Types and Typed Tables

In this exercise¹, a database to store publications of a research group is considered. The conceptual schema is depicted in Figure 1 using an UML class diagram.

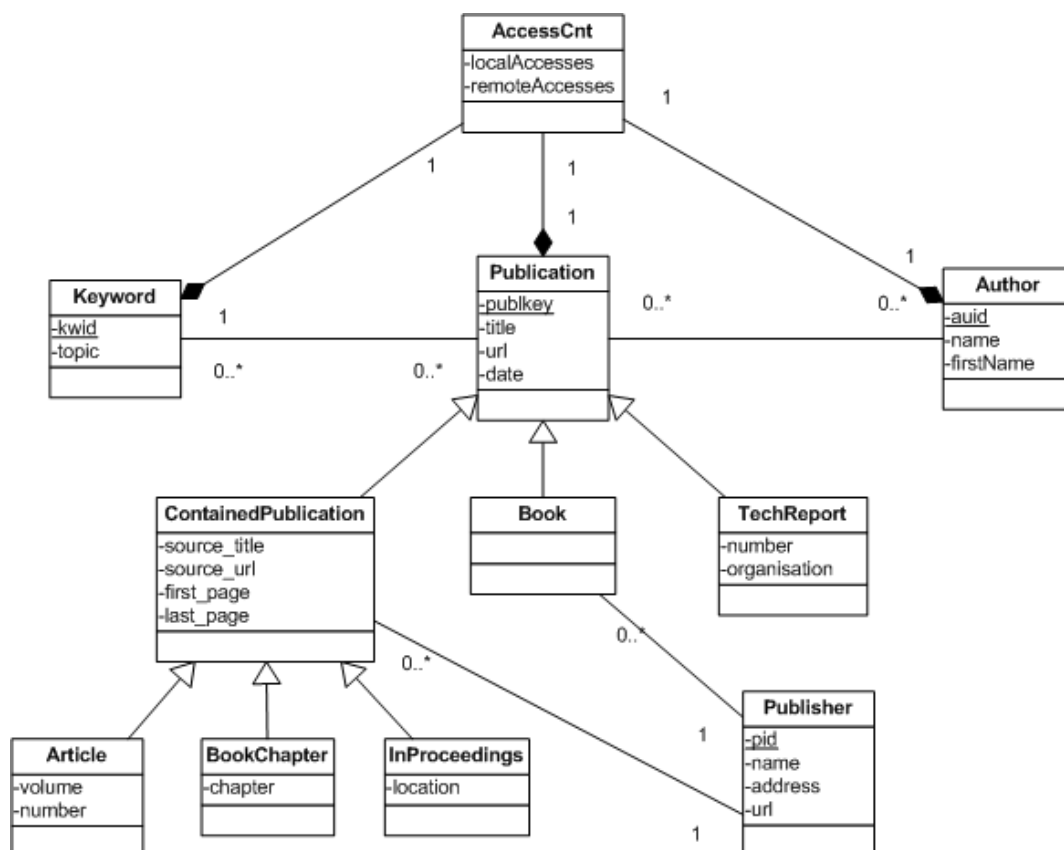


Figure 1: Conceptual schema of the publication database

Different types of publications are to be distinguished: Articles, book chapters, and proceeding papers are contained in collections of publications, whereas books and technical reports are not. Each publication is written by one or more authors. Authors may contribute to multiple publications. Books, articles, book chapters, and proceed-

¹ This exercise is based on *Andreas Geppert: Objektrelationale und objektorientierte Datenbankanzepte und -systeme, dpunkt.verlag, 2002.*

ings are published, whereas technical reports are made available in an informal manner, i.e. a publisher does not exist for this kind of publications. Keywords are used to suggest the subjects touched by publications. Several keywords may be assigned to a single publication and a single keyword may be assigned to multiple publications. The database keeps track of accesses to publications, keywords, and authors. Accesses from within the intranet are distinguished from remote accesses.

- a. Specify SQL DDL statements to create user-defined structured types for the entities given in the conceptual schema! What variations of object identity exist? Specify SQL Reference Types for each alternative!
- b. Specify SQL DDL statements to define typed tables based on the user-defined structured types! Furthermore specify SQL DDL statements to define any necessary (untyped) tables!

2) Querying Typed Tables

Specify SQL queries to retrieve the following information.

- a. The name and address of all publishers.
- b. All authors (name) whose first name starts with an “A”
- c. All articles (pubkey) ordered by year of publication.
- d. All authors (name) and the number of their publications.
- e. All authors (name) with the publications (title) they contributed to.
- f. The author with the highest number of publications.
- g. All authors (name) with the number of local accesses.
- h. The sum of remote accesses for each author (name).
- i. Pairs of authors (names) that have more than one joint publication.
- j. The name of the author with OID 0.
- k. All publications (without its subtypes).
- l. All articles and books (title).
- m. All contained publications (source_title, source_url, first_page, last_page, volume, number, chapter, location).
- n. All publications with a flag to indicate whether they are contained in a collection or not (title, flag).
- o. The highest number of joint publications of two authors.
- p. All publications of authors named “Smith” with a keyword that contains the word “object”.
- q. All publications with their keywords (if any, otherwise with null value).
- r. All publications (title) that have been more often accessed remotely than locally.
- s. Insert a new author named “John Doe” with OID 1.
- t. Increment the remote access counter of the publication with OID “XY2010a” by one.