Application Study of a Global Genealogy Database

Mikael Ronström
Ericsson Business Innovation AB
Mikael.Ronstrom@inn.ericsson.se

1 Introduction

This paper performs a study of the performance impact on central servers given a global genealogy network. In particular it studies the database part of such central servers to see what is possible with current technology and technology brought forth in a near future. It does also present an idea of special link objects to assist in providing references to genealogical results. By using this link object it becomes possible to immediately from a scanned image of a birth record link into that persons complete known genealogy. It is assumed that the process of storing genealogical information is semi-automatic with tools that assist the human users.

There are six groups of objects in the genealogy application. The first group comprises the historical records that have been scanned, either directly or from a microfilm. This group also contains refinements of those objects where computers have been used to make the images clearer, to increase the sharpness, brightness and the contrast, possibly also adding colours to make the text more visible. This means that there could be several versions of the scanned objects.

This represents an immense storage volume. Already today there exist millions of microfilms of historical records, each containing about a thousand images of A3-format. One such image needs around 1 MB of storage in compressed form. Already from the start this will therefore need petabytes of storage space. This will then increase as more and more records are added to this historical database.

1 Petabyte would need 30.000 disks with current technology. With a 60% increase of storage space per year it will be possible to store 1 Petabyte in 5.000 disks in 2005. In 2005 clustering technology of databases will have made a breakthrough making it possible to use a large set of small servers with approximately 10 disks per server. Thus approximately 500 such servers can store 1 Petabyte in 2005 at an approximate cost of 20 k$ * 500 = 10 M$ (including HW + SW). To provide 1 Petabyte today would require much larger servers and many more disks and would probably cost close to 200 M$. Thus the time when petabyte servers are economically feasible is quickly approaching.

The next class of objects are the catalog objects. These objects specify the historical records, and what they represent. This is where the researcher goes when he wants to see what historical records exist in the area of his ancestors.

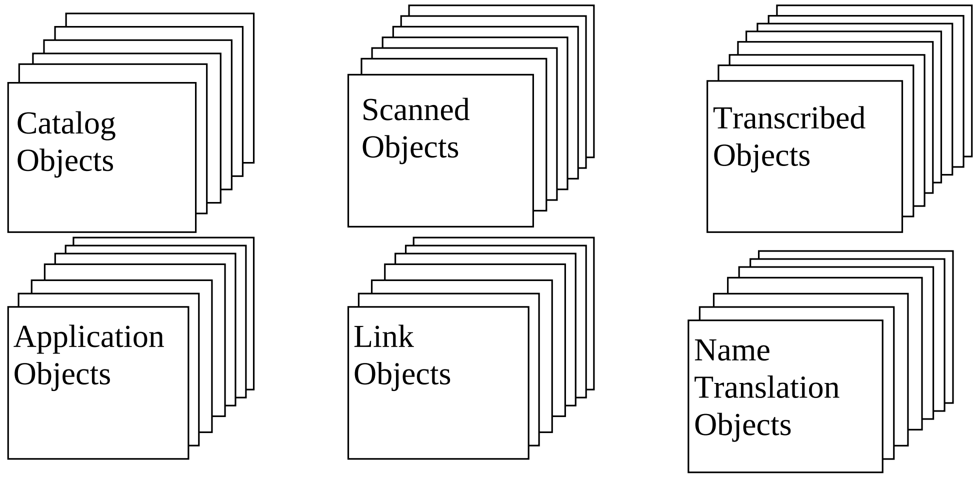

Figure 1-1 *Object Classes*

The third class of objects are transcriptions of the scanned objects. This means that a person or a computer program has read the text and made a note of what was written. There are also many versions of the transcribed objects. Each transcription can exist in many versions; first the original, where one simply tries to read character by character what is written; another version could be where older words have been translated into modern words; yet another version could be where the whole text has been translated into modern language and of course there could be translations to other languages as well. Also there could of course be many transcriptions of the same text. These need not be the same since mistakes could be made, especially when the text is difficult to read.

The fourth class of objects comprises the application objects which represent the information that is actually sought. In genealogy this means information about births, christenings, marriages, deaths and burials, as well as other information that can be used to link persons together into families. There could also be several versions of these objects, representing the findings of different genealogists. Of course the above information could be used for many other applications apart from genealogy. One such application is historical research.

The fifth class of objects comprises link objects which are used to link all the information together. A transcribed object needs to be linked to the original and possibly also to the version from which the transcription was made. Application objects needs to be linked to all sources (scanned objects as well as transcribed objects and even other application objects). The link object contains information about where on the pages that the information objects can be found as seen in Figure 1-1. The link object specifies which part of the page that is covered by the linked object. This could be specified by the x and y coordinates of the upper left corner and the length on the x-axis and the length of the y-axis.

The sixth type of object comprises name translation objects. These are necessary to provide translation to standardised names. In historical records it is not uncommon that the same name can be spelled in many ways. An example is Margareta that can be spelled in at least the following ways in Swedish records: Margareta, Margaretha, Marget, Margeth, Maret, Mareth, Mareta, Maretha, Margreta, Margretha, Marit, Marith, Märet, Märeth, Margit, Margith. All these spellings represent

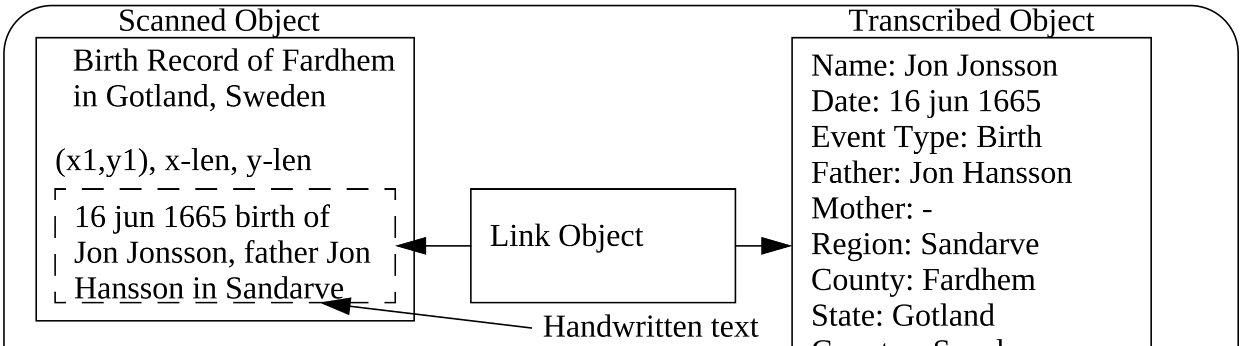

Figure 1-2 *Example of relations between Scanned Objects, Linked Objects and Transcribed Objects*

the same name and can even be used in different records to represent the same person. The different versions of names can differ in different geographical areas. Also provided in the name translation object could be information on the usage of that name, where it is common, history of the name and so forth.

1.1 Conceptual Schema

In Figure 1-3 a conceptual schema of the genealogy application is provided.

Figure 1-3 *Conceptual Schema of Genealogy Application*

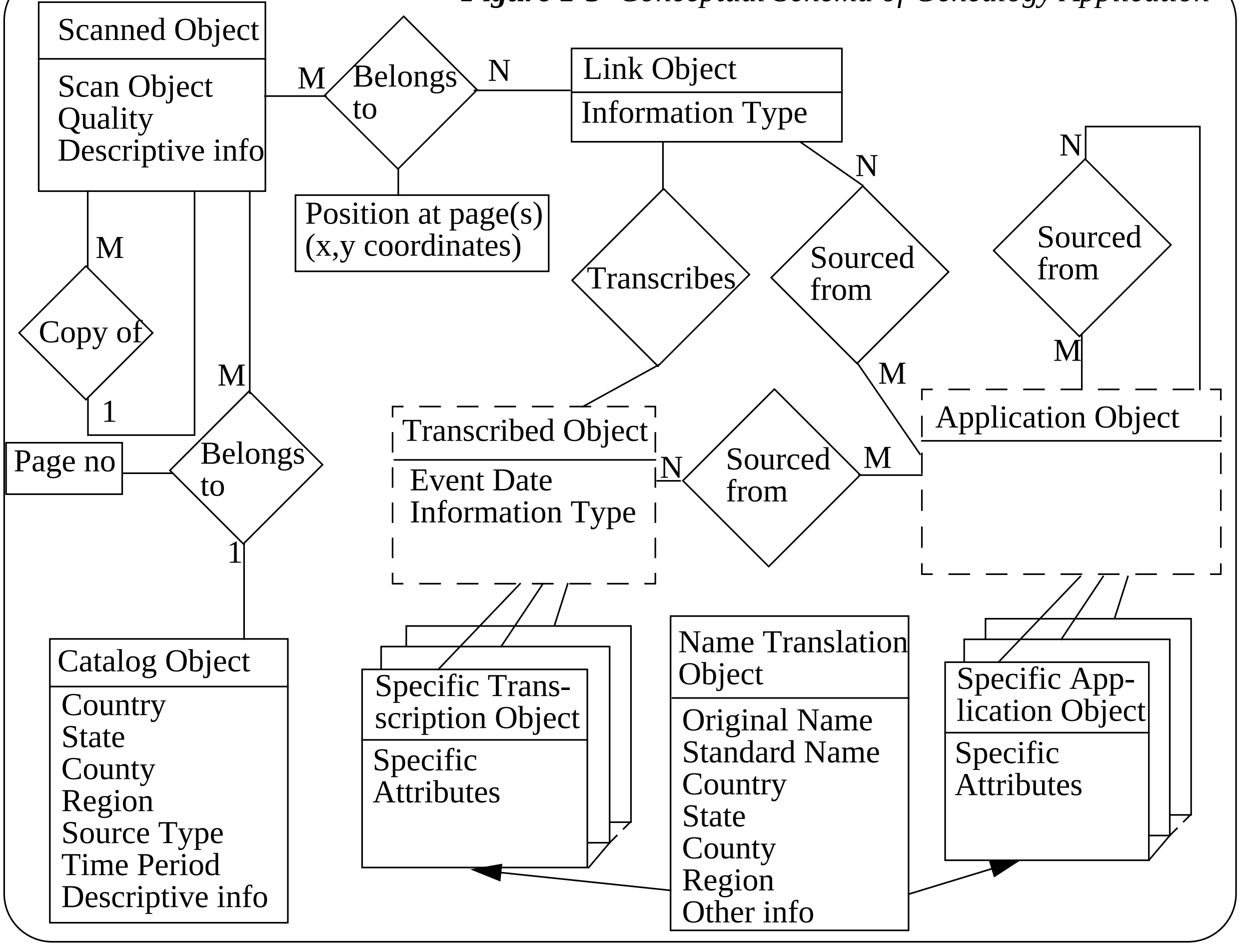

A more detailed description of the relations between the application objects are found in Figure 1-4. The basic application object is an object describing a person. A person has many relations; There are marriage relations, and there are relations to parents and their are relations to children. Other relations such as foster parents, adoptive parents and so forth are also possible. Names in transcribed objects and application objects are often given in the form that they were written in an original document. Many names with similar spellings represent the same name. Therefore there are name translation objects that translate the original names to a standardised name to simplify the search process.

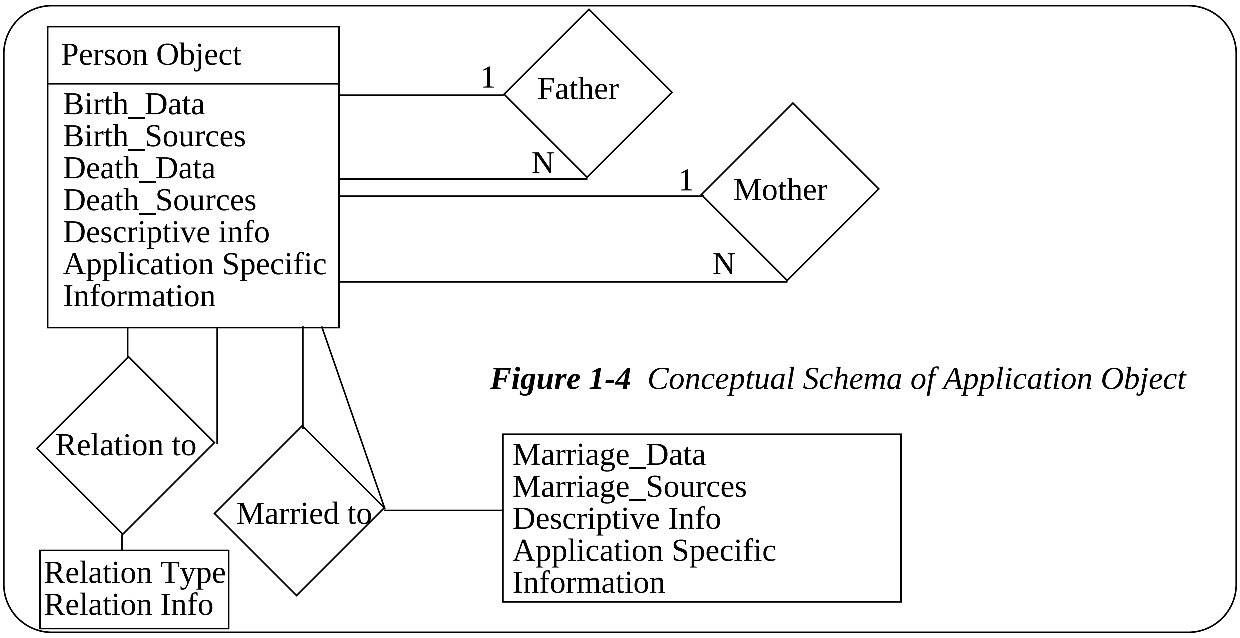

Figure 1-4 *Conceptual Schema of Application Object*

1.2 Traffic Model

To analyse the behaviour of this application we need to make guestimates on the traffic model. In this model we guestimate that the genealogy application is used by 50-100 million users around the globe. The model analyses the total requirements on all servers. In reality there will be a number of servers in central parts of the telecom network (e.g around 10-20 of them). We distinguish between five different activities. The first inserts scanned objects of approximately 1 MByte into the database, this could be the originals or refined versions. The second is when a genealogists is inserting application objects, transcribed objects or link objects. The third is used to search for information by scanning parts of the database. The fourth is used to provide links for objects. Finally one can read the scanned objects.

The traffic model is based on user behaviour and how many users that are active with different activities. Most likely very many people will be interested in browsing the data given that the data is accessible to all people. There is, however, a more limited set of people who are updating the data and providing new research results. These people use much more time than the people who only browse the data and thereby there will be a high percentage of the currently active who are updating the data.

***Table 1-1* Traffic Model**

	Insert Scanned Object	Insert/Update Object	Search Object	Get Links	Read Scanned Object
Events/sec	300	5.000	3.000	15.000	10.000
Read Object		10	50	10	
Write Object	30	10			
Scan Object			50.000		
Read Large Object					1
Write Large Object	1				

By using the traffic model we can derive the processor resources needed in the database part. In addition to the database processors one also needs a set of application processors (most likely a farm of simple web servers at about 1 k$ per server). The cost per object is based on measurements of NDB Cluster a database developed with clusters of small servers in mind and which uses new technology in an efficient way to enable such applications as global genealogy servers. The processor equivalent is an UltraSparc-III. It is assumed that most data will fit into main memory in 2005 with about 500 * 8 GByte = 4 TByte of total main memory in the server. The exception is the scanned objects.

***Table 1-2* Processor Cost Model**

	Cost per Object	Number of events/Sec	Processors
Read Object	50 µs	350.000	17.5
Write Object	100 µs	59.000	5.9
Scan Object	5 µs	15.000.000	75
Read Large Object (1 MByte)	50 ms	10.000	500
Write Large Object (1 MByte)	200 ms	300	60

In 2005 the processors will improve about a factor of 5. Thus the 500 small servers can easily handle the load. However the usage of hierarchical compression techniques is most needed to avoid that the servers need to send 10 GByte/sec through the telecom network. It should be possible to decrease this need to around 1 GByte/sec rather.

2 Conclusion

The conclusion is that there are two important factors which are needed to make global genealogy servers possible. The first is the improvement of disk technology which makes the number of disks needed smaller. The second is the development of small cheap servers combined with database software that enables to use clusters of small servers to build cost-efficient petabyte databases. It is likely that this has occurred in the next five years.