

EXPERIENCES IN USING IMS.2 FOR A CONSORTIUM OF LOCAL AUTHORITIES¹

A.J.O'Brien, Systems Manager, LOLA, (London On-Line Local Authorities)

1. Introduction

When, a year ago, we were approached by PATRAC and asked whether we would provide a paper for last year's Symposium, we declined on the grounds that we would prefer to have an operational on-line database system in existence before we accepted any invitations to major conferences. Happily, this year not only has our system celebrated its first birthday, but the Symposium concentrates on achievements rather than plans.

This paper deals with the path which led to the implementation of what we consider to be the most advanced on-line system in British local government.

2. Background to the LOLA Project

2.1. The original joint schemes

It is necessary to go back eleven years to find the beginning of the trail. At that time the Metropolitan Boroughs O & M Committee, as it was then called, was implementing an integrated batch system for a group of six local authorities in South-East London.

Four years later, after a major re-organisation of local government in London, two Boroughs in North-East London, Hackney & Tower Hamlets, commissioned the London Boroughs' Management Services Unit [LBMSU] (as it had become) to implement and operate a joint computer system.

The computer (a LEO III, roughly equivalent in power to an IBM 360/30 but with multi-programming capability — this in 1963!) was the same as that used by the South-East boroughs, and rapid implementation was achieved by using the standardised programs developed for the earlier consortium.

2.2. The Haringey Survey

In 1967 the London Borough of Haringey (who by now had joined Hackney and Tower Hamlets in the North-East London Computer Scheme in order to cope with their short-term computer requirements) commissioned the LBMSU to investigate the long-term data processing requirements of a local authority. The

investigation began in the autumn of 1967 and resulted in over 10 man-years of work. The end product was a series of reports covering the detailed working and computer requirements of each of the 18 departments, together with technical papers on various aspects of data processing. The Summary report (which was published in January 1969) was a 118-page book entitled "Report on the Initial Study: London Borough of Haringey Long-Term Computer Project"².

2.3. The Haringey Report

The basic concepts which the report identified were:

- (a) the requirement for a database for the local authority which would be consistent with departmental database and which would be of use to all the legitimate interests concerned.
- (b) the general need for increasing emphasis on and use of information as a resource in itself.

These two concepts implied a large computer equipped with mass storage devices and linked to all major departments through a telecommunications network.

To fulfil the centralising function, which is their major purpose, the databases would contain sufficient common items of information relevant to the work of more than one department together with cross-references to more extensive files maintained by individual departments. Furthermore, the system envisaged would not be restricted merely to providing facilities for enquiry, but would be based on the capability of updating records on-line.

These ideas are illustrated in Figure 1.

The proposed system as outlined in the report would have the ability to:-

- (a) interrogate or update records on any database at any time during the working day.
- (b) deal 'simultaneously' with messages from all remote terminals whatever database or record or activity was required.
- (c) accumulate 'queues' of work within the system and present these for action by a particular terminal user.
- (d) index all records of properties, organisations and individuals and thus permit access to such records to be gained by name or address as well as

¹ This document was scanned from a contemporary photocopy, OCR'ed, proofed and formatted similar to the original style, plus additional footnotes, in August 2016 by Alan E Cooper. Unclear and additional words shown in [].

² [Colloquially called the *Yellow Report* on account of its distinctive yellow cover. 2016 note]

by reference number.

- (e) maintain cross-references to non-computer files.
- (f) total or compile selective lists and tables of records meeting criteria defined by terminal users on an 'ad hoc' basis. Such results would be compiled and printed overnight.
- (g) perform normal batch processing runs either separately from, or concurrently with, terminal transactions.
- (h) store an indication of departmental interest in one or more activities related to selected records and to produce for the appropriate department details of any of the stated activities which occur on these records.
- (i) suppress access by unauthorised users to all or part of any database.

The advantages of such a system can be summarised as:-

- (a) standardisation of recording and referencing in those most critical and public areas of information – people and property – within a logical framework and with proper cross-indexing.
- (b) a similar logical framework for internal records to extend the scope of integrated data processing.
- (c) better communications and co-operation between departments, providing a basis for a more efficient and consistent service to the public.
- (d) less effort in acquiring, transmitting, copying, re-organising, extracting and analysing information.
- (e) laying the foundations of a progressive system of reporting on operational functions for managerial and strategic requirements.

2.4. The Growth of the Consortium

The concepts embodied in the Report were readily acceptable to Haringey's Council and Chief Officers. The difficulty was that the scale of investment and resources needed for implementation was well beyond the means of one London borough. Our estimates (even in 1969) suggested that the computer hardware would cost about £1 million³ and that the development of the basic 'nucleus' of computer applications would involve at least 80 man-years.

Our conclusions had included an evaluation that a minimum of four London boroughs was needed for a viable approach. There existed already Hackney and

Tower Hamlets in the consortium. Follow-up studies in these boroughs revealed similar characteristics to those which had given rise to the Haringey recommendations. Three boroughs, however, were still not enough to show a cost justification for the new proposals, particularly because of the modest level of LEO III costs arising from the acquisition of a second-hand central processor and the use of existing programs which eliminated the development element.

Early in 1970 the London Borough of Hillingdon expressed interest in joining any consortium formed to pursue the objectives stated in the Haringey report. The decision to form the London On-Line Local Authorities (LOLA) consortium was taken in March 1970. Its computer staff were obtained by the en bloc transfer of the computer division LBMSU to the employ of the new consortium. This [fully] acknowledged that all the staff concerned were working full-time on computer services to the four boroughs.

In May 1970, after further extensive evaluation of hardware and software, an order was placed with IBM for the central equipment and software programs needed to implement the initial phase of the project.

3. The Choice of IMS/360

It is obvious that in a project such as this, the controlling software is of great importance. Although a number of manufactures could provide the necessary hardware, the software capability was not so easily available. The choice became selecting between the use of software already in existence and proven, or software which was then at the 'drawing board' stage or at best in the process of development. Given, at that time, that we had only two years to produce the first application, it was finally decided to use a proven package - IBM's Information Management System (IMS/360).

IMS is a sub-monitor system running under the control of the Operating System. Its purpose is to control hierarchically structured data-bases on direct access storage devices, and to provide access to the data by means of a network of remote terminals and by batch programs.

The main features of IMS which determined its suitability for the LOLA project are:-

- (a) the ability to construct and maintain hierarchical databases with records organised in segments allowing for the easy addition or deletion of segments and records.
- (b) the ability to specify logical relationships between records which are physically separate and for the relationships to be maintained by the software.

3 [About £15m (US\$ 20m or Euros 17m) in 2015 prices. 2016 note.]

- (c) the isolation of the application programs from the physical structures, storage and access methods of the databases.
- (d) the ability for application programmers to work in conventional high-level languages without having to learn specialised telecommunications or database languages.
- (e) security and recovery facilities which inhibited unauthorised access to data and allowed for reconstruction in the event of data corruption.
- (f) well developed facilities for the routine housekeeping associated with the running of a teleprocessing system.

4. The Integrated Data Bases

The records of local authority activities, and hence the sources of information for a data bank can be classified into three main groups:

- (a) Property; for example rating hereditament, council houses, premises registered for various commercial or industrial uses.
- (b) People and Organisations; for example ratepayers, creditors, council tenants, school children, electors.
- (c) Internal resources; for example revenue and capital budgets, staff, stores holdings, transport, financial investments.

To develop an integrated system of databases requires that access to data can be made other than by means of a referencing system peculiar to one particular office. For example, it should be possible to discover the contexts in which an authority has information about a property using only the address of the property as a reference.

At LOLA we are developing the databases within a general framework of controlling indexes. The concept is illustrated in Figure 2.

Each index contains general purpose information. Thus a property index could contain the full address of the property, description of the type of property, rateable value, age, use etc. In addition, cross-references are required to all databases which contain application-dependent data about this property. By interrogating the property index, using the address as the search key, the basic details of the property can be retrieved, together with the specific reference numbers given to this property in various applications. Authorised users can follow up specific references and access the appropriate application data.

5. Database Design for a Consortium of Users

5.1. One of the basic dilemmas facing anyone developing generalised systems for a group of users is to decide how to solve what we came to call 'the consortium problem'.

In essence, with four users (even with the commitment to a common system) everything can exist four times. Having agreed on a structure for the database, the decision must be made whether to have one database containing data for all four users or to have four databases with identical structures but different names i.e. are the files 'shared' or 'separate'.

5.2. The main advantages of shared files are:-

- (a) fewer files to define within the system.
- (b) fewer files to maintain.
- (c) less computer storage required.

These are important considerations as much time is spent in producing almost identical information and maintaining it. With just our first application, the main storage control requirements for 'shared' or 'separate' files differs by over 25,000 bytes.

5.3. The main advantages of separate files are:-

- (a) smaller and, therefore, more easily recovered in the event of corruption.
- (b) a full system can still be provided to other users following errors on one user's file.
- (c) security of access is simpler to implement.
- (d) with the present version of IMS, system throughput is greater.

There are other factors to be taken into account largely in the technical areas such as disk utilisation and data chaining. LOLA decided to implement separate files and events have justified our approach; I would expect the reasons to be valid for any form of consortium using IMS.

6. Database Structure

6.1. Reference has already been made to the controlling index of Property & People. Figure 3 gives an illustration of the interconnection between these indexes and an application's databases and also within the indexes. Only a subset of the physical segments in the application databases are shown. Figure 4 shows an

example of a logical structure as used by an application program making use of the cross-references (or 'pointers' as they are called in IMS) to obtain data held in different physical files.

6.2. One of the major advantages of the database philosophy in IMS is the ability to extend the structure of a physical database without impacting existing programs. Providing that the changes do not remove segments in a logical structure nor add segments within a logical 'branch', no changes are necessary to an existing program. This is of immense importance in the context of the Property and People indexes where additional interconnections will be made as more applications are implemented.

7. Application Development

7.1. 'Working Parties'.

It has been our practice for many years to tackle the problem of designing a common system for the consortium members by means of 'Working Parties'. A 'Working Party' is a group, usually about a dozen, with representatives of each authority and LOLA. This group determines the objectives of the application, considers problem areas presented in the form of discussion papers by the Systems Analysts, and agrees the specifications. The authorities are, therefore, involved in the design of the system from the very beginning.

7.2 Development timescale.

The development of a major application using IMS with both teleprocessing and batch programs has a long timescale, even in an installation which already has experience of such work.

Our first application included an overhead for creating the IMS system and learning how to design databases and use IMS. The major allocation of resources was:-

Systems Design & Programming	20 man years
Database design & IMS	19 man years
General purpose facilities	6 man years

Total	45 man years

This includes 7 man years of support provided by IBM. About 2,000 hours of machine time was used. For a subsequent application of equal complexity, I would expect these figures to drop by over 50%.

7.3. Program Testing.

There is an enormous problem connected with testing teleprocessing programs which run under IMS. Both the interface to the databases and the telecommunications interface are handled by DL/I.

Having coded the programs using 'modular programming' techniques and tested these with a module testing package, many of the logic problems have been corrected. To test the interaction with databases and the network, there are two possibilities:-

- (a) a testing package which runs as a batch program using test databases and simulating the network interface.
- (b) using IMS itself and entering test data from a remote terminal.

Up until the present we have been forced to use the second alternative due to the absence of any testing package. This approach is very expensive in machine resources and machine utilisation. Typically a testing version of IMS would require 300,000 bytes of main storage and might service only two or three programmers who between them are probably only entering one message a minute.

7.4. Database Design.

It was always intended to divorce the database design and maintenance responsibilities from the teams actively working on applications - that is to have a Database Manager. The justification of this rapidly became clear with a variety of interconnecting databases with different processing requirements.

The problem of database design using IMS is not so much a question of can it be done but what is the optimum way of doing it. There are usually a number of methods that will produce the desired results in terms of data retrieved but can differ very significantly in efficiency of machine utilisation - for example, a simple two-directional logical relationship can be implemented in over 20 different ways. Given the intended development of further relationships incorporating existing data, it becomes even more important to have a specialist group who fully understand the uses to which the data is currently put in order to prevent new applications impacting existing operational work. At LOLA there is a team of seven responsible for database design and maintenance and also for producing programs and facilities of general use.

8. The First Application

8.1. LOLA began its on-line system with a Rate Collection⁴ and Accounting system using video terminals for all aspects of interrogation and updating except rate receipts. The main advantages in taking this as the first were:-

- (a) it would provide a very firm basis of the property index since every rating hereditament would create an entry in the index.
- (b) it was a well-understood application technically in that we had written a batch rating system as far back as 1962 and this had been subsequently implemented in twelve authorities.
- (c) it was a 'self-contained' application affecting primarily one office in the authority.

Its disadvantages were:-

- (a) since rating is geared to either a 6 or 12-month cycle, the deadline (less than 2 years away) could not be missed. _
- (b) it resulted in notifications of money due being sent to about 50,000 people in the borough and, as everywhere with computer-produced bills, errors would cause anger and complaints, together with bad publicity for the authority.
- (c) the nature of the application forced a 'big bang' implementation strategy as opposed to a phased implementation approach.

The advantages, however, were too great to ignore.

8.2. The Conversion Process.

For three of the boroughs, the rates data was already held on the LEO III computer. The conversion process decided upon was to write programs which would operate on this data to:-

- (a) restructure it into the segments required for the databases;
- (b) separate it according to its physical database destination;
- (c) provide symbolic cross-references for all [logical] relationships;
- (d) load the physical databases;

- (e) resolve all the cross-references and provide physical pointers where necessary (using IMS utility programs).

This process took considerably longer than expected. One of the problems being that inconsistencies in cross-referencing were only identified at the latter part of the process whereas correction of the situation required modification to early programs. Other problems related to the incompatibility of the magnetic tapes on the two computers and to the fact that one of the IMS utility programs had an elapsed running time of over five hours with no re-start facilities.

In the end the conversion process, excluding re-runs, took over 17 hours of 360/50 machine time and produced our four main databases containing over 1 million database segments comprising some 50 million bytes of data.

There remains a necessity to periodically re-organise the data for performance purposes. This involves unloading, re-loading and re-referencing all the segments. Since most databases tend to grow in content over a period of time, this process can take well over 20 hours of 360/50 time. The process would presumably be more difficult if one had field-oriented databases.

8.3. User Training.

A direct result of computer systems for a consortium is the necessity to implement each application several times. Recognising that this creates special problems, LOLA has for a number of years used a special type of Systems Analyst whom we call a Systems Advisory Officer. Two of his major functions are to implement and review applications. To this we added a team concerned solely with Education and Training.

There are extra problems of training to be tackled for an on-line system, not least of these being that the computer ceases to be a remote god propitiated by large volumes of data submission documents, but that it intrudes directly into the user's office in the form of a video terminal which places the user in direct contact with the computer. To produce an effective group of users, it is essential that they should learn to treat the computer as a tool to use as a routine item of office equipment. To this end practical experience has no equal as a training aid. We adopted an approach of formal classroom training in the use of the facilities, followed by practice in the user's own office. This took the form of making the operational system available at agreed times for the users to submit realistic transactions against training databases.

Initially the error rate was high but in a surprisingly short time the users adapted to the new techniques. One of the minor by-products of the on-line system is the high proportion of competent typists now to be found in the Rates Office - this

⁴ [Rates were the council's principal source of income, being a tax on properties. 2016 note.]

despite the fact that full-time terminal operators are not used.

One technique that was found to be highly beneficial was the use of a terminal simulator. This is a video display specially modified to function in connection with an ordinary tape recorder. By this means it is possible to pre-record messages in and out of the system and subsequently play them back on the terminal. The users can, therefore, be shown the processes of a transaction without the necessity of having a computer available.

This is also proving a successful device for discussing alternative formats of display at the early stages of system design overcoming the difficulty in attempting to visualise the 'real thing' when looking at a print layout form.

9. Further Development

9.1 We have now four boroughs with an operational on-line rating system. This uses 36 remote video terminals supported by teletypewriters. Centrally we have an IMS system with 28 physical databases and 50 on-line transactions. These databases currently incorporate 390,000 Rating records, 185,000 names and 360,000 properties. Holding this data requires 10 disk packs, each with a capacity of over 30 million characters. All of this runs on a 360/50 with 512,000 bytes of main storage which is dedicated to IMS during office hours. The computer is already overloaded, of course, and will be replaced in August 1973 with a megabyte 370/158 system.

In addition we have the basis for development of further applications in a wide variety of areas, and by virtue of the ability to build on the Indexes we are not restricted to the 'traditional' areas such as payroll, stores accounting, rent collection etc.

9.2. The consortium members are now completing their appraisal of computer development requirements with particular reference to property data and it would be inappropriate in this paper to state specific areas of commitment in advance of their decisions. The expected rate of development should increase the network by 100 terminals before 1976.

Given the long timescale already mentioned for large applications, it is now thought desirable to consider a 'three-pronged' approach:-

- (a) developing major applications;
- (b) use of packages;
- (c) development of 'facilities' rather than applications, that is the maintenance

of basic data with a small number of on-line facilities to interrogate and update the information without a rigid and major requirement for routine processing

9.3. [unknown]

As major projects, LOLA is already well advanced with both a Financial Management application (implementation starts in October 1973) and a Purchase Control and Creditors application on which programming has started. Other possibilities are:-

Work Planning & Control (principally detailed progressing of maintenance work on council property e.g. houses, schools).

Manpower Utilisation.

Co-ordination of Capital Projects.

Land Charges and Property Register.

9.4. A number of packages already exist which could be beneficial to the consortium, for example:-

Planning - the Central London Land Use System (CLUSTER)

Engineers - Survey Analysis, Technical Design

Architects - Critical Path Analysis

9.5. Development of Facilities.

These are based on rapidly developed extensions to the existing databases giving a speedy implementation and quickly realised benefits to the users, for example Rates/Rents interface — To make selected rating data relating to council-owned property available to the Housing Department plus facilities to hold details of council tenants.

Property physical characteristics — A further expansion of the property index to assist the Housing Department in the areas of Lettings and property maintenance.

Property Progressing and Monitoring — facilities to allow a variety of departments to record significant activities against properties for progressing and 'diary' keeping requirements.

Similar possibilities exist relating to the People Index.

9.6. This paper has sketched out what LOLA has achieved in the past two years and indicates the approach which has led us this far. Behind the scenes lie many lost

evenings, nights and weekends in achieving the first on-line IMS system in Britain — which despite all the problems (compounded in February 1972 by electricity cuts arising from industrial action) met its deadlines. The major problems arose from software errors, untried terminals and lack of experience (both IBM (UK) and ourselves) - hopefully either gone, or much reduced, in the future.

We are now at a major point of our development, we have a springboard for rapid development and the next year should [verify] the far-sightedness of our constituent authorities to their courage in supporting innovation

Tower Point North,
Sydney Road,
Enfield,
Middx.

22.3.73.

Figure 1

THE INTEGRATED INFORMATION SYSTEM FOR LOCAL GOVERNMENT

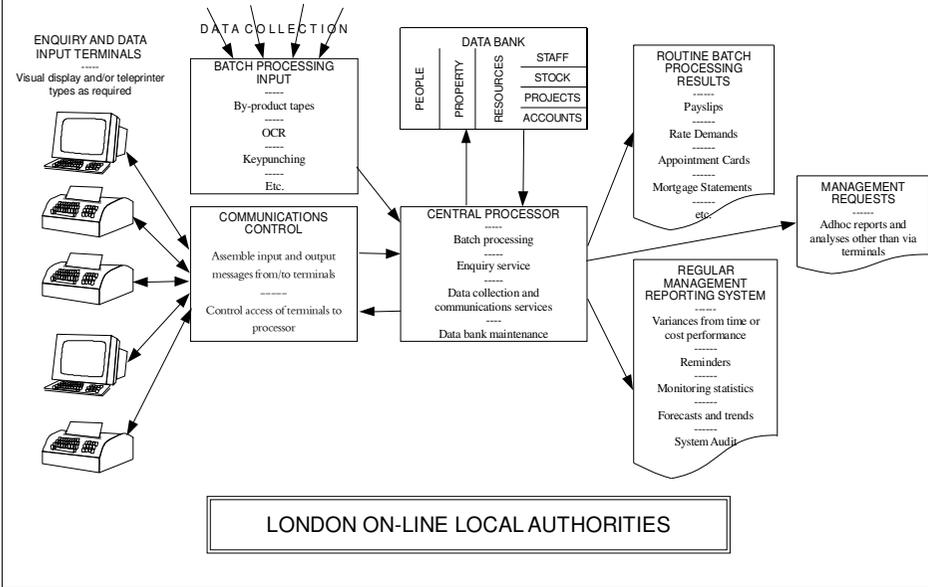


Figure 3

[Controlling index of Property & People and interconnection to two application's databases
Only a subset of the physical segments in the application databases are shown.]

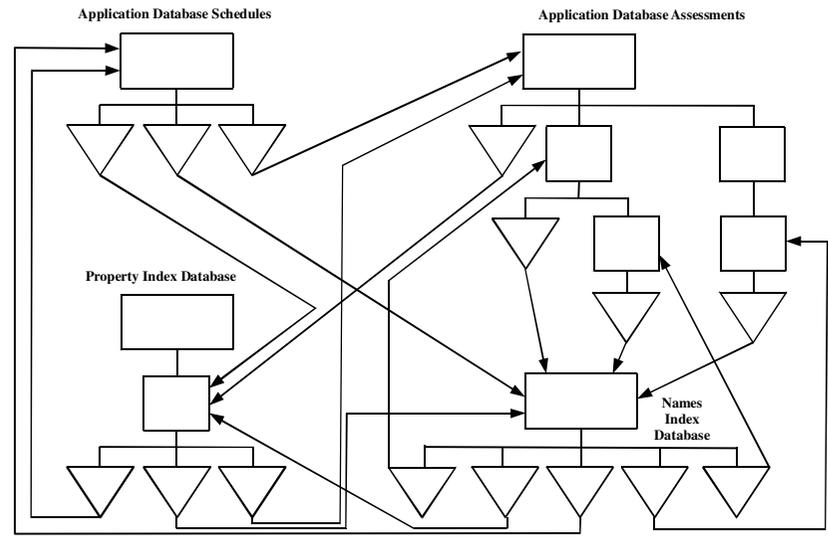


Figure 2

CONCEPTUAL DATA BASE STRUCTURE

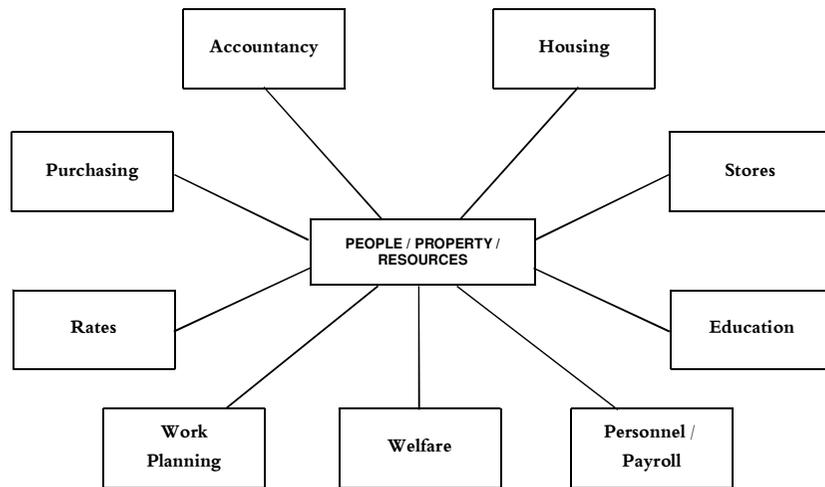
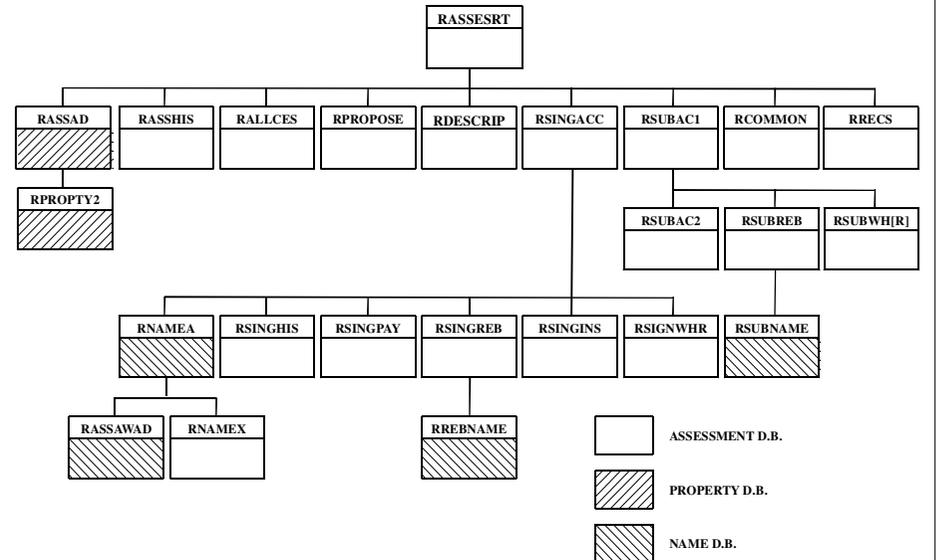


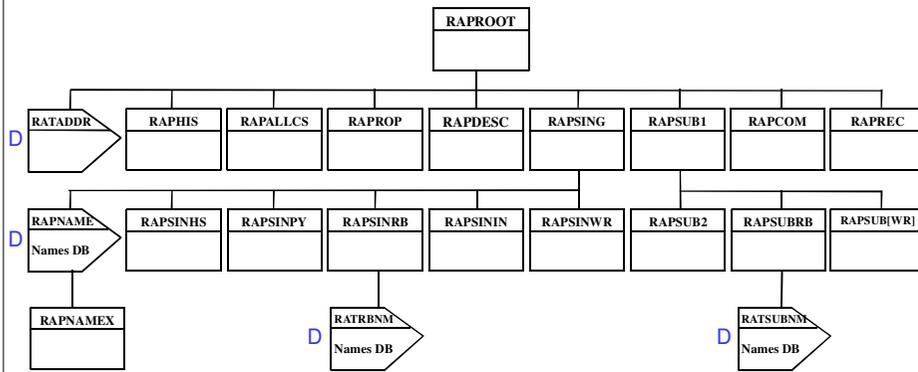
Figure 4

Assessments Logical DB



[Figure 5]

RR001A – Assessment Physical DB

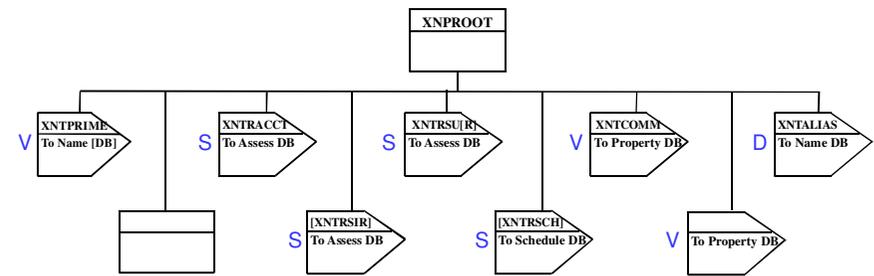


D = Direct

The blue text are items that were hand annotated on the original paper. This page, whilst contemporary with the presentation paper and was attached to a copy, may not have been included in any presentation handout.

[Figure 7]

RX002A – Names Physical DB

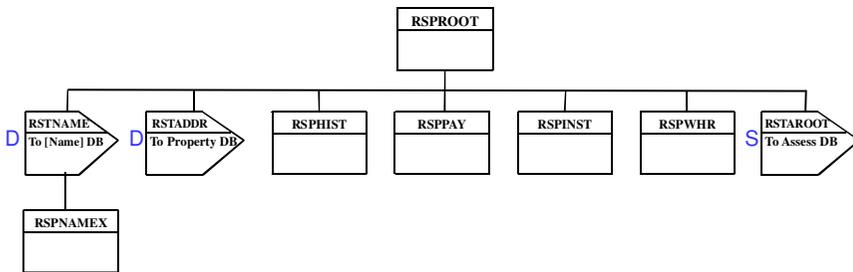


D = Direct
S = Symbolic
V = Virtual

The blue text are items that were hand annotated on the original paper. This page, whilst contemporary with the presentation paper and was attached to a copy, may not have been included in any presentation handout.
The segment names in [] have been taken from *Experiences in Using IMS 2*

[Figure 6]

RR003A – Schedule Physical DB

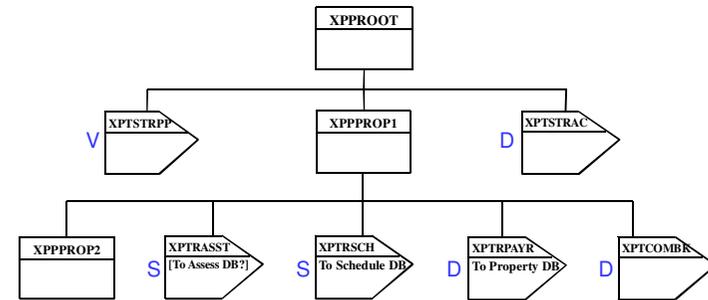


D = Direct
S = Symbolic

The blue text are items that were hand annotated on the original paper. This page, whilst contemporary with the presentation paper and was attached to a copy, may not have been included in any presentation handout.

[Figure 8]

RX001A – Property Physical DB



D = Direct
S = Symbolic
V = Virtual

The blue text are items that were hand annotated on the original paper. This page, whilst contemporary with the presentation paper and was attached to a copy, may not have been included in any presentation handout.