

THE TITLE PLANT OPERATING SYSTEM
A Data Base System of Index Files for Recorded Documents

By: Jerry Koory, January 2007

Author Sketch: U. of Calif. Santa Barbara, Education
System Development Corp.
Planning Research Corp.
H W Systems Inc.
Consulting: The Rand Corp. Ticor Inc.
1990: Retired

Copyright © March 2007, Computer History Museum, All Rights Reserved

In 1965, the Planning Research Corporation of Westwood, California was awarded a fixed price contract to computerize the title plant for a consortium of small title companies doing business in Los Angeles County. The title plant was owned by Title Records Incorporated (TRI) which was owned jointly by four title companies.

TRI had been keypunching cards to index the microfilm that is produced daily by the Los Angeles County Recorder and the Superior Courts of the County, which contains copies of documents pertaining to title and to an individual's ability to hold and pass title of real property. Just over 10,000 documents were recorded by the recorder and the courts on each working day. The pertinent documents were abstracted onto keypunch cards, usually one card per document. Two hundred fifty working days per year X 10,500 documents = 2,625,000 index records (keypunch cards) produced each year. They were using more than a case of cards each day. The punch card salesman loved them.

To use the card plant, each day's batch of index cards had to be sorted by property description or person's name (encoded) and then merged into the plant. It could then be searched to find where, in the store of microfilm, were documents that had to be examined as part of the process of issuing a new title policy when the property in question was sold. The person doing the document search would go to the card files, find the appropriate tray containing index cards for the property in question, and let his fingers do the walking down the tray until he found the cards he wanted. He would then pull up the cards and write down the document numbers and dates for documents to be copied off the microfilm for examination by the title officer.

The title companies knew that they had to find a better and more efficient way to store and search this ever growing plant. The answer, according to IBM, was the IBM S/360; in particular, the 360 model 30 with data cell drives as the primary storage facility for the data base: that is, the index files to the microfilm records.

PRC was a company with experience in working on defense related projects doing system analysis on large information projects for various elements of the Armed Forces under cost plus contracts. The company had no experience with fixed price commercial business. This made for some interesting discussions and nervous times by the senior management. Furthermore the technical staff had no experience with

the IBM 360, but then, at that time, neither did anyone else, except IBM. My background, as well as everyone else's, was with word computers; the IBM 704, 709, 7090 and Control Data 1604. The 360 was a new animal. To make matters worse, the data cell drives were also new, and none had been delivered at the time we started the project.

Stan Wong and I moved into a small office at California Land Title in downtown Los Angeles so we could learn what a title plant was and how it was used. While we were doing that, Bill Brehm was trying to learn the 360 and its Disk Operating System, as well as the characteristics of the data cell drives. Understanding the data cell drives and how they could be used was a key to the later design of the Title Plant Operating System.

At the same time we were getting our feet wet, IBM was trying to encourage us to use Index Sequential; a software file access system they thought would do the job. Some of the IBM personnel (system engineers) were also working for the big competitor of our four title companies, Title Insurance and Trust (TI). TI provided 85% of all the title insurance policies issued in Los Angeles County while our four title companies combined, did about 5% of the business. TI was the 900 pound gorilla for title insurance in Los Angeles County. It was a big day in the life of California Land Title when Executive Vice President Jack Edwards came into the office Stan and I shared and announced that there would be a celebration after work, because they had hit the great milestone of 2% of the business the previous month.

IBM had sold TI a 360 model 40 and they were busy trying to do the same thing that the three of us at PRC were doing. The IBM system engineers were planning to use Index Sequential and wanted us to use it also. Once we understood that software and how it worked, we decided that it would be too slow in processing the very large data base we were looking at. The 360 model 30 was just not fast enough, along with the data cells, to do the job using Index Sequential. So we had to build the data base system from scratch.

As it turned out there were actually two title plants. The property related records from the county recorder (the property index (PI)), and the individual related records from the courts, (the general index (GI)). The property index was divided into subsections by property type, of which there were 6 or 7 different types (section land parcels could be broken into quarter sections, quarters of quarter sections, and lots within quarter/quarter sections). The GI was indexed on the names of the party or parties affected by the document. The title companies had been using Russell Soundex to encode the last name of individuals so the card sorting equipment could handle the processing.

Russell Soundex produced a coarse index, but was adequate in the punched card era. One of the advantages afforded by the 360 was that we could use the person's full last name as the indexing key. That, of course, brings up the question of the spelling of names. Many names can be spelled differently but sound the same, for example; Swensen, Swenson. Documents recorded by the courts are not checked for these sorts of spelling problems. There could be two different judgments entered against the same John Swenson with one of the documents containing the name spelled Swensen. Fortunately, the title companies were willing to put the burden of complete searching onto the searcher by having them provide search requests for various spellings of the name.

As we came to understand the data cell drives and their benefits and limitations, we realized that we could not expect to process, sort and merge the day's batch into the entire title plant in one night's processing. We chose to mimic what TRI was doing with the card plant. Today's batch was processed and merged with the previous day's batch. This was known as the Current Plant. The older and much larger plant was known as the Main Plant. TRI would continue building the Current Plant daily until it became too large to be processed in the time available. It would then be called the Intermediate Plant, and a new Current Plant was begun. TRI would then schedule a weekend (or more) to merge the intermediate plant into the main plant creating an even larger main plant. We designed the system to operate in the same way without the Intermediate Plant step.

Stan and I wrote a Functional Design of the Title Plant Operating System to define for TRI what the system would do. This was presented to TRI along with a question and answer session. We asked them to read the Functional Design document preparatory to signing off their approval. They were asked to note questions and problems that they saw with the design. After a week, we met with them, got their comments, questions and suggestions, and modified the document accordingly. This was important to us, as it now defined what we were going to build, with the idea that when we demonstrated that it worked, we would get paid.

Stan, Bill and I got to work on a system specification including flow charts and data layouts to guide the programming. TRI was not asked to review and approve this work as they did not have the capability to do so. The characteristics of the data cell drives and the disk drives were the key to using the capacity efficiently. Our design involved creating a low key index of the sub cells and disk tracks so that with a binary search we could quickly obtain the address of the portion of the data base where the indexes to the document in question resided. Each index record was just 60 characters (bytes) in length, which we wrote onto the storage tracks to maximize the capacity of the tracks. Programming got underway with us not having access to a 360.

About the time we were ready to begin testing some software, IBM announced a delay in the delivery of the 360 and the data cell drives. How could we proceed? IBM was willing to give us computer time at their data center on a 7094 using software to simulate the 360 and the data cell drives. That was not very appealing to me. I had some prior experience with IBM delays and was concerned that the six month delay might turn into a 12 month delay. Since this was a fixed price contract, I couldn't charge TRI for this waiting time, and they were in no position to support what was now a staff of four analysts and programmers sitting around waiting for IBM to deliver.

TRI called for a meeting with IBM and myself to try to work something out. I took a contract administrator from PRC with me. As we were driving to the meeting I told him that he was not to say anything unless I asked him a question. He had not met any of the people in the meeting before and would not understand the subject. I told him that if I asked him a question his answer was "no". No conversation, just "no". That worked very well. When IBM or TRI would suggest something I didn't like I would say "Jim, can we do that?", and Jim would say "no!". The result was that we pulled our crew out of TRI, and TRI paid us a penalty because of the unexpected idle time for the staff. As it turned out, it was close to a year before we were able to test any of our software on a 360/30 with data cells.

Almost 10 months after the announcement of the delay of the 360, we were at last able to test a program on a 360/30 with data cells on a system that IBM had just installed in another customer's facility. This was the first installation of this configuration within the territory of TRI's IBM Branch Office. We were allowed a couple of hours to run a special test program that Bill had written to demonstrate the writing and then reading of data with data cells. The test was successful. We now had a delivery date that was close enough to put us back to work on the project.

During our design work Stan and I had been concerned about how to incorporate backup files into the process. To make copies of the plant on data cells would be a very time consuming process as well as expensive for TRI, as the sub cells were expensive. We solved the problem by using the tape drives. During file reorganization, the process of merging the current plant into the main plant, we would write the merged data out onto tape. When the reorganization was complete, the new main plant would be loaded into the data cells while creating the low key index. The tapes that held the merged plants would be saved and moved off site, thereby becoming a backup file for recovery purposes.

After TRI received their 360/30, we were able to complete our testing. The last problem was to load the data bases from the more than 13 million punched cards that TRI had been accumulating over a period of more than 4 years; that is 1,300 cases of cards. It took a week or so working two shifts a day to finish.

Initially, search requests were keypunched into cards, fed into the system via the card reader and the results, lists of documents in the microfilm files to be examined, were printed. Later, remote searching was made possible by using terminals and printers at branch offices connected by telephone lines. The system was expanded to include other counties in California along with upgrades to the computer(s) and storage devices. It was a great day when the data cells were retired, having been replaced by 3330 disk drives.

As of this writing, the end of 2006, the Title Plant Operating System is still in use, albeit with many modifications and improvements.