QUESTION-ASKING QUESTION-ANSWERING SYSTEMS

by

Laurent Siklossy
Department of Computer Sciences
The University of Texas at Austin

May 1977 TR-71

To be presented at and published in proceedings of: International Seminar on Intelligent Question-Answering and Data Base Systems, June 21-30, Bonas, (Gers), France. Sponsored by: AFCET, ASCEB, DRME, IP, IRIA. Director: Prof. J-C Simon, Universite de Paris VI. Co-director: Prof. L Siklossy, University of Texas at Austin.

ABSTRACT

We describe conceptual and initial implementation work carried out to investigate a new class of question-answering systems operating on large, complex data-bases. These systems have the following capabilities:

1) they may suggest to the user new questions, similar to those asked by the user, but that may not have occurred to him since he has limited knowledge of the data-base.
2) they answer questions in complex, non-linear data-bases. Hence, this is an effort to investigate data-bases other than the typical inventory data-bases, and their related set-theoretic cum elementary statistics questions.
3) they make suggestions to the user about information related to the question, but not explicitly contained in the question.

Alternate questions are suggested when the user's question is near a discontinuity in the data-base. In this case, the "Well, you didn't ask me..." syndrom is avoided. The data-base is organized along several dimensions to make such discontinuities apparent. A discontinuity results if a small change in one dimension of a question could result in a great discontinuity in the answer.

Illustrative examples from tax laws, freight transporation, design and in particular air fares are given.

1. INTRODUCTION

The manipulation of data bases, in particular of large data bases, will probably become an increasingly greater task of computers in the years to come. The data bases will be subjected to queries, often from human beings, with the help of what is known as a question-answering system (QAS).

Our research has its origin in our dissatisfaction with the apparently universally accepted scheme for the design of a question-answering system, as well as in the present limitations of data base design and in the capabilities of question answering systems operating on such data bases.

The present scheme for the design of a question-answering system is as follows:

a) the QAS has access to a data base. This data base can be dynamic or static, and its organization is not of importance at this level of design.
b) a user (man or machine) makes a request to the QAS. The request could be a request for a retrieval or could involve comparisons and statistics (How many female employees of Company X earn more than the average salary of male employees?)
c) the request is manipulated by a number of subsystems in the QAS. If the request is given in a natural language, a linguistic component must be present. Retrieval procedures, query optimizers, security checkers, statistical packages, etc. may also be invoked.

If the request is ambiguous, the question-answering system may request additional information (from the user), to disambiguate the question. If the question cannot be answered, the user may be asked for additional information that may help to answer the question. For example, in the last example, the QAS may reply: "Do you mean ... male employees of Company X?"

d) eventually and finally, the QAS has a specific question that it answers.
e) this answer is given to the user in a form which (hopefully) he can use.

We now wish to show that the above scheme for a QAS is unsatisfactory. We first give a hypothetical anecdote as illustration.

About a year ago, while walking in the streets, I met Mrs. A., the wife of my friend Mr. A. Mrs. A, without a doubt, was pregnant. I congratulated her. I see Mr. A fairly frequently, and at our next meeting, I congratulated him, and mentioned:
-You hadn't told me that your wife was pregnant.
-Well, YOU DIDN'T ASK!
That was true, I hadn't asked. On the other hand, it is not considered particularly proper to be constantly asking friends and acquaintances whether they or their wives are pregnant. Nor is it considered in good taste to ask constantly a great number of other questions. However, it could be socially acceptable to ask the question under circumstances when it would be meaningful to ask it:
-Mrs. B told me that your wife was pregnant. May I congratulate you?

Here is another example closer to some domains of interest to us. Suppose that on February 14, 1976 you had called up the office of your travel agent to reserve a passage from New York to London and return. You and the soft voice at the other line (Mrs. Eglesby or so it sounded) have the following conversation:
-When would you like to leave, Mr. C?
-On June 2.
-And on which date would you like to return?
-On June 15.
-Very good. Which class of service do you wish to use?
-Economy class.
-There is a 10 PM departure from Kennedy airport on TWA non-stop to Heathrow airport in London on June 2. On June 15, there is a PANAM flight from Heathrow to Kennedy departing at 11 AM, also non-stop. Would these flights be

convenient?
-Yes, I guess so.
-Fine, I shall make the reservations, etc. The price of the ticket is N.

This conversation is rather typical. Yet, a number of interesting points are worth noting. Mrs. Eglesby did not mention departures on a variety of other airlines. It may have been an effort to avoid further balance of payments difficulties to the USA. It may also come from the fact that the U.S. carriers will give her a larger commission, or that whe will receive this commission sooner. After all you didn't ask her if there were other airlines!

On June 2, on your flight, you sit next to Miss D (she is really too young to be Ms D; remember, there were no youth fares then). She also goes to London, but returns on June 17. You also find out that she is paying considerably less for her ticket than you. In London, you meet Mr. E. He left New York on May 31, and returns to New York on June 13. He paid less for his ticket than you did.

On the return flight, your luck has improved. You sit next to Ms. F (you saw it coming!). She has left New York on May 31, and she is paying even less for her ticket than either D or E (and, naturally, you!).

Back on the farm, you call Mrs. Eglesby. The conversation is monotonous, but here it is:
-Why didn't you tell me that if I stayed a couple of days longer in London my airfare would have been less?
-YOU DIDN'T ASK.
-Why didn't you tell me that if I moved my entire travel plans a couple of days forward, my airfare would have been less?
-YOU DIDN'T ASK.
-Why didn't you tell me that if I had left a couple of days earlier, returning on the same date, my airfare would have been even less?
-YOU DIDN'T ASK.
-Why don't you ... (deleted).

What has happened? There is no way to fault Mrs. Eglesby, of course, strictly speaking. You made a request to her, and she answered it, as you would expect any smart QAS to answer it. Or perhaps it is not smart, but simply stupid. Clearly, she should have informed you of the rules, which (incompletely specified) in this case involve two components
-trips for an absence of 14 to 21 days benefit from a special lower excursion fare.
-April and May are "shoulder" season. June is "high" season.

Fares are lower in "shoulder" than in "high" season, both for regular fares and for excursion fares. (All of this changed in 1977, hence our choice of 1976 for this story!) You didn't ask, simply because you didn't know. You simply had no idea that such a question would be relevant. There just wasn't any reason to ask the question, that's all. Clearly, Mrs. Eglesby should have suggested some alternative days of departure or return for your trip. If your travel plans had been inflexible, good, she would have answered your original request. If not, you might have saved some funds.

2. DISCONTINUITIES

How can we understand in more general terms what is happening here? Our theory is that a QAS should go into the question-asking mode when the request is in the neighborhood of a discontinuity. A discontinuity is present if a small change in one dimension results in a great change in another related dimension.

In our example, the request was close to two discontinuities. One discontinuity is the May 31/June 1st discontinuity. A small change in departure (technically even a few minutes!) may move you across the discontinuity barrier: May 31, midnight. But such a small change can result in a large change in the airfare. The second discontinuity is the 13-or-fewer-days-stay/14-to-21-days-stay barrier. Again, a small change in the duration of the trip near the barrier could result in a large change in the airfare. In the example given, the effects of both discontinuities were additive.

More formally, we can express the criteria for a discontinuity. Given a function H of the variables v1, v2, ... , vn; H(v1, v2, ... , vn), there is a discontinuity of H near V ≡ (V1, V2, ..., VN), where Vi is a value of the variable vi, if the value of H near V can be much different from the value of H at V. Notice that we say "can be" and not "is", the value of H need not be much different for all points near V.

The terms "near" and "much different" are not absolute. They depend not solely on the data base, but also on the user. To return to our example, if the departure date had been June 14, then it would hot have been "near" the May 31/ June 1 discontinuity for many travelers, although it could still be "near" for some, for example a retired captain who is going museum hopping.

In the same way, "much different" is also relative to the particular user. In the late fifties, one could travel by propeller planes across the Atlantic, or by jet planes. The difference in fares (some $20 at some time) would be small for many travelers, while large for others. In this example, we have in fact one discontinuity (the type of aircraft) affecting several functions H: comfort (for example, noise); cost of airfare; duration of air travel; safety; etc. An informed user would be able to choose according to his own criteria.

It follows from our analysis that our desirable QAS should be user oriented. It sould ascertain reasonable values for "near" and "much different" to be used near discontinuities.

2.1 What if you asked the "right" question?

Suppose that the QAS has determined that the request is near some discontinuity (or discontinuities), but that it is on the "right" side of the discontinuity, i.e. a change near the request would only lead to a worsening of the answer. That would be the case of Ms F: the cost of her airfare would be higher if she left a little bit later, or stayed a little bit less time in London. We claim that the discontinuities should still be pointed out, as well as the attached dangers and costs, at stepping over them.

2.2 Multi-dimensional discontinuity.

As the example of the propeler/jet discontinuity showed, a discontinuity can result in "much difference" along several dimensions. Under these circumstances, the QAS cannot usually decide on an answer to a request. Even such requests as:
-I want the cheapest flight from New York to London.
should not result in automatically supposing that the user wants a propeller flight. The question should be interpreted as:
-Given the types of flights that I know about, or am willing to endure, I want the cheapest flight from New York to London.
Hence the QAS should not blindly go ahead and assume that cheapest must mean propeller. The QAS should ask from the user:
-There are some propeller flights. They make stops in Gander, Shannon and Manchester, take 5 hours longer than jet, and cost $20 less one way. Would you want a propeller flight?
The answer would, of course, depend on the user.

We shall return to other examples of discontinuities, but want to first investigate the types of data bases under question, and the algorithms that query them.

3. DATA BASE TYPES

Developers of data base management systems have designed three principal types of data base systems: hierarchical, network and relational data bases. It has been shown that all three designs are theoretically equivalent. They differ in their ease of use, ease of implementation, proximity to the user's concepts and organizations, etc. As a result, from the study of relational data bases, one can characterize the requests that may be issued to such data bases. The requests are all expressed as (non-recursive) operations on sets. Notice that in relational data bases, the data base is a set of tables representing relations. A relation in turn is a subset (of some appropriate cartesian product). The permissible operations on the relations are all simple set-theoretic operations (or amenable to these operations): set membership, projection, intersection, union, join, etc. To these set-theoretic operations one usually adds some elementary statistics (counting, averaging, etc.).

The above mentioned data base designs have been quite successful in representing and querying a large variety of data bases. But due to the nature itself of the data base operations allowed -set operations- all the data bases are essentially equivalent to "inventory" data bases, since the fundamental question -set membership- is present. There can be many sorts of inventories. Inventories of people in airplanes (airline reservation systems); employees in an organization; goods shipped or received; classifications of objects, etc.

3.1 Non-linear data bases.

Other data bases, for example the data base of airline tariffs, do not fit easily into the "inventory" frame. Although it could be argued that functions could be all defined starting from sets, the types of question-answering activities carried out have no resemblance to tests of set membership cum statistics.

It may be enough to consider simply the organization of a data base of airline tariffs that would include information on the various costs of a New York - London round trip fare for the various cases considered in section 1, and ways to retrieve these questions. We doubt that present set-oriented data base management systems could successfully cope with such a data-base.

In fact, the airline tariffs are much more complicated than the previous examples would illustrate. Let us consider an example. We are personally aware of a Stockholm to Stuttgart to West Berlin to Toulouse trip with stopovers in each city. The initial fare (as calculated by Air France) was obtained as follows: add the three separate fares Stockholm to Stuttgart, Stuttgart to West Berlin and West Berlin to Toulouse to obtain the total fare. A very clever European travel agent calculated the fare differently. He wrote a Stockholm to Alicante (Spain) ticket, with stopovers in Stuttgart, West Berlin and Toulouse! The mileage of the trip was within the maximum 25% increase on the allowable Stockholm to Alicante mileage. The ticket cost about $75 less (in 1974).

Notice that the calculation is certainly not obvious. Alicante is not on the itinerary. A point closer to Stockholm (for example Barcelona) would not have provided enough mileage. A point much farther than Alicante, for example Dakar, would have resulted in a fare larger than the initial one. It is not even certain that the fare calculation using Alicante was minimal. By law, the fare charged a passenger must be the minimum fare calculable. We would want our system to be good at applying existing fare calculation algorithms.

Although the rules on how fares can be calculated using "hidden cities", such as Alicante in our example, are clearly and unambiguously defined, the implementation of these rules in actual algorithms is far from trivial. Which candidate hidden cities should one select? One could have also chosen a hidden city prior to the point of departure: for example, a Leningrad to Toulouse ticket, with intermediary stops in Stockholm, Stuttgart and West Berlin. One could choose both a hidden city before Stockholm, and another one after Toulouse, Or one could have selected intermediary points for breaking the fare. The various restrictions on fare calculations must also be taken into account.

The discontinuities mentioned in section 2 indicate that such data bases as the airline tarifs are non-linear. (We shall discuss other non-linear data bases in the Appendices.) Algorithms that process non-linear data bases must often resort to case analysis, and could be combinatorially explosive. Their efficient implementation may require the use of multiple representations of and classifications in the data base.

4. ADDITIONAL NON-REQUESTED INFORMATION

We have seen in section 3 that we want our QAS to possess powerful algorithms to operate in large non-linear data bases. We saw in section 2 that the QAS should not necessarily simply answer the user's request, but may help the user formulate an alternate improved request. We shall now see that the QAS should not limit itself to answering the request that the user has finally agreed upon. The QAS should also give additional information that may be relevant.

Let us return to the New York - London trip. Assume that the city of London wishes to encourage travelers to spend the first few days of their stay in Europe in London (while they still have some money) so it designed the SPEND-HERE-FIRST promotion. To that effect, it offers very low hotel fares for exactly the first four nights to an incoming passenger. Hotel vouchers must be purchased before arrival in London.

Did Mrs. Eglesby mention such arrangements? No. Should she have? Probably yes. And that should be our QAS's behavior: it should mention the SPEND-HERE-FIRST program unless it is positively known that such a program would not interest the user.

Let us consider another example. The cost of a Frankfurt - Lisbon - Oporto ticket is the same as a Frankfurt - Lisbon ticket (in 1976). Hence a user wishing to purchase a Frankfurt - Lisbon ticket should probably be informed of her possibility to travel on to Oporto free (the magic word!).

Some limits must be set on what will be suggested. In Europe (but not in North America) there are many possibilities for making stopovers in intermediate cities when travelling between two cities by air. Clearly, the QAS system should not suggest all such possibilities, but perhaps suggest only the possibility in principle. If the user is interested, some of the major intermediary cities could be mentioned.

From our examples, we can try to define the attributes of the additional non-requested information that the QAS system should volunteer. In fact, we can define them in terms of discontinuities! The QAS should volunteer additional non-requested information if a "null or very small" change in one dimension of the request would be consistent with a "large" change in another dimension. As in section 2, "very small" and "large" are relative to the user. In our examples, a null or small change in the cost of the trip would be consistent with "large" changes in a variety of dimensions: cost of overnights; additional city visited beyond the original itinerary; additional cities visited along the original itinerary.

5. CONCLUSIONS

We have discussed question-answering systems (QAS). We have shown that present designs studied in data base management systems were inadequate to represent data bases and their associated QAS which had the properties:

a) the user's query is not to be treated as sacro-saint. Instead, neighboring queries should be suggested if they might still translate the user's intentions, and if they result in large changes in some dimensions (such as cost).

b) the data base is complex, and algorithms that search the data base are themselves complex.
c) the user should be informed of additional information not necessarily contained in the finally agreed upon query.

The unifying concept overall is that of discontinuity. We are interested in non-linear data bases; i.e. data-bases which contain discontinuities. The proximity of a discontinuity to a query indicates that alternate questions may be suggested. Finally, the presence of a discontinuity near the answer may suggest additional information to be conveyed by the QAS to the user.

Usually, non-linear data bases are large, and hence size introduces an additional level of difficulty for implementation. We have considered non-linear data bases of various types: airline tariffs, freight transportation, taxes, design problems, etc. Our examples were taken from the domain of airline tariffs since we would expect it to be the one best known by the reader. In addition, the airline tariff domain has been partly implemented, as described in section 6. The Appendices include information on the data bases mentioned, emphasizing their discontinuities.

A problem with these data bases which also represents a serious implementation problem is their volatility. The data bases and the algorithms that process them are changed with great frequency - often by governmental fiat. - The updating problem is theoretically without interest if the data base and algorithms remain sufficiently similar. Practically, though, the updating problem is serious.

6. ACKNOWLEDGMENTS

During the Fall of 1976, the author was Richard Merton Professor at the University of Karlsruhe, West Germany, with funding provided by the Deutsche Forschungsgemeinschaft (DFG). During that time, with the great help of Dr. Peter Raulefs of the Institut für Informatik, Universität Karlsruhe, and German, Mexican and Peruvian student co-workers, an implementation of a QAS satisfying the criteria discussed previously in this paper was started. The initial implementation would consider only Western Europe, including parts of North Africa, the Middle East and Asian U.S.S.R. This domain consists of over 500 airports. In the first stage, only regular fares would be considered. Routing, distance and mileage information were incorporated.

In 1977, the implementation efforts were continued under the direction of Dr. Raulefs. The most recent results of the implementation will be given at the time of the seminar. It is expected that the following capabilities, at least, can be exhibited:

- hidden city faring, such as the Stockholm-Stuttgart-West Berlin-Toulouse fare. In fact, we should then be able to state whether the Alicante fare was the lowest or not.
- rail-air connections. Example: instead of leaving from Strasbourg, take a train to Stuttgart, and fly from there. It's much cheaper.
- multiple stops, including stops beyond the final destination, such as the Frankfurt - Lisbon-Oporto example.

APPENDIX A. AIRLINE TARIFFS

Anyone who has traveled is aware of the great number of fares for travel along certain itineraries, and of the great variation among fares calculated by the airline and travel agents for the identical itinerary. Consumer Reports and other consumer oriented magazines have regularly reported on these facts.

As an example, between New York and say London the excursion rates on regular I.A.T.A. carriers include a 14 to 21 days excursion and a 22 to 45 days excursion fare and for each type of excursion, three tariffs (in 1976): high season (summer), low season (winter) and shoulder season (in between). The non-linearity of the fares becomes immediately apparent: by extending one's trip from 21 to 22 days abroad, and leaving one day earlier (switching departure from June 1, start of summer season, to May 30 which is still shoulder season) and assuming no stopovers in Europe, one can save a considerable amount since shoulder season is cheaper than summer, and the 22 - 45 days excursion is itself lower priced than the 14 - 21 days excursion. The difference is several hundred dollars.

The Stockholm - Stuttgart - West Berlin - Toulouse fare is an example of the use of a hidden city, i.e. a city beyond and not on the itinerary chosen. We give some additional examples which we have actually used personally.

There are no published fares between Great Falls, Montana, and Austin, Texas. One could calculate the fare from Great Falls to Denver, Colorado, and then from Denver to Austin. It is cheaper to calculate a Great Falls to San Antonio, Texas (which is beyond Austin), which is a published fare, and get off in Austin. San Antonio then becomes the hidden city.

Some hidden cities are more fun and less obvious. The U.S.A. to Europe 22 to 45 days excursion allows no stopover in Europe, i.e. you cannot fly from New York to London, stay in London, fly to Zurich, stay there, and then fly back on the same ticket. You can have an "open jaw" which allows you to fly New York to London on the 22 - 45 days excursion fare (one half of the round trip fare) and then fly Zurich - New York again on the 22 - 45 days excursion fare (again one half of the round trip fare). But the London - Zurich transportation would have to be arranged separately. For someone having a "large" open jaw, i.e. where the point of arrival from the U.S.A. and the point of departure back to U.S.A. are far away, it may become worthwhile to look for a hidden city. The hidden city must be outside Europe and North Africa, which allow no stopovers, and must allow some stopovers. Recently, we have made a trip which had open jaws in London (arrival) and Marseille (departure) with an intermediary stopover in Bucharest, Romania. The hidden city was Tripoli, Libya! And the alternate fare calculation using Tripoli was indeed lower. Hidden cities can be found that get around restrictive lengths of stay. Inside Europe, Budapest, Hungary, is often a convenient hidden city for Vienna, Austria. Excursion fares to Vienna usually require fairly long stays in Vienna before the return can start (perhaps two weeks). Excursions to Budapest have minimal stay requirements (perhaps only one overnight), and from many points in Western Europe are cheaper than the regular fare to Vienna.

Tariff tables give allowable fare mileages between cities, and also give actual mileages. The allowable mileage is at least as great as the actual mileage, and sometimes much greater. Actual mileages are computed from point to point on an itinerary, and added up. If the total mileage (usually calculated separately on both the onward and the return leg of a round-trip) is within the allowed mileage, there is no fare increase. Otherwise, the allowable mileage is incremented by steps of 5%, up to a maximum 25%, until it is higher than the actual mileage (a number of other details must be taken into account). Since allowable mileages are not a fixed increment above actual mileages - they take sometimes into account the difficulty of reaching a point, which may require several changes of airplane - it becomes possible to choose itinerary points for which the allowable mileage is very high. For example, from Scandinavia, the allowable mileage for

Rotterdam and some Greek islands is very much greater than the actual mileage. A good fare calculator might therefore hold in his head a distorted map of the world based on allowable miles. Notice that this distorted map is relative to some point of focus (Scandinavia was just cited). The map may change if focus is on another point, for example New York.

Occasionally, even on simple direct itineraries, fares to farther points may be lower than fares to less distant points. In Texas, Dallas, Austin and San Antonio are approximately North to South, in that order. Austin to San Antonio is about 60 miles. The airline Braniff International flew the Dallas - Austin - San Antonio route, while Southwest Airways flew Dallas - San Antonio but not Dallas - Austin or Austin - San Antonio. To compete with Southwest (and in an effort to eliminate the competition, probably) Braniff lowered its Dallas - San Antonio fare below the Dallas - Austin fare. The clever Dallas - Austin passenger would buy a Dallas - San Antonio fare on a Braniff Dallas - Austin - San Antonio flight and get off in Austin. It is rumored that Braniff crew were instructed not to allow such passengers off the aircraft, nor would they check luggage only to Austin. How nasty can you be? Needless to say, Braniff never recommended to its Dallas - Austin passengers that they buy the lower Dallas - San Antonio fare!

Special promotional fares make optimal fare calculations even more difficult. If there are special fares to India, they could be used on a trip from the U.S.A. to Sri Lanka, for instance, instead of using the regular U.S.A. to Sri Lanka fares.

Switching now to cases of additional information volunteered by the QAS, here is another example. While I.A.T.A. carriers allow no stopovers on their 22 - 45 day excursions on New York to Glasgow, the airline Loftleidir allows a stopover in Reykjavik. Moreover, Loftleidir offers a very inexpensive hotel/meals/excursions package for such a stopover. It is doubtful that a user would know and therefore ask about the existence of the stopover privileges or the package. The system must volunteer the information. Of course, the suggestions should not be out of place. On a New York to Glasgow itinerary, it is meaningful to mention the stopover and package in Reykjavik, but not stopover privileges in Bombay.

Let us mention finally a last interesting aspect of travel schedules. On excursion rates, travel must begin back within a minimum and maximum number of days, for example 7 to 28 days on some excursion plans between the U.S.A. and some parts of South America. It would be possible for someone to increase his stay in South America by choosing on his return an itinerary through smaller cities having infrequent air connections. He could leave on day 28 from city A, fly to B. Assuming that the next flight to the next stop, city C, does not occur for five days (he arrived Wednesday, and flights from B to C operate only on Mondays), he may legally stay in B for that time, thereby extending his trip. He might even gamble and perhaps receive a free vacation in city B as follows: he may have a legal but tight connection in B (i.e. his arrival in B may be 65 minutes before the scheduled departure from B to C, and the legal connection time is 60 minutes). If, as is often the case, the flight from A to B is late and if the flight from B to C leaves on time, (which is not often the case, unfortunately!) he has missed his connection and the airline flying him from A to B is responsible for his hotel room and board for the five days until the next flight (he need not accept any alternate routings). Recommended only for those who have large amounts of time on their hands!

APPENDIX B. FREIGHT TRANSPORTATION

The freight rates in the United States may even be more complex than the airfares. One can ship goods (depending on their size and nature) by the Postal Service, United Parcel Service, various airlines, air express shipping companies, barges, trucks, trains, etc. It is not unusual for a package shipped from Oregon to arrive in Austin, Texas, bearing a Postal Service stamp from Dallas: it was somehow shipped from Oregon to Dallas by some means (but not the Postal Service) and then sent from Dallas to Austin by the Postal Service.

The shipping rates are again highly non-linear. Only a few miles might separate two towns, but goods sent from say New York to the two towns by the Postal Service may be billed at quite different rates if the two towns are not in the same postal zone (as viewed from New York). Sending a truck that is 90% full costs essentially the same as one which is 100% full, and therefore if the shipper can wait and send another 10% in the same truck, he sends these 10% at almost no cost.

Convenience, speed, security and lack of damage are also considerations in shipping. During transfers, time may be lost, merchandise set to the wrong place or damaged. As an example, in about 1972 the French newspaper Le Monde changed its shipping plans for sending its weekly condensation to readers residing in the United States - a fairly sizable readership. Previously, the copies were sent individually by air mail from Paris to their destination. In the new system, the copies were all shipped by air cargo to New York and then shipped from New York by the U.S. Post Office using second class mail rates. The total costs were lower, but the second class service was found to be impossibly slow. So many readers complained that the initial service - air mail from Paris - was reinstated.

APPENDIX C. INCOME TAXES.

The income tax regulations constitute another very large data base with non-linear algorithms. Perhaps best known is the capital gains rule. If a stock is sold at a profit, the full profit is taxable if the sale occurs within six months* of the date of purchase. If the sale occurs six months* plus a day, or longer, after the date of purchase, only half of the profit is taxed. A child's birth on December 31 of year NNN will result in a deduction for year NNN to the parents, who would not receive this deduction if the child were born one day later on January 1 of year (NNN + 1).

The income tax preparation industry is perhaps best tuned to its data bases and algorithms. Tax attorneys and lawyers are available (but are usually affordable only by the rich) who will help suggest changes to the inputs to the tax computation algorithms so as to reduce taxes. Stock portfolios are often changed towards the end of a year for just such purposes. In addition, the income tax algorithm often offers several alternate branches. It is not always easy to find the best branch which will minimize taxes. Again, tax lawyers and accountants apply their knowledge of tax laws to try and select the best branch which will result in the lowest tax.

If a professional (teacher, lawyer, etc.) buys a car in year 19NM, he can depreciate part of the value of the car during 19NM (by more than 60% according to some calculations). If the car is used R% for business, that fraction of the depreciation is tax deductible. One can make R equal to 100% by buying the car very late (say December 29, 19NM) and driving it purely for professional purposes for example to a professional meeting. So here is the scenario. The car is bought December 29, 19NM, and departure to the meeting is immediate. Return from the meeting is next year. Then in 19NM, the car was used 100% for business, and the full allowable depreciation can be applied (times 50%, the car being bought during the second half of 19NM). Notice how non-linear the algorithm is. If the car is bought December 28, 19NM and the trip is started December 29, then the car was used 3 days for business (December 29, 30 and 31) out of four days of possession, and the fraction used for business is R = 75% only, instead of 100% as before. (Warning: the Internal Revenue Service may harass you if you

---

*Before 1976. In 1977, the length is nine months, another example of the volatility of the data base!

follow the law too closely, as described above!)

APPENDIX D. HOUSING DESIGN AND CONSTRUCTION.

In the design process, be it of houses or machinery, etc., the designer must simultaneously satisfy a number of constraints (building codes, material strengths, etc.) while at the same time optimizing some factors (long durability, low cost, etc.). The constraints result again in a highly non-linear data base. The installation of a central humidification system in a house may be simple if the house is to be heated by a fanned-air heating system. It would be very costly if the house were heated by radiators. The addition of a bathroom in a multi-storied house may be relatively inexpensive if the new bathroom is to be situated over an already existing bathroom, and the water pressure is adequate. Otherwise, the cost of the new bathroom may be prohibitive. Building codes often specify minimal support structures for a building. A slightly larger building may necessitate a new set of support elements, making the marginal cost of the small size increase very high. Similarly, costs often jump when one switches from one grade of material (for instance, considering stress capabilities) to the next higher grade. Small changes in the design of a building may necessitate shifting from one grade of material to the next higher grade, again resulting in large additional costs.

The customer who wants a building (the user of the system, presently made up typically of a team of architects) has again some freedom in the design constraints (the inputs to the system) that the building must satisfy. He will welcome suggestions for changes, modifications, etc. to his original constraints, while at the same time he would want a good problem-solver to satisfy his design constraints.