SDC
System Development Corporation
A Burroughs Company

Deductively Augmented Data Management: An Example of Its Use

Antonio Leal

TM-(L)-7328/000/00

30 March, 1983

SDC
System Development Corporation
A Burroughs Company

Deductively Augmented Data Management: An Example of Its Use

Antonio Leal

TM-(L)-7328/000/00

30 March, 1983

CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

The Deductively Augmented Data Management (DADM) System is the result of a research and development program carried out at System Development Corporation (SDC) in the area of deductive inference and data management. Charles Kellogg has been project leader since the inception of DADM. Others contributing to the design, implementation, and application are Dr. Larry Travis (consultant from the University of Wisconsin, Madison, Wisconsin); David Bosley, Dr. Darrel Van Buer, Dan Kogan and Richard Whitney (at SDC); and the author, Dr. Antonio Leal, a consultant from Tetrax Corporation, Los Angeles, California.

1. INTRODUCTION

The purpose of a Data Base Management System (DBMS) is to serve as an interface between the information stored in a data base and the customers who wish to access that data. Most DBMS architectures are able to access virtually any subset of the data and come equipped with utility programs for sorting, filtering, formatting, etc. However, the interpretation of the data is left up to the user. That is, the user is expected to analyze relationships among the retrieved data and deduce conclusions about what those relationships mean.

Recent advances in Knowledge Base Systems in the field of Artificial Intelligence have provided automated tools for the interpretation of data over and above its mere retrieval. As a data base stores facts, a knowledge base stores relationships. A data base user must know in advance which subsets of data will be relevant to the intended analysis and ultimate conclusions. In contrast, a knowledge base user is able to generate queries in terms of the desired outcomes and results knowing that all information in the data base relevant to the query will be considered in the formulation of the response. Thus, a Knowledge Base Management System (KBMS) can be thought of as a "front end" to a Data Base Management System which (1) accepts user-generated queries, (2) translates them into a series of data base access queries, and (3) correlates the retrieved data into a relevant response. (See Figure 1).

This report describes the "deductive" approach to Knowledge Base Management. This approach has its foundations in formal mathematics and attempts to "prove" conclusions about the stored data using the rules of deductive logic. The associated knowledge base contains data relationships as well as rules called "premises" which connect one set of relationships to another forming deductive chains of "reasoning". Although the data relationships can be derived from the data base, the premises must be generated with the assistance of experts in the application domain. Thus, Knowledge Base Systems are often called "Expert Systems".

Over the past decade, System Development Corporation has performed research in Knowledge Base Systems and has developed the Deductively Augmented Data Management (DADM) system. DADM has been applied to areas such as Naval task force deployment; corporate managerial decision making; dissemination of scientific information; and, as this report describes, counter-terrorism. DADM is implemented on a Xerox 1100 scientific processor (Lisp

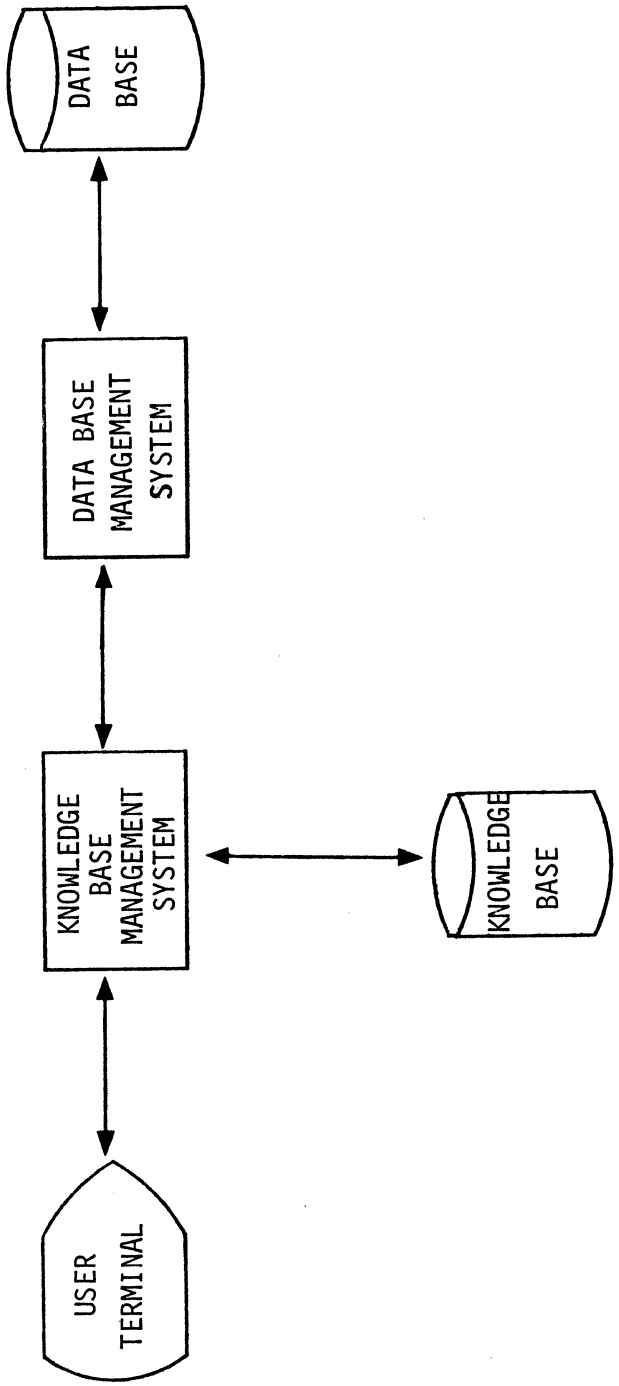

Figure 1. Knowledge Base System

Machine). DADM uses the built-in window-graphics technology of the Xerox 1100 along with a mouse-driven cursor to provide a user-friendly interface. A user interacts with the system via the Xerox 1100 display, keyboard, and mouse. (See Figure 2). The DADM deductive processor uses the information stored in its knowledge base (KB) to create low-level database (DB) access strategies from high-level user requests. They are sent to a distributor module that automatically routes retrieval commands to either the local relational data management system (RDMS), or to an external data management system such as a Britton-Lee IDM-500 (a relational database machine). In the latter case, data returned from the IDM-500 are automatically combined (aggregated) with data obtained locally. The data are then used by DADM to construct answers and explanations. Lisp procedures, to be applied automatically by DADM before or after database search, are stored in a local procedure library. DADM can remotely access one or more of the up to 50 databases that may exist within an IDM-500. The VAX 11/780 is used to obtain hardcopy printouts of the graphics screen with a Versatec V-80 printer.

Throughout this explanation of deductive processing, and of DADM in particular, one sample application will be developed from the domain of counter-terrorism. This unclassified example was selected because of its similarity to real-world problems. The test data base was constructed both to simulate realistic analysis conditions and to demonstrate the power of a Knowledge Base Management System such as DADM. It is assumed that a small data base exists which contains information about suspected terrorists and terrorist organizations, their use of weapons and methods, their locations, and some historical data on past terrorist incidents and events. The data assembled for this example is contained in Section 7. However, the sample data base is for illustration purposes only. Many of the names, organizations, and events have been fabricated to support the examples of logical deduction. It should not be considered factual or representing the activities of actual persons or organizations. Furthermore, in a full application of DADM, the data base would be considerably larger.

The development of the counter-terrorism example begins with a descriptive scenario showing the use of DADM to answer example queries (Section 2). The definition of the domains of discourse for the model is given in Section 3. The domains categorize the contents of the data base and provide names for referencing sets of objects. These domains are then connected by groups of "relations" (Section 4). The relations are of two basic types. Some of them correspond to information stored directly in the

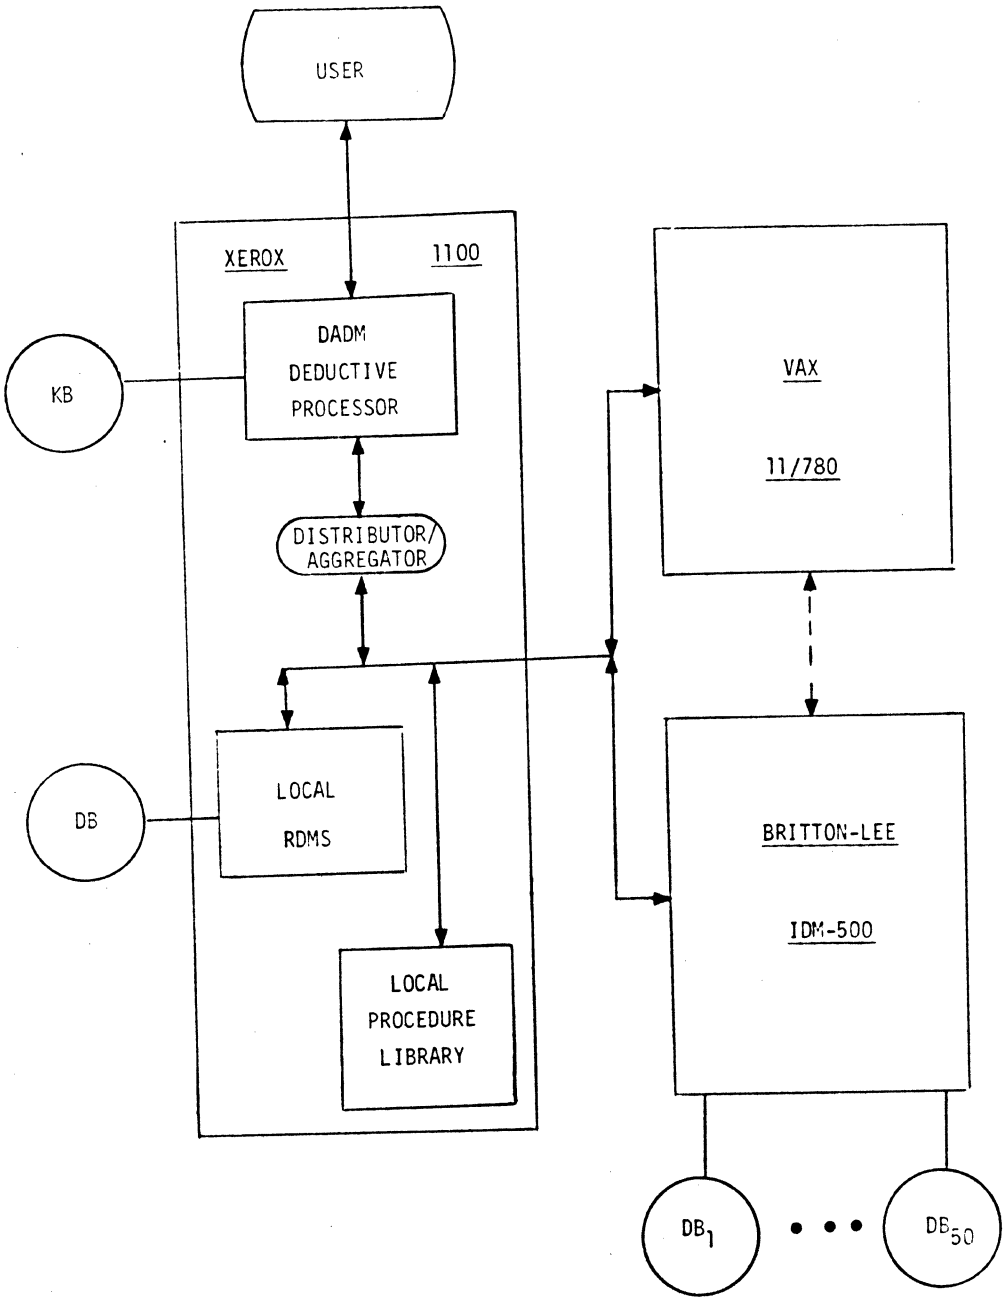

Figure 2. DADM Application Development Environment

data base. Others, however, are used to deduce new information. The precise manner in which this new information is obtained is captured and defined in the set of "premises" (Section 5). Premises are rules that permit conclusions to be made using deductive logic (Section 6). The contents of the example counter-terrorist data base are presented in Section 7 and Section 8 contains the summary and conclusions including a description of how DADM can be applied in a real environment.

2. APPLICATION

The application of DADM to the area of counter-terrorism provides the analyst with a powerful deductive tool for discovering new relationships among stored data base facts about terrorist organizations and individuals. Questions presented to DADM are translated into a series of data retrievals which are connected together in an attempt to deduce the answers. Thus, the questions may involve the use of high level concepts which are not part of the data base at all. These high-level concepts represent relationships among the data items which would be of interest to the analyst.

For example, the current prototype counter-terrorism data base contains information about (1) suspected terrorists; (2) known or suspected terrorist organizations; (3) major cities and countries in which terrorist activities have taken place; (4) various skills and weapons used by terrorist organizations; (5) nationalities that have been victimized in past terrorist events; and (6) a list of previous terrorist incidents with associated information such as the date, place, weapon used, skill used, and the responsible organization if known. The data base also contains information such as (1) the membership of individuals within organizations, (2) the location of cities within countries, (3) the base of operations of organizations, (4) the weapons requirements for the various skills (i.e., "aircraft attack" requires "surface-to-air missiles"), (5) the possession of weapons by individuals or organizations, (6) the intention of one organization to supply weapons to another, plus (7) known organizations that are responsible for past terrorist incidents.

In contrast to these data base facts, DADM contains a knowledge base of rules called "premises" which allow it to make deductions about new relationships. Some of these relationships concern (1) the areas of operation of terrorist organizations, (2) aid and assistance one organization can give another, (3) protection organizations can give to individuals, (4) transfer of skill and weapons knowledge from one organization to another, (5) weapons supply networks among terrorist organizations, and (6) the possibility that a particular organization is responsible for a terrorist incident where such responsibility was not previously known.

An example premise used in formulating deductions would be the following. "If one organization supplies a specific weapon to another organization, then the latter must possess that weapon." This type of rule could be used in deducing weapons supply

networks and lead to the conclusion that a particular terrorist organization possesses a certain weapon where this information was not known before (i.e., was not stored in the data base). By compounding rules, further deductions can be made. For example, the discovery that a particular organization possesses a critical weapon may contribute evidence to concluding that the organization may be responsible for a past terrorist incident.

The following examples show how DADM can be used to answer questions that require logical deduction. Figures 3-8 depict the actual operation of DADM and are taken directly from the Xerox 1100 display by means of the Versatec V-80 printer.

It is assumed that an analyst wishes to find organizations that possess surface-to-air missiles (SAMs). Thus, the analyst would type in the following statement.

(FIND ORG-1 POSSESSES SAMS)

Figure 3 shows the entry of this query. "ORG-1" is a variable name which stands for any organization in the data base that satisfies the request. After processing, DADM produces an "inference plan" that is used to derive an answer. The inference plan is composed of a connected combination of logical rules (premises). Since the premises can be combined in many different ways, a number of different plans could be created in the attempt to answer the question. In Figure 3, plan number 2 produced the successful result. The response "Yes" is then entered to the DADM prompt "Answers?" and data base search (and possibly compute) requests are displayed. A "search request" is a data base query generated by DADM that retrieves the necessary information to answer the user question. Then, DADM responds with "PLO", shown in the "ANSWER SUMMARY".

How did DADM arrive at this conclusion? The search listing for the data base is not enough to explain the deductive reasoning process used to answer the question. Thus, DADM prompts the analyst with "Explain?" asking if more detailed information is desired. A "Yes" response produces the evidence display shown in Figure 4. From this display, and the accompanying graphical representation, it can be seen that the PLO is receiving SAMs supplied by Libya.

The analyst can also ask DADM to display the rules or "premises" that were used in formulating the conclusion. Figure 5 shows one of these premises. This premise says that if a person or

```
MODE:  Query:
./FIND ORG-1 POSSESSES SAMS)

Query analyzed on 25-MAR-83 13:44:09
```

```
PLAN NUMBER: 2
2 search problems.
Answers?  Yes
              SEARCH/COMPUTE PLAN:
    SEARCH       *WOULD-SUPPLY PERSON-ORGANIZATION-1 ORG-1 SAMS
    SEARCH       *OWNS PERSON-ORGANIZATION-1 SAMS

ENTERING DATA BASE

DATA-BASE SEARCH SUCCESSFUL

***************
ANSWER SUMMARY --
VARIABLES:
(ORG-1)
ANSWERS:
(PLO)
***************
```

Figure 3. Which Terrorist Organizations Possess SAMS?

```
Explain?    Yes

EVIDENCE CHAIN  FROM PLAN 2

====================


  **2
  FACT libya owns sams
  CONCLUDE libya possesses sams

 **1
 FACT libya would supply sams to plo
 CONCLUDE libya supplies sams to plo

**0
CONCLUDE plo possesses sams

====================


                                                    >>
```

Figure 4. Facts and Conclusions Display

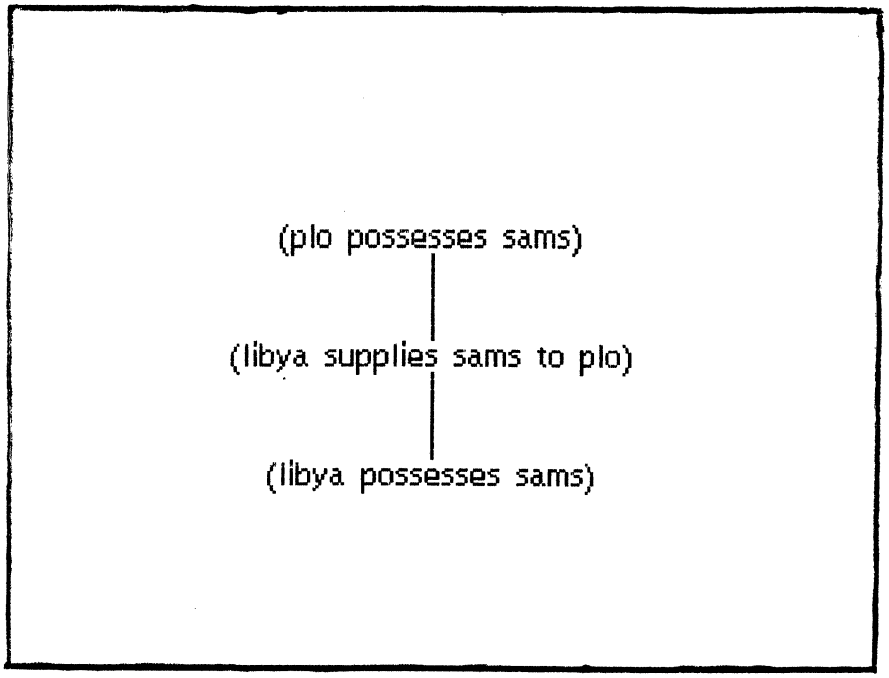

Figure 4. Facts and Conclusions Display (Cont'd)

```
(PREMISE FOR (SUPPLIES (LIBYA PLO SAMS)) WT 88)
If person-organization-1 would supply weapon-1 to
                          person-organization-2
                             and
                          person-organization-1
                                 possesses weapon-1

then person-organization-1 supplies weapon-1 to
                          person-organization-2
```

Figure 5. A Sample Premise

organization has the intention to supply weapons ("would-supply") and it "possesses" the weapons, then it "will" supply them.

Premises, used in combination with data base facts, allow DADM to find new facts and relationships toward answering the initial query. The relations use English words ("would supply", "will supply", etc.) in a technical sense to describe connections among data base objects (in this case, organizations and weapons). However, they do not necessarily carry all of the connotations of their English counterparts. The only real definitions that the relations have are those contained in the listing of the associated objects in the data base and those supplied by the rules. For example, in the current small demonstration application, the "possession" of a weapon means that the organization could "use" the weapon in a terrorist action or "could supply" the weapon to another organization. In contrast, the relation "would supply" denotes a willingness to supply the weapon to another organization if that weapon was possessed. Combining these two notions, a premise can be written which says, "If organization A 'possesses' a weapon, and if organization A 'would supply' the weapon to organization B, then organization A 'will supply' it to organization B." This implication of a definite action can be made only within the context of the currently defined premise set. In other words, the premise set constructed for the examples described in this report is not complete enough to take into consideration all of the other factors that could influence the decision of one organization to supply weapons to another. In real applications of knowledge base systems, four or five hundred premises (rules) are usually required to describe the environment adequately. The rules are abstracted, synthesized, and approved by experts in the field and are always updated periodically to reflect changing situations, directives, procedures, or scope of application.

The next series of example displays from DADM (Figures 6-8) shows a much more complicated deduction. The analyst is interested in finding an organization that is "probably responsible" for a particular terrorist incident: I20. The incidents are stored in the data base under reference numbers (I1, I2,...) and information about the date, place, weapon used, etc., is known about incident number I20 but not which organization was responsible. Figure 6 shows the initial entry of the query.

(FIND ORG-1 IS PROBABLY RESPONSIBLE FOR I20)

```
Query:
.(FIND ORG-1 IS PROBABLY RESPONSIBLE FOR I20)

Query analyzed on 30-DEC-82 13:42:50

PLAN NUMBER: 1
14 search problems.
Answers?  Yes
              SEARCH/COMPUTE PLAN:
     SEARCH       *WAS-USED-IN SKILL-1 I20
     SEARCH       *HAS-KNOWLEDGE-OF ORG-1 SKILL-1
     SEARCH       *REQUIRES SKILL-1 WEAPON-2
     SEARCH       *OWNS ORG-1 WEAPON-2
     SEARCH       *VICTIMS-OF NATIONALITY-1 I20
     SEARCH       *VICTIMS-OF NATIONALITY-1 INCIDENT-1
     SEARCH       *IS-RESPONSIBLE-FOR ORG-1 INCIDENT-1
     SEARCH       *OCCURRED-IN I20 CITY-1
     SEARCH       *LOCATED-IN CITY-1 COUNTRY-1
     SEARCH       *OCCURRED-IN INCIDENT-2 CITY-2
     SEARCH       *LOCATED-IN CITY-2 COUNTRY-1
     SEARCH       *IS-RESPONSIBLE-FOR ORG-1 INCIDENT-2
     SEARCH       *WAS-USED-IN WEAPON-1 I20
     SEARCH       *OWNS ORG-1 WEAPON-1

ENTERING DATA BASE


DATA-BASE SEARCH UNSUCCESSFUL
More?
```

Figure 6. Facts Unavailable for First Plan

Upon responding with "Yes" to the DADM prompt "Answers?", the analyst can view the data base searches. This time, however, the data base search was unsuccessful. This means that there was insufficient information in the data base to reach a conclusion based on this particular inference plan. However, there is often more than one way to find the answer. Thus, as shown at the bottom of Figure 6, DADM asks "More?" which indicates that further inference plans have been found. A response of "Yes" to this prompt produces a second plan (not shown) which also results in an unsuccessful data base search. However, the third plan is successful. (See Figure 7). Plan number 3 required 16 separate data base access commands involving information concerning skills and weapons as well as the locations of various cities within countries. At the bottom of Figure 7, the "ANSWER SUMMARY" provides the results of the deduction showing that the PLO is probably responsible for incident 120. ("Probably responsible" is used linguistically and does not refer to a numeric value of probability.)

Figure 8 shows the explanatory evidence chain along with the corresponding graphical display. Within the evidence chain are intermediate deductions which lead to the final conclusion. These deductions use a number of related facts from the data base including weapons supply from Libya, the operational range of the PLO as deduced from a different past incident (13), knowledge of how to use bombs supplied by an independent terrorist (Carlos), and other pertinent information. A more detailed explanation of how this particular result is deduced is described in Section 6.

Although the above examples are somewhat simplistic in the sense that the conclusions would probably be known to an experienced analyst anyway, they serve to illustrate the deductive power of DADM and how it can bring together many related data base facts in forming a conclusion. In a large data base, the complexity would easily rise beyond human capability to correlate all relevant facts into a single unified result. It is in this environment that DADM can be of immense assistance to even the expert analyst.

The following sections describe in detail the application of DADM to the counter-terrorism model including a listing of all the premises and a description of how they interact. The experimental data base is also included.

```
More?  Yes
New Inference plan


PLAN NUMBER: 3
16 search problems.
Answers? Yes
             SEARCH/COMPUTE PLAN:
     SEARCH       *WAS-USED-IN SKILL-1 I20
     SEARCH       *TRAINS-FOR PERSON-1 ORG-1
     SEARCH       *HAS-KNOWLEDGE-OF PERSON-1 SKILL-1
     SEARCH       *REQUIRES SKILL-1 WEAPON-2
     SEARCH       *OWNS ORG-1 WEAPON-2
     SEARCH       *VICTIMS-OF NATIONALITY-1 I20
     SEARCH       *VICTIMS-OF NATIONALITY-1 INCIDENT-1
     SEARCH       *IS-RESPONSIBLE-FOR ORG-1 INCIDENT-1
     SEARCH       *OCCURRED-IN I20 CITY-1
     SEARCH       *LOCATED-IN CITY-1 COUNTRY-1
     SEARCH       *OCCURRED-IN INCIDENT-2 CITY-2
     SEARCH       *LOCATED-IN CITY-2 COUNTRY-1
     SEARCH       *IS-RESPONSIBLE-FOR ORG-1 INCIDENT-2
     SEARCH       *WAS-USED-IN WEAPON-1 I20
     SEARCH       *WOULD-SUPPLY PERSON-ORGANIZATION-1 ORG-1
                  WEAPON-1
     SEARCH       *OWNS PERSON-ORGANIZATION-1 WEAPON-1

ENTERING DATA BASE

DATA-BASE SEARCH SUCCESSFUL

**********
ANSWER SUMMARY --
VARIABLES:
(ORG-1)
ANSWERS:
(PLO)
**********

Explain?
```

Figure 7. Successful Database Search

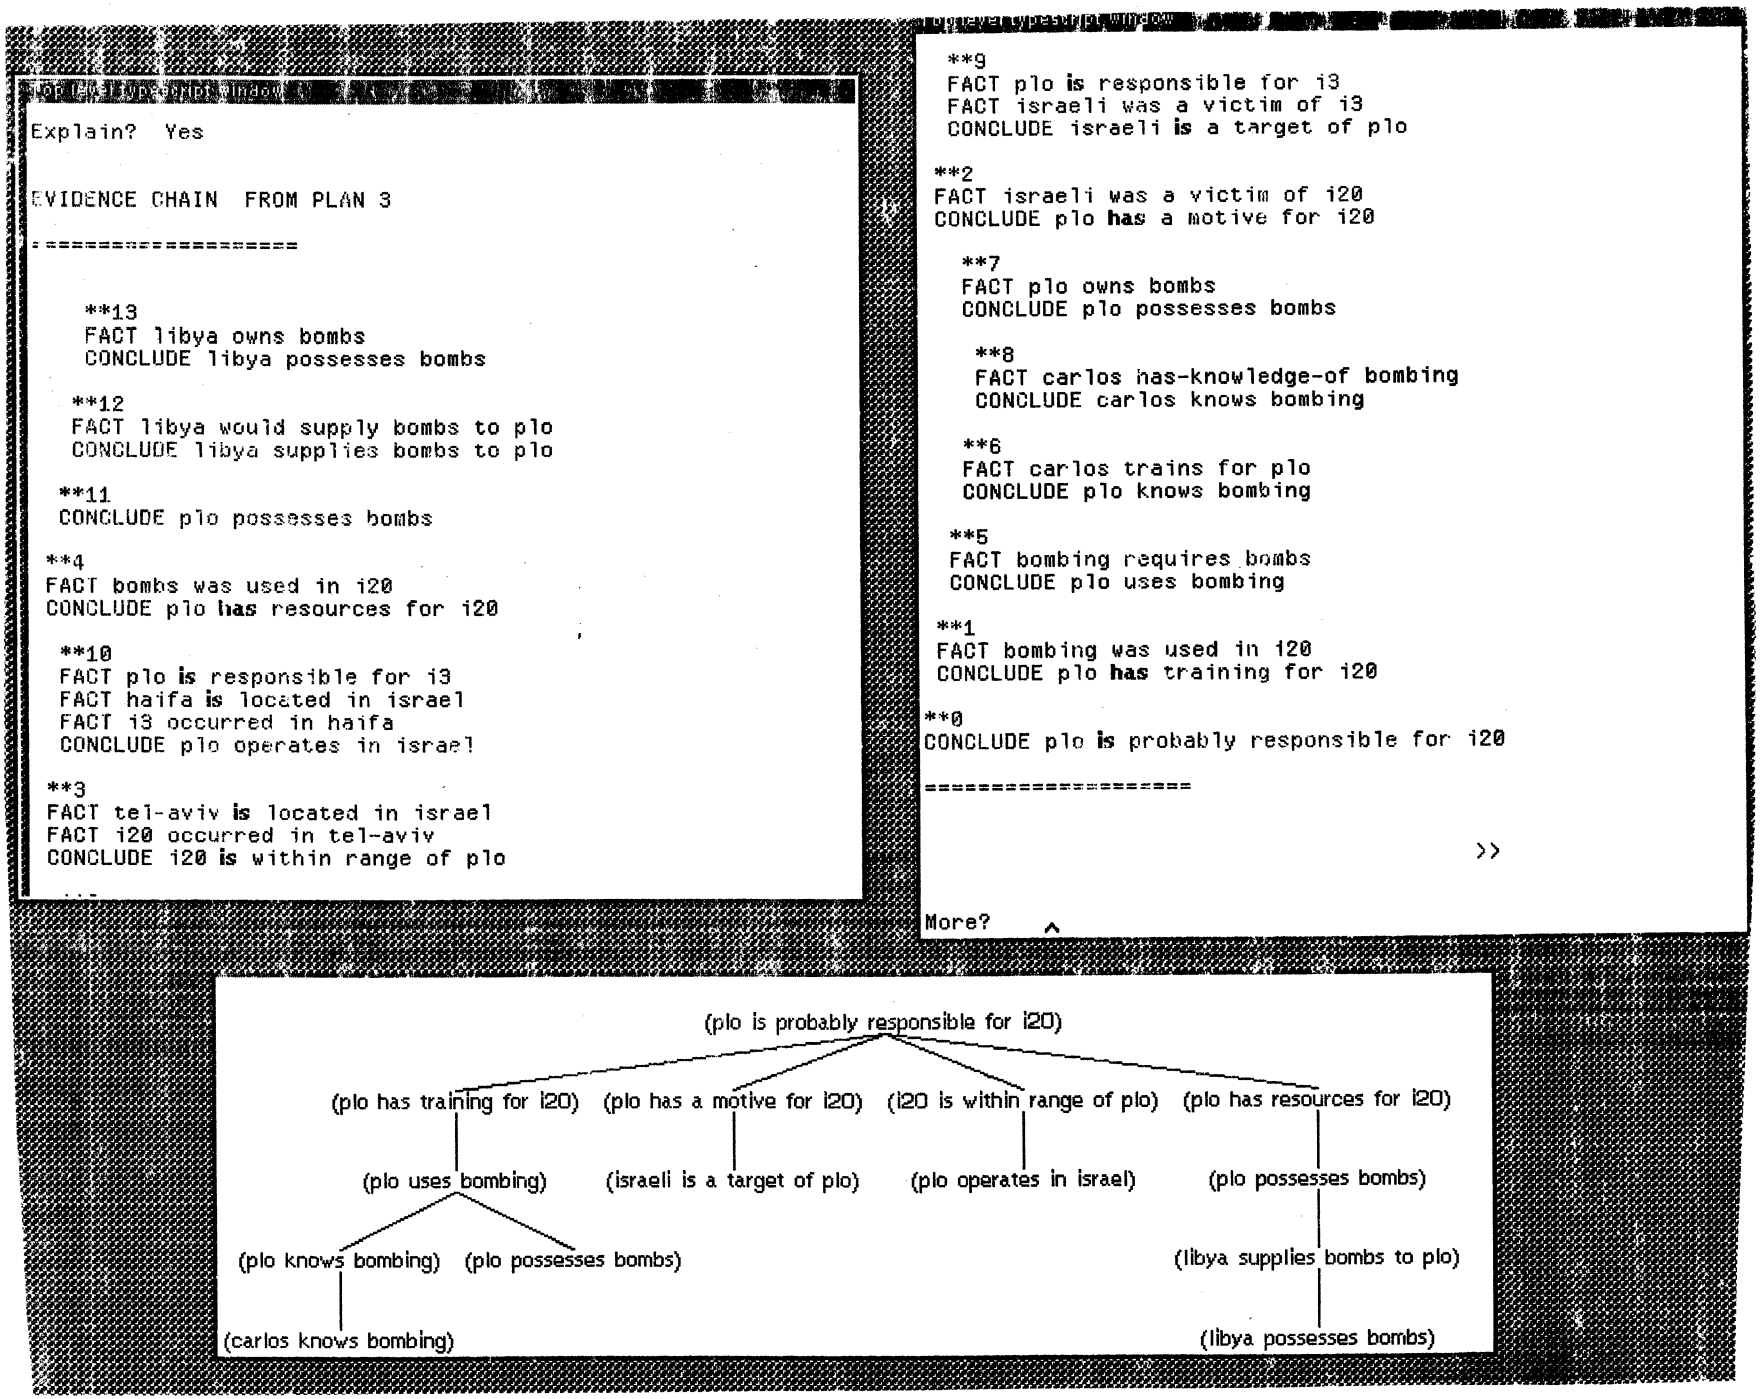

Figure 8. PLO is Probably Responsible for Incident I-20

3. DOMAINS

A "domain" is a set of objects that is associated with the area of application. The objects themselves comprise the contents of the data base and are used in logical deductions. The domain sets will, of course, be different for each application area. For example, in the area of management, typical domains might be employees, projects, departments, salary rates, hourly assignments, budget expenditures, etc. Each domain is thus composed of a collection of objects (names, numbers, codes, etc.) which may appear one or more times in the data base.

For the counter-terrorism example, the following hypothetical domains are defined: persons, organizations, countries, cities, skills, weapons, nationalities, and incidents. Each domain is assigned an abbreviation letter. Since it is possible to combine two or more domains together to make a larger one, the code letters help to identify them.

Persons (P)

The "persons" domain (abbreviation: P) contains the names of known or suspected terrorists.

Nikolai Aksenov	Simone Losada
Yasar Arafat	Abdul Mohammad
Carlos	Mario Moretti
Jock Clancy	Bill O'Toole
Juan Delado	Patrizio Peli
Vallejo Desanchez	Gerado Peraso
Bernadette Devlin	Shalim Sahid
Bonet Erzin	Jose Sanchez
Manual Esponza	Bobby Sands
Fatabyle Fawzial	Tomaso Santini
Otto Frenzel	Richardo Santos
Karl Freund	Antonio Scarini
Jamid Gammal	Olga Schmidt

Organizations (O)

The "organizations" domain (abbreviation: O) contains the names of known or suspected terrorist organizations.

ASALA	(Armenian Secret Army for the Liberation of Armenia)
BJO	(Black June Organization)
ETA	(Basque Fatherland and Liberty)

FMLN	(Farabundo Marti National Liberation Front)
GAP	(Guerrilla Army of the Poor)
MLAPU	(Marxist-Leninist Armed Propaganda Unit)
Montoneros	
M-19	(April 19th Movement)
RAF	(Red Army Faction)
Red Brigade	
PIRA	(Provisional Irish Republican Army)
PFLP	(Popular Front for the Liberation of Palestine)
PLO	(Palestine Liberation Organization)

Countries (C)

The "countries" domain (abbreviation: C) contains the names of world countries in which terrorist organizations are located or in which terrorist activities have taken place. The countries form a single set which can be used for both of these concepts.

Argentina
Belgium
Columbia
Cuba
El Salvador
England
France
Guatemala
Ireland
Israel
Italy
Lebanon
Libya
Mexico
Peru
Russia
Spain
Switzerland
Syria
Turkey
United States
West Germany

Cities (T)

The "cities" domain (abbreviation: T) contains the names of relevant cities located within the listed countries.

Ankara
Bagdad
Beirut
Belfast
Berlin
Bethlehem
Bilbao
Bogota
Brussels
Buenos Aires
Cannes
Lima
Madrid
Manchester
Medellin
Mexico City
Milan
Moscow
New York
Padua
Paris
Rome

Damascus
Frankfort
Haifa
Lemoniz
London
Ramstein
San Salvador
Tel Aviv
Verona

Skills (S)

The "skills" domain (abbreviation: S) contains the names of some of the methods known and used by terrorist organizations.

Aircraft Attack
Armed Attack
Assassination
Bombing
Burglary
Hijacking
Hostage Seizure
Kidnapping
Sabotage
Underwater Demolition

Weapons (W)

The "weapons" domain (abbreviation: W) contains the names of special weapons used by terrorist organizations. It is assumed that all terrorist organizations possess small arms such as pistols, rifles, etc.

Bombs
Booby Traps
Grenades
Machine Guns
Plastic Explosives
Surface-to-Air Missiles (SAMs)

Nationalities (N)

The "nationalities" domain (abbreviation: N) contains the names of groups of people who have been the victims of terrorist attacks.

American
Argentinian
British
Columbian
French

Guatemalan
Israeli
Italian
Lebanese
Palestinian
Peruvian
Salvadoran
Spanish

Incidents (I)

The "incidents" domain (abbreviation: I) contains code numbers identifying documented terrorist incidents. It is actually a collection of smaller sub-domains which contains information about each of the incidents. Except for the identification code and the date, the incident sub-domains borrow items from the other domains listed above. For example, the "place" must be one of the cities listed in the cities domain. Thus, the "place" sub-domain is nothing more than a subset of the cities domain. It consists of all those cities which have been the site of terrorist incidents (up to the limit of the available information in the data base). For a complete listing of the example incidents, see Section 7.

Incident Sub-Domains	Example
Identification Code	19
Place (City)	Belfast
Date	6/5/82
Skill Used	Bombing
Weapon Used	Bombs
Victim (Nationality)	British
Weapon Stolen	None
Responsible Organization	PIRA

It is not necessary for information on every incident to be known. DADM can be used to make deductions concerning the missing information.

Domain Unions

A number of single domains may be combined into a larger domain forming a "domain union". The abbreviations for the large domains are taken from their components.

PO:	person or organization
OC:	organization or country
POC:	person, organization, or country
CT:	country or city
SW:	skill or weapon

The above domains can now be connected to form a set of "relations" which describe various situations and associations among the data base objects.

4. RELATIONS

A "relation" is an association among objects from two or more domains. A single relation may connect objects from the same domain or objects from different domains. For example, the "belongs to" relation would connect persons with the organization of which they are members. Each relation, then, has two or more "arguments" which are like blanks that must be filled in with names of objects from the proper domain. The "belongs to" relation would be written as follows:

```
 P  belongs to O
---           ---
```

where "P" is the abbreviation for the "persons" domain and stands for the name of a person, and "O" is the abbreviation for the "organizations" domain and stands for the name of an organization. Any person-name and any organization-name can be filled in but the relation is true only if that combination appears in the data base or has been deduced. "True" is not used here in an absolute sense but in a logical sense. That is, a relation is true if it appears in the data base or if it can be deduced by DADM. Consequently, the validity of the relations, and ultimately the conclusions, depend upon the validity of the data base and of the premises.

If a domain union appears as one of the arguments in a relation, then objects from any of the specified domains may be used. For example, the "possesses" relation permits persons, organizations, or countries to possess weapons.

```
 POC  possesses  W
-----           ---
```

DADM requires each relation to be classified in one of three types: base, deduced, or computed. A "base" relation is one that appears in the data base. The combinations of objects that make it true can always be looked up. On the other hand, a "deduced" relation does not appear in the data base. DADM attempts to find "true" instances of deduced relations by logical chains of reasoning using the available base relations and instances of other previously deduced relations. A "computed" relation is one that involves mathematical computation in one

form or another. The current counter-terrorist model does not contain any computed relations.

The separation between base and deduced relations assures that the data base will not be altered and will maintain its integrity. That is, DADM does not change the information in the data base in any way. The trade-off for this integrity is that some relations must have dual definitions. For example, while the "possesses" relation is deduced, the "owns" relation is base and they both have the same meaning. Such a distinction is necessary in those cases where DADM is to (1) use information from a stored data base relation and (2) deduce information about the same relation. Referring to the hypothetical data base in Section 7, it is assumed that certain organizations are known to have specific weapons. This information is stored under the "owns" relation (i.e., the PIRA owns grenades). Because of the premises dealing with supply networks, DADM could deduce similar information about other organizations where the new information is not stored in the data base (i.e., the Red Brigade owns grenades). However, DADM is not allowed to change or update the data base. Therefore, it cannot store the newly deduced information under the "owns" relation. A complimentary relation called "possesses" has been defined to accommodate this new information. It is then only necessary to create one linking premise: "If organization A owns weapon W, then organization A possesses weapon W." Now, all other premises can refer to "possesses" with the confidence that both stored and deduced information will be available. The "has knowledge of" and "knows" relations are defined in the same manner.

With respect to relations, DADM is capable of operating with either one of two distinct "world views". If an instance of a relation appears in the data base, it is considered "true". For example,

Patrizio Peli belongs to the Red Brigade

The above "fact" appears in the data base. (See Section 7.) Thus, DADM would respond with "Yes" to the query, "Does Patrizio Peli belong to the Red Brigade?" The difference in world views occurs with instances of relations that do not appear in the data base. For example, the following relation is not contained in the sample data base.

Patrizio Peli belongs to the PLO
---------- -----

The "closed world" view assumes that any instance of a relation that is not in the data base is "false". In other words, because "Patrizio Peli belongs to the PLO" is not in the data base, the closed world response would be "no" to the query, "Does Patrizio Peli belong to the PLO?" Thus, the closed world view gives a sense of completeness to the data base. However, under the "open world" view the system would respond, "I don't know", to this query.

Although DADM is capable of using either world view, it usually operates under the "open world" assumption. The decision to use open or closed world can be made individually for each relation. In the example data base, all relations are assumed to be "open world".

The remainder of this section lists all of the relations for the example counter-terrorism model. Each relation is categorized as "base" or "deduced" along with a template specifying the applicable domains and their positions. The combinations of objects in the data base for each base relation can be found in Section 7.

Belongs to (base)

The "belongs to" relation specifies membership in organizations. A single person can belong to more than one organization, and, of course, organizations can have more than one member.

P belongs to O
--- ---

Is located in (base)

The "is located in" relation applies strictly to the location of cities within countries.

T is located in C
--- ---

Is based in (base)

Organizations are assumed to be "based" in a particular city or country. Each organization can be based in only one place. The activities of organizations in other cities and countries are defined using the "operates in" and the "within range of" relations.

O is based in CT
--- ----

Operates in (deduced)

If an organization has conducted terrorist activities in a particular city or country (or is based there), it is said to "operate" there.

O operates in CT
--- ----

Aids (deduced)

Organizations can "aid" each other by supplying weapons, etc. Countries can also aid organizations.

OC aids O
---- ---

Assists (deduced)

"Assisting" is a weak form of aiding. For example, supplying knowledge about skills constitutes assistance.

OC assists O
---- ---

Sends advisors to (base)

Countries or organizations can "send advisors to" other

organizations for the purpose of training in the use of skills and weapons.

OC sends advisors to O

Would protect (deduced)

Countries or organizations can provide sanctuary for individuals or even whole organizations.

OC would protect PO

Trains for (base)

The "trains for" relation refers to an individual who belongs to one organization and provides training for members of another organization.

P trains for O

Knows (deduced)

The "knows" relation refers to the possession of knowledge about how to use particular skills. It is assumed that an organization cannot use a skill unless the knowledge of how to use the skill has been acquired. The current example set of relations and premises does not cover "specialists" within an organization. Consequently, the organization as a whole is permitted to "know" skills and "possess" weapons independently from its members. Premises are defined which transfer knowledge and weapons among members. While this approach is not the best, it is adequate for a small model and does provide a sense of existence to an entire organization which can then be held accountable for the actions of its members.

PO knows S

Has Knowledge of (base)

This relation is the data base counterpart of the deduced "knows" relation. The two are synonymous except for their origin. The "has knowledge of" relation allows skill knowledge to be stored in the data base and associated with organizations known to possess it. The "knows" relation (see above) allows DADM to deduce new information about transferred skill knowledge.

PO has knowledge of S

Uses (deduced)

A person or organization "uses" skills and weapons to carry out terrorist activities. Their ability to conduct terrorism depends upon their resources.

PO uses SW

Requires (base)

Certain skills "require" specialized weapons. For example, an aircraft attack requires surface-to-air missiles. Although most terrorist activities can be conducted with a number of different weapons, for purposes of this example, each major skill will be associated with only one special weapon. (See Section 7).

S requires W

Needs small arms only (base)

This relation has only one argument and specifies those skills that do not require special weapons. A single-argument relation is called a "predicate" and is "true" only for those objects actually listed in the data base. This relation lists those skills not covered under the "requires" relation.

S needs small arms only

Owns (base)

"Owns" is a base relation that provides data on the known weapon resources of organizations, countries, or persons.

POC owns W
----- ---

Possesses (deduced)

The "possesses" relation has the same meaning as "owns" except that it is deduced. This separation permits new information about weapons resources to be discovered through logical deductions concerning weapons supply networks.

POC possesses W
----- ---

Produces (base)

This relation lists countries that "produce" weapons.

C produces W
--- ---

Would supply (base)

The "would supply" relation shows the intention or desire of a person, organization, or country to supply weapons to another. Since the relation is weapon-specific, it has three arguments.

POC would supply W to POC
----- --- -----

Supplies (deduced)

This relation specifies an actual weapons supply agreement that is continuous over time. It is also weapon-specific.

POC supplies W to POC

Are targets of (deduced)

Nationalities that are victims of a particular terrorist incident are assumed to be "targets" of the responsible organization. Such information helps in locating the organizations responsible for other incidents.

N are targets of O

Occurred in (base)

The "occurred in" relation refers to the location of a particular terrorist incident stored in the data base. This relation refers to the "place" sub-domain of incidents.

I occurred in T

Was used in (base)

This relation allows access to information in the data base about particular skills or weapons used in past terrorist incidents. It refers to the "skill used" or "weapon used" sub-domain of incidents.

SW was used in I

Were stolen during (base)

Weapons that were stolen during a particular terrorist incident (such as burglary) can be used for further terrorist activities.

Thus, premises link the weapon stolen to the possession of the weapon.

W were stolen during I

Were victims of (base)

Nationalities that were victims of particular terrorist events can be referenced using this relation.

N were victims of I

Has resources for (deduced)

A number of different pieces of information are needed to deduce the organization that is possibly responsible for an incident. The first component is the possession of required resources.

O has resources for I

Has training for (deduced)

The second major component for deducing responsibility is the training necessary for carrying out specific terrorist activities.

O has training for I

Has motive for (deduced)

The third major component is motive.

O has motive for I

Within range of (deduced)

Finally, the operational range of a terrorist organization provides evidence for potential responsibility of particular events.

0 is within range of I

Could be responsible for (deduced)

Even with only partial or incomplete information, DADM can formulate a list of organizations that could be responsible for a particular incident.

0 could be responsible for I

Is probably responsible for (deduced)

This relation is used to find the organizations which are most likely to be responsible for a particular incident.

0 is probably responsible for I

Is responsible for (base)

This relation is not deduced. It is simply a way of accessing the data base in order to obtain information about past incidents for which responsibility is known.

0 is responsible for I

The use of the word "probably" in the relation "is probably responsible for" does not refer to mathematical probability. DADM operates by finding objects in the data base that satisfy some specified criteria. However, DADM does not impose an order on those objects nor does it assign individual numeric values of

comparison. Thus, for example, in answer to the query, "What organization is probably responsible for incident 120?," DADM could respond with more than one organization, each one satisfying the set of premises that lead to the relation "is probably responsible for."

DADM can, however, provide a degree of "plausibility" on the answer as a whole. Each premise can be assigned a numeric plausibility or "degree of belief." When DADM processes a query, these plausibility factors are mathematically combined to produce a degree of belief on the final conclusion. Plausibility assignments have not been made on the premises for the counter-terrorism model.

5. PREMISES

The premises are the heart of the DADM logical deduction system. They constitute the knowledge necessary to draw conclusions above and beyond the static facts formed by the relations in the data base. Each premise describes in detail the information necessary to conclude something new--a "true" fact that is not stored in the data base. Since many premises are based on human experience and intuition, they are open to argument and controversy. However, once a consistent set of premises has been agreed upon, all conclusions reached by DADM will be valid within that set. In other words, DADM proves its results within the defined system. Thus, if DADM reaches a conclusion that is obviously not true, either the data base is in error or the premise set must be modified.

A "premise" is a logical statement and is composed of two major parts: an antecedent and a consequent. The antecedent, or "if" part, is made up of a collection of relations connected by logical words such as "and", "or", "not", etc. The consequent, or "then" part, is usually a single relation. The logical rule that associates the two parts is as follows.

> If the antecedent relations are true,
> then the consequent relation is true.

Here is an example.

```
 IF  O  is responsible for  I
    ---                     ---
AND  I  occurred in  T
    ---             ---
AND  T  is located in  C
    ---              ---
THEN O  operates in  C
    ---             ---
```

This premise is general because the domain blanks have not yet been filled in with actual objects. The above premise says, "If an organization is responsible for a particular past incident, and the incident occurred in a particular city, and that city is located in a certain country, then the organization operates in that country." For example, if (1) "O" is replaced by "PIRA", (2)

"I" is replaced by incident code "19", (3) "T" is replaced by "Belfast", and (4) "C" is replaced by "Ireland", the premise could be used to deduce that the PIRA operates in Ireland.

```
  IF  the PIRA  is responsible for  19
     ----------                     ----
AND  19  occurred in  Belfast
     ----             ---------
AND  Belfast  is located in  Ireland
     ---------              ---------
THEN  the PIRA  operates in  Ireland
     ----------             ---------
```

When the domain blanks have been filled in, the premise is said to be "instantiated". Since, in the above case, all of the antecedent relations occur in the data base, the conclusion can be considered valid. Each of the three antecedent relations is "base". That is, each can be looked up directly in the data base. However, this need not always be the case. The consequent of one premise could be used in the antecedent of another premise forming "chains" of deductions. The final conclusion may be derived by a long deductive path that accesses many different and seemingly un-related data base facts. Section 6 describes how these deductive chains work.

Notice, in the general statement of the example premise above, that like domain designators refer to the same object. For example, the "T" in "I occurred in T" refers to the same city as the "T" in "T is located in C", etc. Sometimes, the same domain designator must appear in a premise and refer to two or more different objects. In this case, a numeric tag (subscript) is appended to the domain abbreviation for clarity. For example,

```
 POC-1  supplies  W  to  POC-2
-------           ---   -------
```

The attached digit simply keeps the two distinct "POC" objects from being confused when they are used in other relations within the same premise.

This section lists all of the premises for the example counter-terrorism application. Some of them will appear obvious and even trivial. However, these premises are absolutely necessary for

completing logical chains of reasoning. Much of the intelligence of humans derives from background information gained through experience. This background knowledge is usually taken for granted but it must be explicitly defined in a computer model expected to perform logical reasoning. The set was chosen for the purpose of demonstrating the deductive power of DADM and not necessarily to be a definitive statement about the characteristics of terrorist organizations. Consequently, the premises should be taken as a preliminary model of the application environment.

The premise set is presented in more or less a "bottom-up" order. That is, simple premises dealing with "aiding" and "assisting" are listed first, followed by those concerning "knowing" and "possessing". The highest level premises, concerning "responsibility", are listed last. These are farthest from the data base in the sense that they require many previously defined premises to complete their meaning. Premises are always linked via deduced relations. Thus, in order to follow a chain backwards from a high-level premise to the data base, the antecedent relations in the high-level premise must be matched with their occurrences in the consequent position of lower-level premises.

Premise 1. "Assisting" is a weak form of "aiding".

IF OC-1 aids O-2

THEN OC-1 assists O-2

Premise 2. Assistance is transitive.

IF OC-1 assists O-2

AND O-2 assists O-3

THEN OC-1 assists O-3

Premise 3. "Sending advisors" constitutes "aiding".

IF OC-1 sends advisors to O-2

THEN OC-1 aids O-2

Premise 4. Organizations tend to protect their own members.

IF P belongs to O

THEN O would protect P

Premise 5. Providing sanctuary is a form of assistance.

IF OC-1 would protect O-2

THEN OC-1 assists O-2

Premise 6. Protecting an organization means protecting any of its members.

IF OC-1 would protect O-2

AND P belongs to O-2

THEN OC-1 would protect P

Premise 7. If a country or organization aids another organization, then it would protect that organization.

IF OC-1 aids O-2

THEN OC-1 would protect O-2

Premise 8. Sending a trainer constitutes aid.

IF P trains for O-1

AND P belongs to O-2

THEN O-2 aids O-1

Premise 9. A person belonging to a terrorist organization knows everything that the organization knows. (Premises 9 and 10 permit the transfer of skill knowledge from one person to another by way of the "organizational knowledge".)

IF P belongs to O

AND O knows S

THEN P knows S

Premise 10. An organization knows everything that any of its members know.

IF P belongs to O

AND P knows S

THEN O knows S

Premise 11. A person belonging to a terrorist organization will use any skill that the organization as a whole uses. (This glosses over the concept of "specialists".)

IF P belongs to O

AND O uses S

THEN P uses S

Premise 12. An organization will, as a whole, use any skill used by its members.

IF P belongs to O

AND P uses S

THEN O uses S

Premise 13. Assistance implies the transmission of knowledge. (Again, due to the simplicity of the premise set, subtle relationships are not accounted for. This premise does not provide for the situation in which an organization hires an individual to perform a special, one-time action.)

IF OC-1 knows S

AND OC-1 assists O-2

THEN O-2 knows S

Premise 14. Obviously, using a skill implies knowing how to use it.

IF PO uses S

THEN PO knows S

Premise 15. Organizations learn from their trainers.

IF P trains for O

AND P knows S

THEN O knows S

Premise 16. The use of a skill implies the use of its required special weapon.

IF PO uses S

AND S requires W

THEN PO uses W

Premise 17. "Having knowledge" of a skill implies "knowing" that skill. This premise allows skill knowledge to be stored in the data base (has knowledge of) as well as to be deduced (knows).

IF PO has knowledge of S

THEN PO knows S

Premise 18. Using a skill demands knowledge of the skill as well as its required special weapon.

IF PO knows S

AND S requires W

AND PO possesses W

THEN PO uses S

Premise 19. "Owning" a weapon implies "possessing" it.

IF POC owns W

THEN POC possesses W

Premise 20. Terrorist organizations which know skills that demand only small arms can use those skills.

IF PO knows S

AND S needs small arms only

THEN PO uses S

Premise 21. Countries which produce weapons would also possess them.

IF C produces W

THEN C possesses W

Premise 22. Terrorists can obtain weapons from the country in which they are based. This premise simply means that organizations based in a particular country have convenient access to weapons produced in or possessed by that country. However, it does not take into consideration illegally based organizations that may not have access to special weapons such as, for example, nuclear weapons. A more detailed collection of related premises is required to accommodate this type of concept.

IF P is based in T

AND T is located in C

AND C possesses W

THEN P possesses W

Premise 23. Obviously, if persons or organizations use weapons, they must possess them.

IF PO uses W

THEN PO possesses W

Premise 24. Members of an organization possess the weapons that the organization as a whole possesses. ("Possession" here can be thought of as having access to the weapons of the organization.)

IF P belongs to O

AND O possesses W

THEN P possesses W

Premise 25. Organizations possess the weapons possessed by their members. (Like "knowing", this premise permits the transfer of weapons among members of an organization.)

IF P belongs to O

AND P possesses W

THEN O possesses W

Premise 26. If a person, organization, or country has the intention to supply a certain weapon and it possesses that weapon, then it will supply the weapon.

IF POC-1 would supply W to POC-2

AND POC-1 possesses W

THEN POC-1 supplies W to POC-2

Premise 27. Obviously, the supplier of a weapon must possess it.

IF POC-1 supplies W to POC-2

THEN POC-1 possesses W

Premise 28. Continuing from above, the receiver also possesses the weapon. (Premises 27 and 28 imply that the "weapon" transfer is continuing over time. It does not consider the one-time transfer of a single unique special weapon which may be difficult to obtain.)

IF POC-1 supplies W to POC-2

THEN POC-2 possesses W

Premise 29. Supplying weapons constitutes aid.

IF OC-1 supplies W to O-2

THEN OC-1 aids O-2

Premise 30. If a skill or weapon was used in a particular terrorist incident, then the responsible organization must know how to use that skill or weapon.

IF SW was used in I

AND O is responsible for I

THEN O uses SW

Premise 31. Organizations "operate" in those countries in which they have engaged in terrorist activities.

IF O is responsible for I

AND I occurred in T

AND T is located in C

THEN O operates in C

Premise 32. "Targets" of an organization are those groups of people (nationalities) which have been victims in past incidents for which the organization is responsible.

IF N were victims in I

AND O is responsible for I

THEN N are targets of O

Premise 33. Organizations possess the weapons they steal.

IF W were stolen during I

AND O is responsible for I

THEN O possesses W

Premise 34. Knowing skills contributes to an organization's ability to carry out a terrorist activity.

IF S was used in I

AND O uses S

THEN O has training for I

Premise 35. Organizations usually have a motive for carrying out terrorist attacks and the motive usually involves victims.

IF N were victims in I

AND N are targets of O

THEN O has a motive for I

Premise 36. The "range" of an organization is the set of countries in which it operates. An incident is "within range of" an organization if it occurred in one of those countries.

IF I occurred in T

AND T is located in C

AND O operates in C

THEN I is within range of O

Premise 37. Organizations need weapon resources to carry out terrorist attacks.

IF W was used in I

AND O possesses W

THEN O has resources for I

Premise 38. Ability is evidence for the responsibility of an organization for an incident.

IF O has training for I

AND O is within range of I

THEN O could be responsible for I

Premise 39. Possessing relevant resources is also evidence for responsibility.

IF O has resources for I

AND O is within range of I

THEN O could be responsible for I

Premise 40. Motive is further evidence of responsibility.

IF O has a motive for I

AND O is within range of I

THEN O could be responsible for I

Premise 41. It takes a lot of evidence to conclude that an organization is probably responsible for an incident.

IF O has a motive for I

AND O has training for I

AND O has resources for I

AND O is within range of I

THEN O is probably responsible for I

The fact that the highest level premise answers queries about responsibility for past terrorist incidents does not mean that only one type of question can be posed to DADM. Questions concerning any of the defined relations are possible. For example, the following queries (in English form) are typical of those that could be answered within the scope of the above premises and the experimental data base.

What organizations are aiding M-19?

Is the Red Brigade sending advisors to any other organizations?

Who would protect Jamid Gammal?

Does the ETA know underwater demolition?

What organizations use SAMs?

What organizations possess plastic explosives?

Who is supplying plastic explosives to the PIRA?

Does the PIRA operate in Spain?

In which past terrorist incidents were machine guns stolen?

What organizations would have a motive for incident 12?

Who could be responsible for incident 15?

List the incidents for which M-19 is responsible.

DADM can also respond to questions with assumptions. These stated assumptions temporarily over-ride the contents of the data base.

Assuming that the PLO sends advisors to the Red Brigade, could the Red Brigade be responsible for 120?

Could the ETA possess SAMs if Libya produced them?

If the PLO supplies weapons to ASALA, what weapons would ASALA possess?

What skills would the ETA know if Carlos is training for them?

Even with the relatively small set of premises constructed for this model, a great variety of questions can be asked. With a few hundred well selected and integrated premises along with a large data base, DADM can make a significant contribution to situation analysis as well as saving time in data base searches.

6. DEDUCTIONS

By using the base relations, DADM can answer questions concerning the information stored directly in the data base. However, by combining the base relations and the deduced relations with the premises, DADM can answer questions that require deductive reasoning. This section describes that deductive process starting with direct data base questions and progressing toward more complicated deductive-related questions. Each question submitted to DADM must contain references to known "base" or "deduced" relations. If the question contains only base relations, it requires no deductive processing and can be answered by accessing the data base directly. For example,

In what country is Damascus located?

Abdul Mohammad belongs to what organizations?

What organizations are based in Beirut?

Although the above questions are written in English for readability, the present implementation of DADM requires a special format for query input. For example,

In what country is Damascus located?

would be entered into DADM as follows:

(FIND DAMASCUS LOCATED IN COU-1)

This notation makes it easier for the DADM processor (written in LISP) to decode the query. ("COU-1" is a variable which stands for a country name.) A more natural English-like query language can be obtained by preprocessing the query and translating it into LISP format. In order to gain more freedom in query expression, a more complex English language processor is required. Recent experiments have integrated DADM with the End-User Friendly Interface to Data Management (EUFID) system. EUFID, designed and developed at SDC, is capable of accepting queries similar to the three examples shown above and translating them into a format acceptable to DADM.

In order to ask a high-level question, it is only necessary to reference high-level relations in the query. A question that contains deduced relations causes a deductive reasoning "chain" to be built which connects the relation mentioned in the question to existing base relations. This is accomplished by searching

through the list of premises for possible connections. For example,

Does the Red Brigade possess surface-to-air missiles?

This fact is not contained in the data base, but can be logically deduced from a number of other facts which are in the data base. The chain of reasoning is as follows. (Refer to Section 5 for the premises and Section 7 for the data base facts.)

From the data base, it is known that Russia produces SAMs. Thus, premise 21 permits the inference that Russia possesses SAMs.

Premise 21: IF Russia produces SAMs
THEN Russia possesses SAMs

This new fact, coupled with Russia's willingness to supply SAMs to Libya (see "would supply" in the data base), fits premise 26 which establishes a supply-line from Russia to Libya.

Premise 26: IF Russia would supply SAMs to Libya
AND Russia possesses SAMs
THEN Russia supplies SAMs to Libya

Once the supply line is established, premise 28 can be used to conclude that Libya possesses SAMs

Premise 28: IF Russia supplies SAMs to Libya
THEN Libya possesses SAMs

Now, Premises 26 and 28 can be used over again to establish the supply connection between Libya and the Red Brigade.

Premise 26: IF Libya would supply SAMs to the Red Brigade
AND Libya possesses SAMs
THEN Libya supplies SAMs to the Red Brigade

Premise 28: IF Libya supplies SAMs to the Red Brigade
THEN the Red Brigade possesses SAMs

The final premise has, as its consequent, a relation which precisely matches the query and thus forms the answer to the question.

Deductions concerning other defined relations can also be made in this manner. Some of them can be quite extensive and could

involve many steps including some premises which can be used in a recursive manner.

The following two example queries concern a particular incident, 120, about which some facts are known but not the responsible organization. The analyst is interested in finding out which terrorist organizations could be responsible for 120.

What organizations could be responsible for 120?

The query is entered into DADM in LISP format (see Figure 9):

(FIND ORG-1 COULD BE RESPONSIBLE FOR 120)

DADM responds with a data base search plan containing 7 separate retrieval queries. (See Figure 9.) Figure 10 shows the results of the deduction indicating that the PLO or the BJO could be responsible for 120. Also shown in Figure 10 are the evidence chain (left window), and two relevant premises used in the deduction (lower right windows). Information from the data base is shown after the word "FACT". Deduced relationships are shown after the word "CONCLUDE". The numbers after "**" are reference numbers for the inference plan and do not correspond to premise numbers.

DADM begins the deductive reasoning process in the section of the evidence chain marked **4 (see Figure 10, left window). Premise #19 (see Section 5) is used to conclude that the PLO possesses bombs from the data base fact that the PLO owns bombs.

Premise 19: IF the PLO owns bombs
THEN the PLO possesses bombs

The next step (**2) uses the above conclusion to establish that the PLO has sufficient resources to carry out incident 120. This requires premise #37 plus the additional data base fact that bombs were used in that particular incident.

```
MODE:  Query:
.(FIND ORG-1 COULD BE RESPONSIBLE FOR I20)

Query analyzed on  4-MAR-83 14:58:38

PLAN NUMBER: 1
7 search problems.
Answers?   Yes
              SEARCH/COMPUTE PLAN:
    SEARCH      *OCCURRED-IN I20 CITY-1
    SEARCH      *LOCATED-IN CITY-1 COUNTRY-1
    SEARCH      *OCCURRED-IN INCIDENT-1 CITY-2
    SEARCH      *LOCATED-IN CITY-2 COUNTRY-1
    SEARCH      *IS-RESPONSIBLE-FOR ORG-1 INCIDENT-1
    SEARCH      *WAS-USED-IN WEAPON-1 I20
    SEARCH      *OWNS ORG-1 WEAPON-1
```

Figure 9. What Organizations Could be Responsible for Incident I-20?

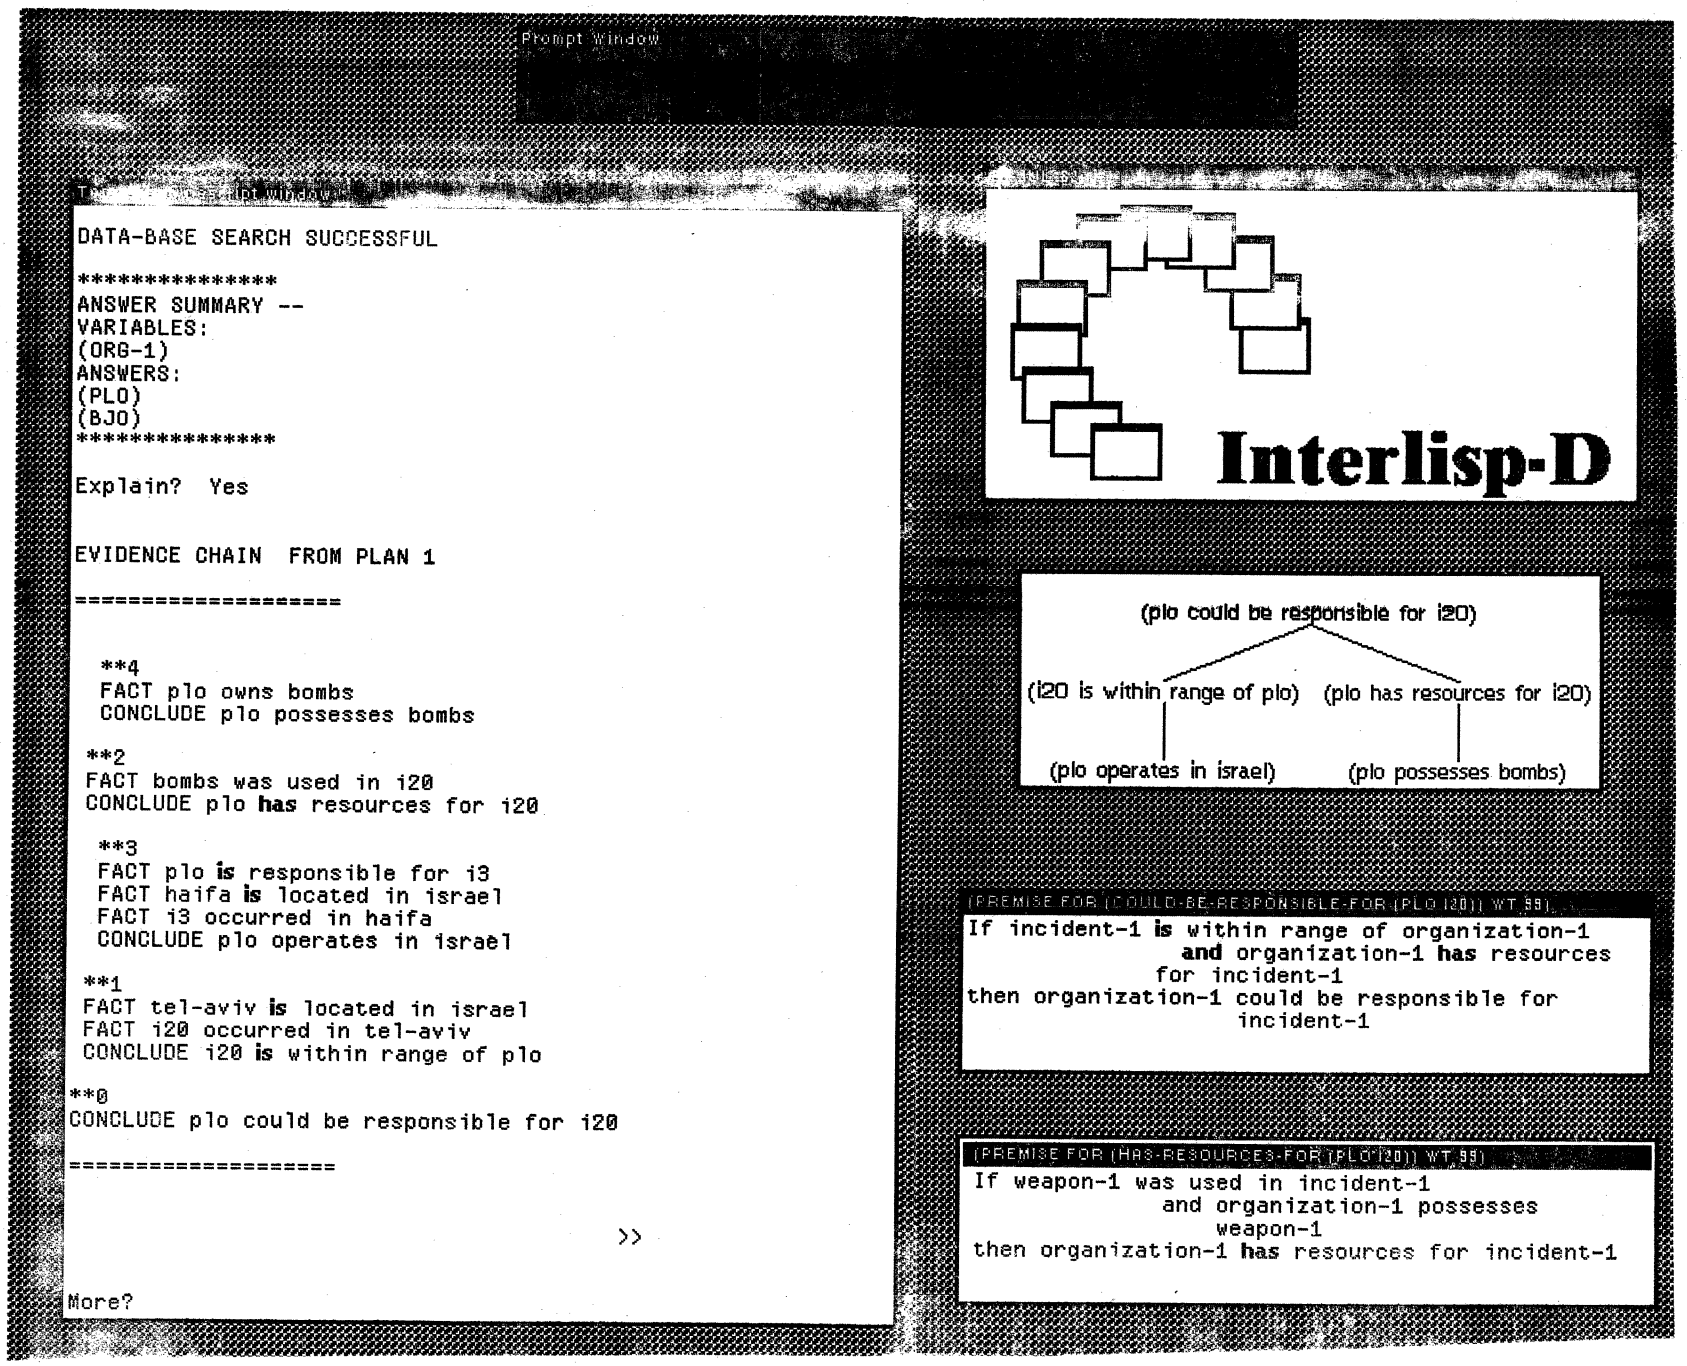

Figure 10. Both the PLO and the BJO Could be Responsible for Incident I-20

Premise 37: IF bombs were used in I20
AND the PLO possesses bombs
THEN the PLO has resources for I20.

The graph window on the right side of Figure 10 shows this line of reasoning and how it provides part of the information necessary to conclude that the PLO could be responsible for I20. However, an additional piece of information is needed. In step **3, DADM uses three data base facts to establish that the PLO operates in Israel (necessary because I20 occurred in Israel). However, a previous incident, I3, is used to draw this conclusion since it is already known that the PLO was responsible for I3 which also occurred in Israel.

Premise 31: IF the PLO is responsible for I3
AND I3 occurred in Haifa
AND HAIFA is located in Israel
THEN the PLO operates in Israel

Using this new information along with data concerning the location of incident I20, DADM is able to envoke Premise #36 to show that incident I20 is "within range of" the PLO (**1).

Premise 36: IF I20 occurred in Tel-Aviv
AND Tel-Aviv is located in Israel
AND the PLO operates in Israel
THEN I20 is within range of the PLO

In Figure 10, Section **1, there is no need to reproduce the line stating that the "PLO operates in Israel" since it was just concluded in the previous section (**3).

Finally, the original query can be answered using Premise #39 and the deduced information from previous premises (**0).

Premise 39: IF the PLO has resources for I20
AND the PLO is within range of I20
THEN the PLO could be responsible for I20

The chain of reasoning which leads to the conclusion that the BJO could also be responsible for incident I20 is identical to that

given in Figure 10. It is not necessary for DADM to create a totally new inference plan for the BJO. All possible answers to a query are carried simultaneously through a deductive chain and eliminated one-by-one as necessary when they do not satisfy the premise requirements.

The highest-level premise in the example set concerns the identification of an organization that is probably responsible for a particular terrorist incident. Such a query could be formulated as follows.

What organization is probably responsible for I20?

In order to complete this deduction, and arrive at a conclusion (if the data in the data base are sufficient for an answer), a number of different premises are required. Figure 8 in Section 2 shows the entire evidence tree for the above query. Again, information from the data base is shown after the word "FACT", and deduced relationships are shown after the word "CONCLUDE". At the bottom of Figure 8, a graphic display of the plan is shown.

Tracing through the plan of Figure 8, we can see the logical reasoning behind the conclusion that the PLO is probably responsible for incident I20. In **13, DADM makes the deduction that Libya possesses bombs by using Premise 19 along with the data base fact that Libya owns bombs. Then, in **12, a bomb-supply network is established from Libya to the PLO by using Premise 26. This permits DADM to conclude that the PLO possesses bombs by using Premise 28 (**11). Finally, this line of reasoning ends with a conclusion that the PLO has resources to carry out incident I20 (shown in **4). This conclusion is obtained from (1) the previous deduction which discovered that the PLO possesses bombs, (2) from the data base fact that bombs were used in I20, and (3) from the connecting Premise number 37.

At this point (**10), a different line of reasoning is started since the conclusion in **4 is to be used later. With three direct data base facts, DADM can conclude that the PLO operates in Israel. These facts refer to a previous incident (I3) in which the PLO is known to have been responsible. The fact that I3 occurred in Israel means that the PLO operates there. This connection is established with Premise 31. Using this new conclusion, DADM can establish that the incident in question (I20) is within the operational range of the PLO (**3), since it occurred in the same country (Premise 36).

Similarly, **9 and **2 establish a motive for 120 which stems from the fact that the Israelis were victims in a previous incident (13) for which the PLO was responsible. The conclusion in **9 uses Premise 32 and **2 uses Premise 35.

Finally, it must be established that the PLO has the capability to use bombing tactics. This comes from the fact that they are being trained by an outside agent (Carlos). Sub-plans **7, **8, **6, and **5 establish this fact using Premises 19, 17, 15, and 16 respectively. This allows the conclusion that the PLO has the necessary training for incident 120, proved in **1 using Premise 34.

With these four major conclusions (**4, **3, **2, **1), DADM can use Premise 41 and deduce that the PLO is probably responsible for 120 (**0).

7. EXAMPLE DATA BASE

The following sample data base contains entries for each of the defined base relations. In any real application, the data base would be much larger and more complete. Remember, the "open world" view is in effect (see Section 3). Thus, missing entries attain an "I don't know" status rather than "false".

P	belongs to O
Yasar Arafat	PLO
Jock Clancy	PIRA
Vallejo Desanchez	ASALA
Bernadette Devlin	PIRA
Bonet Erzin	ETA
Fatabyle Fawzial	ASALA
Otto Frenzel	RAF
Karl Freund	RAF
Jamid Gammal	PLO
Simone Losada	M-19
Abdul Mohammad	PLO
Mario Moretti	Red Brigade
Bill O'Toole	PIRA
Patrizio Peli	Red Brigade
Gerado Peraso	M-19
Shalim Sahid	PLO
Bobby Sands	PIRA
Tomaso Santini	Red Brigade
Richardo Santos	ETA
Antonio Scarini	Red Brigade

P	is based in T
Nikolai Aksenov	Mexico City
Carlos	Tripoli
Juan Delgado	Buenos Aires
Manual Esponza	Bogota
Jose Sanchez	Frankfort
Olga Schmidt	Moscow

O	is based in T
ASALA	Ankara
Black June	Bagdad
ETA	Bilbao
FMLN	San Salvador

MLAPU	Ankara
Montoneros	Buenos Aires
M-19	Bogota
RAF	Berlin
Red Brigade	Milan
PIRA	Belfast
PFLP	Beirut
PLO	Beirut

O	is based in	C
GAP		Guatemala

S	requires	W
Aircraft Attack		Surface-to-Air Missiles
Armed Attack		Machine Guns
Bombing		Bombs
Hostage Seizure		Grenades
Sabotage		Booby Traps
Underwater Demolition		Plastic Explosives

S	needs small arms only
Assassination	
Burglary	
Kidnapping	
Hijacking	

T	is located in	C
Buenos Aires		Argentina
Brussels		Belgium
Bogota		Columbia
Medullin		Columbia
San Salvador		El Salvador
London		England
Manchester		England
Cannes		France
Paris		France
Belfast		Ireland
Bethlehem		Israel
Haifa		Israel
Tel Aviv		Israel
Milan		Italy
Padua		Italy
Rome		Italy

Verona	Italy
Beruit	Lebanon
Tripoli	Libya
Lima	Peru
Moscow	Russia
Bilbao	Russia
Lemoniz	Spain
Madrid	Spain
Damascus	Syria
Ankara	Turkey
New York	United States
Frankfort	West Germany
Ramstein	West Germany
Mexico City	Mexico
Berlin	West Germany
Bagdad	Iraq

P trains for O

Carlos	PLO
Carlos	PIRA
Carlos	M-19
Carlos	Red Brigade
Abdul Mohammad	PIRA
Abdul Mohammad	Red Brigade
Shalim Sahid	PIRA

O owns W

PIRA	Plastic Explosives
PLO	Plastic Explosives
BJO	Bombs
PIRA	Bombs
PLO	Bombs
PIRA	Grenades
PLO	Grenades
M-19	Grenades
PLO	SAMs

P owns W

Carlos	Plastic Explosives
Carlos	Bombs

C	owns	W
Libya		Bombs

O	sends advisors to	O
PLO		Red Brigade
PLO		PIRA
PLO		Montoneros
PLO		M-19
PIRA		Red Brigade
PIRA		RAF
PIRA		ETA
PIRA		PLO

C	produces	W
Libya		bombs
Russia		SAMs

O	would supply	W	to	O
PLO		grenades		PIRA
PLO		grenades		ETA
PLO		grenades		Red Brigade
PLO		plastic explosives		PIRA
PLO		plastic explosives		M-19
PLO		SAMs		Red Brigade
PIRA		bombs		RAF
PIRA		bombs		ETA

C	would supply	W	to	O
Libya		bombs		PLO
Libya		SAMs		PLO
Libya		SAMs		Red Brigade

C	would supply	W	to	C

Russia	SAMs	Libya

P	has knowledge of	S

Carlos	Bombing
Carlos	Underwater demolition

Figure 11 lists the hypothetical incidents used in the counter-terrorism example. Each incident has eight entries:

(1) Identification code
(2) Location city
(3) Date
(4) Skill used
(5) Weapon used
(6) Victim involved
(7) Weapon stolen
(8) Responsible organization

I1	I6	I11
Manchester	Paris	Bethlehem
8/15/80	7/14/82	6/8/82
Bombing	Assassination	Sabotage
Bombs	Machine Guns	Booby Traps
British	Israeli	Israeli
none	none	none
PIRA	PLO	PLO
I2	**I7**	**I12**
Cannes	Haifa	Bilbao
6/30/82	7/4/82	4/15/82
Burglary	Hostage Seizure	Assassination
Plastic Explosives	Grenades	Bombs
French	Israeli	Spanish
Grenades	none	none
Red Brigade	BJO	ETA
I3	**I8**	**I13**
Haifa	New York	Damascus
7/16/82	10/10/81	4/18/81
Hijacking	Assassination	Bombing
Grenades	Grenades	Bombs
Israeli	Palestinian	Palestinian
none	none	none
PLO	BJO	unknown
I4	**I9**	**I14**
Medellin	Belfast	Tripoli
4/30/82	6/5/82	5/20/80
Burglary	Bombing	Bombing
unknown	Bombs	Bombs
none	British	Palestinian
Machine Guns	none	none
M-19	PIRA	BJO
I5	**I10**	**I15**
Frankfort	Rome	Buenos Aires
8/20/82	5/15/82	12/15/81
Bombing	Aircraft Attack	Burglary
Bombs	SAMS	Machine Guns
American	Italian	Argentinian
none	none	Plastic Explosives
RAF	Red Brigade	Montoneros

Figure 11. Data Base of Hypothetical Terrorist Incidents

I16	I21	I26
London	Lima	Madrid
6/10/81	9/1/80	8/1/82
Kidnapping	Burglary	Bombing
unknown	none	Bombs
British	none	Spanish
none	Grenades	none
unknown	Montoneros	ETA
I17	**I22**	**I27**
Ramstein	San Salvador	Milan
8/31/80	8/19/81	7/13/82
Bombing	Hijacking	Aircraft Attack
Bombs	Small Arms	SAMS
American	Salvadoran	Israeli
none	none	none
RAF	FMLN	PLO
I18	**I23**	**I28**
Rome	Bogota	New York
11/12/81	12/10/81	5/3/81
Burglary	Burglary	Bombing
Machine Guns	Bombs	Bombs
Italian	Columbian	Israeli
Machine Guns	Small Arms	none
unknown	M-19	unknown
I19	**I24**	**I29**
Verona	Ankara	Paris
3/13/80	6/8/81	8/12/82
Assassination	Assassination	Assassination
Booby Traps	Plastic Explosives	Machine Guns
American	American	Israeli
none	none	none
Red Brigade	MLAPU	PLO
I20	**I25**	**I30**
Tel-Aviv	Lemoniz	Belfast
5/24/80	5/12/82	2/6/83
Bombing	Underwater Demol.	Bombing
Bombs	Plastic Explosives	Bombs
Israeli	Spanish	British
none	none	none
unknown	ETA	unknown

Figure 11. Data Base of Hypothetical Terrorist Incidents (Cont'd)

8. SUMMARY AND CONCLUSIONS

This report has presented an experimental application of the deductive knowledge base approach to data base question answering, specifically, the application of DADM to the analysis of a hypothetical terrorism model. By using deductive inference techniques, DADM is able to automatically access relevant information in a data base and answer high-level queries concerning relationships among the data.

DADM possesses a number of sophisticated features that are not shown in the example terrorist model. Premises can be considerably more complex, using the logical connectives "OR" and "NOT" as well as embedded implications. They can also be recursive, using their own names within their definitions. Although most inference plans are generated in a "top-down" or "backward chaining" direction (from the final deduced relation in the query toward the data base), DADM can also use "forward chaining" and "middle chaining". This means that the answers to queries can be made by different logical search procedures which insure that all possible attempts to find answers will be taken.

In an actual application of DADM to an existing data base in a new environment, a number of procedural (and possibly software) tasks must be accomplished. These tasks include (1) domain identification, (2) relation development, (3) premise development, (4) premise integration, and possibly (5) query language translator implementation.

Domain identification requires an analysis of the target data base to determine the types of data elements present. The data elements are categorized and a "dictionary" of relevant terms is built. If the data base includes numeric entries, a specification of their uses is established (i.e., numeric ranges, codes, etc.).

Relation development is accomplished in two phases. First, all "base" relations must be identified. These come directly from the data base and reflect its internal structure. Second, a list of "deduced" relations must be constructed. Deduced relations also connect data base entries, but at a higher level. All words and phrases used to name deduced relations must be carefully chosen and approved so that their technical meaning is clearly understood.

Premise development follows relation development. The premises provide the exact definition of all of the relations and

precisely determine the conditions when each deduced relation can be considered "true" or "false". Both the deduced relations and the premises originate from the knowledge and background experience of experts in the application area. Consequently, expert interviews usually provide the foundation for their construction. If it is not possible to arrive at a consensus of opinion on the premises, DADM can be used with alternate premise sets, tailored to fit certain circumstances or individuals.

Premise integration is an important procedure that assures completeness and consistency within the premise set. Additional linking and supporting premises are usually necessary to cover areas which are often taken for granted by experts. In order to be effective, each premise must be carefully constructed so that it captures a relevant piece of expert knowledge and is consistent with all of the other premises. A set of controlled experiments finalizes the premise integration procedure and exercises the system before it becomes operational.

Since DADM constructs data base access queries internally, it must be able to communicate with the target database in its standard query language. DADM currently works best with a relational data base structure. If the target data base is not relational, or if it has a specialized query language, a DADM translator module must be implemented which can generate the proper queries. Thus, the target data management system does not have to be altered.

With proper application and system integration, DADM can provide the data analyst with a powerful tool for formulating conclusions that may otherwise take a great deal of tedious data searching and correlating. The inferencing system can easily be altered to reflect new situations or directives and, thus, provides a dynamic, usable product for aiding detailed analysis as well as high-level decision making.

9. REFERENCES

[1] Kellogg, C., Klahr, P., Travis, L., A Deductive Capability for Data Management, in Systems for Large Data Bases, Lockemann, P.C., Neuhold E.J., (eds), North-Holland Publishing Co., 1976.

[2] Kellogg, C., Creating Inference Plans for Deductive Access to Relational Data Bases, in Proceedings of Workshop on Logic and Data Bases C.E.R.T., Toulouse, France, 1977.

[3] Kellogg, C., Klahr, P., Travis, L., Deductive Methods for Large Data Bases, Proc. Fifth International Joint Conference on Artificial Intelligence, MIT, 1977.

[4] Kellogg, C., Klahr, P., Travis, L., Deductive Planning and Pathfinding for Relational Data Bases, in Logic and Data Bases, Gallaire, H., Minker, J., eds., Plenum Press, 1978.

[5] Kellogg, C., Travis, L., Reasoning with Data in a Deductively Augmented Data Base System, in Advances in Data Base Theory, Volume 1, Gallaire, H., Minker, J., eds., Plenum Press, 1981.

[6] Kellogg, C., Knowledge Management: A Practical Amalgam of Knowledge and Database Technology, in Proceedings of National Conference on Artificial Intelligence, Carnegie-Mellon University, 1982.

[7] Travis, L., and Kellogg, C., Deductive Power in Knowledge Management Systems (abstract), in Proceedings of Workshop on Logic and Data Bases, C.E.R.T., Toulouse, France, 1982 (Full paper to appear in Advances in Data Base Theory, Volume II, forthcoming).

[8] Leal, A., Levin, S., Johnston, S., Agmon, M. and Weltman, G., An Interactive Computer Aiding System for Group Decision Making, Preceptronics, Inc., report PQTR-1046-78-2, February, 1978.

SDC
System Development Corporation
A Burroughs Company
Corporate Headquarters, 2500 Colorado Avenue, Santa Monica, CA 90406