Managing Heterogeneity in Inter-Operating Medical Information Systems*

Maurice Dixon ¶ Jana Kohoutková ‡ Stephen Cook †
Keith Jeffery † Brian Read †

November 18, 1996

¶Computing Information Systems and Mathematics, London Guildhall University, London EC3, UK

‡ Institute of Computer Science, Masaryk University, Brno, Czech Republic
kohoutkova@ics.muni.cz

† Department for Computation and Information, Rutherford Appleton Laboratory, Chilton, Didcot, UK, OX11 0QX
s.cook@rl.ac.uk
kgj@inf.rl.ac.uk
b.j.read@rl.ac.uk

Abstract

This paper reports the progress made in managing the heterogeneities that arise in interoperating medical information systems. The approach taken is to define a domain ontology through the use of a reference model which allows for system evolution. Use is made of a five layer information expressivity model to segregate areas of concern.

The Hypermedata system is being developed to provide transfers of medical information between hospitals in the Czech and Slovak Republics. In addition to conventional data types, there is a need to transfer images, curves, and groupings of text and numeric data to represent test results.

The project is concerned with information exchange, not distributed update transactions. Semantic heterogeneity is identified as a key concern. Local systems may only partially overlap in their areas of concern and hence have only partially overlapping data models. Within those overlapping areas heterogeneity arises at least through:

1. Different procedures and equipment being used for the same tests
2. Version changes in procedures
3. Independent local system evolution
4. Addition of other local systems
5. Additional requirements

*The Hypermedata project is funded by EC Copernicus Grant 94-0943. The partners are AMIS (Czech Republic), ELAS (Slovak Republic), Masaryk University (Czech Republic), Rutherford Appleton Laboratory (UK), Logis (Czech Republic).

6. Access security controls
7. Different attribute names and semantics
8. Different constraints
9. Different attribute types for same attribute
10. Different data structures
11. Different relations among equivalent attributes (roles)

One of Hypermedata's techniques for managing this heterogeneity is to define a reference model for the domain. This will be an input to a knowledge-assisted process of schema resolution. Different modelling languages have been studied and their relative effectiveness for expressing the reference model are critically assessed.

1 Introduction

Hypermedata is a COPERNICUS project concerned with sharing medical records between hospitals within the Czech and Slovak Republics.

The healthcare systems will be autonomous and heterogeneous with the interoperation mechanism based upon data exchange. The records contain images, curves, and groupings of text with numeric data as well as conventional datatypes. The information that is held in medical information systems has great semantic richness and complexity [GGJ93]. This work is proceeding in the context of plans for patient records to be available in Europe at the point of need.

Although Hypermedata is conceptually a federated database, it does not use a global schema in the usual sense of that term. Consequently the system has a number of unusual features, particularly in its data exchange mechanism. However, the focus of this paper is on understanding and managing the heterogeneities of the participating systems, rather than the mechanics of the data exchange process.

1.1 Issues of Heterogeneity

There are many different types of heterogeneity that can occur within interoperating systems. We propose to group the various types into four broad categories [SL90, HDR96, KS91].

Systems Heterogeneity This concerns the differences in the platforms of component database systems eg hardware configuration, operating and communication systems.

Syntactic Heterogeneity This concerns the differences in the underlying data models of the component DBMSs e.g. differences in data structure primitives, differences in the way constraints are expressed, differences in query language.

Schematic Heterogeneity This concerns the situation where equivalent or related data concepts are present in different database systems and have conflicting structural representations. However, the canonical model must use just one basic structure so that the translation to the local structure can be unambiguous [SCGS91]. Saltor states that the OO, Functional and Relational models have this one basic representation (object, function, and relation respectively) but ER type models do not since they have both entity and relation as their modelling primitives and the same thing can be modelled in more than one way. We believe that the labelled graph modelling primitive which will be used in Hypermedata is closest to the functional model.

Semantic Heterogeneity This concerns the situation where mismatches and conflicts arise because of differences in meaning, interpretation and usage of data. These heterogeneities can be subtle and they represent a major research challenge. The problem has been succintly expressed by Heiler [Hei95]. She warned that the results of such mismatches can be catastrophic since the results can be wrong with no indication that they are wrong. This presents obvious dangers in a healthcare system. We seek to identify specific techniques within Hypermedata for identifying and resolving these conflicts. Such techniques can of course impact on the autonomy of the local health care systems; there is a balance to be struck here with the converse problem of losing information which cannot be exchanged. It is semantic heterogeneity that provides the major obstacle to system interoperation.

Issues of semantic and schematic heterogeneity are often entangled together. For example, a particular feature of concern is the definition of access control within the federation since medical records are subject to an expectation of high confidentiality. This control raises substantial issues of metadata heterogeneity at both the schematic and the semantic level.

1.2 Interoperating Systems

The taxonomy of the architectures of interoperating systems has been well characterised [SL90, HBP94]. The effective combination of relational and object databases has been demonstrated in the IRO-Db project [Kla94]. The combination and presentation of multimedia data from different systems has been demonstrated in the MIPS project [JHK+94].

1.3 The Hypermedata Project

1.3.1 Project Vision

The starting point for the Hypermedata project is the idea that hyperlinked multimedia documents are an ideal exchange mechanism and that the informality of such documents could be improved by graph theory. In short, representing hyperlinked multimedia documents as graphs should provide:

- a canonical representation (Cg)
- the basis for topological comparison (using an Analyser tool)
- the basis for the GUI representation for human interaction

1.3.2 Architecture

The architecture of Hypermedata is based on a two-step data transformation:

1. From proprietary form to graph form;
2. From graph form to canonical graph form;

for both schema and instances, which can be interpreted as either a topological structure (graph), a data structure (objects) or a document structure (hyper-linked multimedia). By combining this transformation with its converse, we can achieve exchange of data between any two participating systems in a way which is flexible, extensible and resilient to change.

Given the following components

- two or more participating Healthcare Information Systems (HISa, HISb) and their graph representations (HISag, HISbg), both as schema/intension (HISag-i, HISbg-i) and instances/extension (HISag-e, HISbg-e);
- a canonical representation in graph form (Cg) of both intension (Cg-i) and extension (Cg-e);

the project intends to provide:

1. A meta-level definition of (Cg) — an Exchange Format which controls evolution from Cg to Cg'.

2. Tools to convert instances under a given HIS from proprietary form (HISn-e) to graph form (HISng-e) to canonical graph form (Cg-e) and vice versa — Convertor Parameter Tables.

3. To provide tools to analyse HISag and HISbg together (at least at schema level, possibly at instance level) to propose Cg' — the Analyser.

The provision of the Analyser (which allows additional knowledge to be gathered from a data administrator) and Convertor Parameter Tables makes it unnecessary for the local schemas to be strict subsets of the canonical schema. In principle, the canonical schema can be generated for each exchange although in practice we expect that most exchanges will reuse previous versions of the conversion tables. Thus Hypermedata is loosely federated in the sense of the classification of [BG92].

1.3.3 Modelling Environment

We are presented with a modelling environment which is in a high state of flux. This arises in part from the dynamic state of the federation of systems composing Hypermedata. New systems with their own data models may be joining the federation in the future. The individual systems in the Hypermedata federation will themselves evolve. In addition new requirements may be added once the system is in production.

Ventrone and Heiler have suggested that the development of standardised reference models for domains would provide useful building blocks to support federation [VH94]. They propose that a repository holds the metadata which incorporates the data definitions and their semantics including constraints, rules, and procedures. Experience from MIPS has shown that knowledge based assistance is required to support interoperation [JHK+94]. In order to use automatic knowledge base tools, we shall follow the five layer information expressivity formalism developed for MIPS [JHK+94, DKJ95].

1.4 Purpose of this report

In this paper we report our approach to tackling the heterogeneities that arise in the interoperation of medical health care systems. Our actual experience of the analysis process with a federation of systems is different from the conventional system analysis paradigms. For example, during the short lifetime of the project one system has withdrawn from the federation already and has been replaced by another system which is significantly different. A key aim of this work is to establish the role of a reference model in reconciling the heterogeneities that arise through the evolution of the component systems and the requirements of the Hypermedata federation. We report investigations we have made into the use of different modelling techniques. In particular we discuss how these affect the design strategies.

2 Medical Records Systems

2.1 Why we need to interoperate

The initial concern of Hypermedata is with the operational use of patient records. The transfer of physical documents can be slow and possibly unreliable. We have identified a variety of scenarios where Hypermedata could be useful:

- Patients transfer between hospitals and departments
- Patients need urgent treatment at another hospital
- Patients who have had accidents will be treated at an Emergency Hospital
- Medical staff in different departments want to see the same data
- Remote consultations can be carried out with experts

Initially, the requests for information will be made informally and outside the scope of Hypermedata. Normally, they will be requests for the records of a particular patient. However, our investigations also found a need to produce aggregate records of economics and performance data for the purposes of planning of resources. This could be achieved by extending Hypermedata with a global query capability.

2.2 Principal Data Entities

Essentially, we are dealing with a person who has a medical history. The person becomes a patient of the medical centre. Medical staff ask for certain investigations to be carried out. The results of those investigations are reported. The results are interpreted by medical staff and these interpretations are recorded. The patient receives treatment. Finally the patient is discharged. The systems also allow tracking of the progress of the patient through the period of their treatment.

2.3 Heterogeneities in Hypermedata System

Systems Heterogeneity We do not expect systems heterogeneity to prove a major obstacle since intersystem communication will be via standard protocols.

Syntactic Heterogeneity We will not be requiring transactional updating across systems within Hypermedata so syntactic heterogeneities at the DML level are not expected to be a problem.

Data exchange will be via an exchange file which is neutral (in the sense that it permits various interpretations) and contains instances of the canonical model for that exchange. The exchange file contains the intensional information which tools need to interpret it. Each component database will be responsible for translating its represention to or from the common format when receiving or issuing an exchange file.

Syntactic heterogeneity at the datatype level (e.g. 6-character dates or 8-character ISO numeric dates) will be reconciled by the Analyser module.

Schematic Heterogeneity We are expecting schematic heterogeneity to require significant attention for the case of the same data having different rendering in different systems e.g. a chart represented as a bit map or a fitted curve. A more conventional example, such as medical staff responsibility for a treatment decision, should be straightforwardly resolvable.

The representation of schemas as graphs provides a means for comparing them topologically. However, where the Analyser detects ambiguity, it will have to seek additional knowledge from the data administrator. This can be considered for inclusion in a new version of the reference model.

A particular feature that concerns us is version control of investigation results since a hospital may change its equipment or its reporting practice. The need here is to build up a taxonomy of the schematic heterogeneities that may occur. We are using the reference model to help us explore the implications of schematic heterogeneity both in terms of datatypes and constraints.

Semantic Heterogeneity Semantic heterogeneity represents a major challenge for Hypermedata. The problem that we face is:

> "Can we ensure that a receiving hospital can interpret the data about a patient in the same way as the issuing hospital if we assume that the receiver depends only on the information issued by Hypermedata? Such a decision may involve life or death."

We are excluding from this the consideration of correct medical semantics being present in the issuing system record [GGJ93].

Based upon the work of Kim and Seo [KS91] we have been developing a taxonomy of semantic heterogeneity[HDR96]. Examples of the kinds of issues that arise are:

1. Procedure codes and death codes are used in the health care systems; when a code is interpreted at the receiving site there is a need to ensure it is correctly understood; in principle this kind of problem can be handled by using translation lists although we would not expect all the mappings to be binary 1:1.

2. Result data may be expressed in different units or use different calibrations or precisions, which may in turn depend on the kind of equipment in use in the hospital.

3. Different institutions may have different manual practices so that what they record or omit from recording is different.

4. The reduction of time dependent series data by averaging over time windows.

Further consideration needs to be given to the types of semantic heterogeneity that can arise for images, graphs and tabular data.

The approach we are taking is to define a semantically rich reference model. The data structures and constraints will be defined in terms of first order logic. We also need to represent temporal expressions and complex objects.

As systems evolve we expect to introduce more schematic/semantic heterogeneity since even if the same enhancements are introduced they will be introduced at different times.

2.4 European Developments

This work is taking place at a time when there is significant research effort being invested in the development of relocatable medical records. One approach is to have a portable personal medical record using a chip card based medical information system [KKE94]. For example, a prototype for diabetes patients is being developed. Such a card would raise significant heterogeneity problems at all four levels of our heterogeneity classification. A study aiming at the design of a European patient record concluded that

> "the complexity of a telematics patient record and the clinical work needed to produce it may be underestimated." [Lav94].

These and other work occuring in Europe force us to conclude that Hypermedata will face major challenges to adapt to heterogeneous sources of data during its operational life. Our approach to this conclusion is to seek more general solutions to overcoming the heterogeneities rather than taking a superset approach.

3 Analysis

3.1 Modelling Approach

In the analysis phase of this work the existing system investigation was carried out by one team and the model refinement was carried out by another team. This split arose because the project was required to be undertaken by a consortium and not all the consortium were fluent in Czech. Initial investigations were carried out on the AMIS*H and CARESYS systems focussing at the schema definition level; the first reports on these systems suggested that a canonical schema comprising a superset would be appropriate.

However, it became obvious during the construction of the canonical model that significant problems existed. For example, based on the schema definitions, there appeared to be incongrueties in admission date/time and discharge date/time (apparently AMIS*H admitted patients only as In-Patients and discharged them only as Out-Patients). Similarly, requests for examination did not appear to be related to the results of the examination.

3.2 Role of the Reference Model

Rather than pursue through schema, frames and code a resolution to these possible disparities, we decided to follow the domain ontology approach of Ventrone and Heiler[VH94].

We required a Reference Model which was not just a superset of the existing schemas, so that we could achieve unification at the semantic level and not just at the logical data structure level. This required us to understand and formulate the business rules of the sector of medical health care that we were studying in order to ensure consistent data semantics. We have been particularly keen to identify and segregate constraints.

This domain reference model enables us to evaluate the characteristics of a contributing system. It can be used to identify those components which appear to be implausible, inconsistent, incomplete or missing. Equally, it can highlight additional information topics which are candidates for inclusion in extensions of the reference model. We believe there is a need for formal methods to express these semantics clearly and to manipulate the semantics.

We also believe it is essential to study real and fully representative data, possibly in substantial quantities and with some tool assistance. This may reveal (or at least suggest) dependencies and constraints which are otherwise undocumented. Both semantics and constraints can be buried in code and can occur through manual procedures; they may lie hidden until the semantic heterogeneities they represent lead to serious errors of judgement. We suggest that some form of directed data mining can help to reveal these relationships.

In terms of the five layer model of information expressivity, we have studied various approaches to expressing requirements at the conceptual level, e.g. using Halpin's Object Role Modelling [NH89]. We intend to provide a semantic level by expressing data structures, constraints and other semantics in terms of a first order logic formalism; this, we hope, will allow us to use the knowledge based assists advocated in the MIPS project [JHK+94].

3.3 Reference Model Scope

Our domain model distinguishes between three key areas

- patient information
- medical investigations
- treatment tracking

It is concerned primarily with identifying medical responsibility for decisions and interpretations. Our model currently omits features such as financial considerations and departmental organisation of the hospitals but is easily extended to accommodate them. We give our basic Object-Role Model for Hypermedata in the Appendix.

3.4 Language Requirements

The intensions of each model (Graph, Data, Document) must be represented in some language. Ideally, we would like to use a single syntax with three semantic interpretations, corresponding to the models. On the basis of this single syntactical representation of symbols, one can build semantics which can be interpreted for a purpose or processing intent. The semantics are different for each model:

- Graph model: more-or-less no semantics, only topological relationships;
- Data model: rich semantics - O-O (Object-Oriented) or EE-R (Extended Entity-Relationship) level of expressivity (but ideally DOOD (Deductive Object-Oriented Database));
- Document model: medium-rich semantics, expressing document elements and their presentation ;

For exchange and transformation purposes, the richer semantics of the Data or Document models can be reduced to the skeletal, near-semantic-free Graph model.

3.4.1 Language Properties

Such a language would require the following properties:

- it must be formal and (ideally) provably correct;
- it must be succinct and represent accurately the information that is to be expressed;

- it must be executable (or interpretable for execution) on a variety of hardware / operating system platforms;

This leads towards the functional (terse, formally correct) / logical (expressive and provides deductive and inductive capabilities) class of languages.

3.5 Strategy for Adaptation to Change

Hypermedata needs a design that is stable under change due to

- clarification of current systems details
- upversioning of procedures
- adaptation of performance criteria
- evolution in use after interoperation
- additional requirements e.g. European record standards or work flow applications

We considered an approach based on requiring participating systems to conform to a federal schema. However, we concluded that as the federation grew, the management of schema evolution and propagation would become too great a burden. Hypermedata will use an innovative approach in which the canonical schema to be used for each exchange will be determined at runtime (as a knowledge-assisted and rule-constrained process). The reference model will play an important role by providing the first version of the canonical schema, which may then be adapted and/or reused as the federated systems evolve.

We believe this approach provides the best opportunity for stability because we will provide tightly controlled domain semantics with meta-rules to control reference model schema evolution. Local systems will take responsibility for defining their schema in terms of the Hypermedata graph definition syntax.

4 Data Modelling Formalisms

The language requirements identified above are challenging. We have therefore investigated whether existing data modelling formalisms can provide both sufficient expressivity to describe the kinds of heterogeneities we have discovered, and correspondences with the graph definition language which forms the data modelling primitive.

4.1 NIAM Modelling

Various extensions of E-R modelling have been proposed. One of the most persuasive is NIAM, particularly its derivative Object-Role Modelling (ORM) [NH89]. This technique improves both the formality and expressiveness of E-R. Non-binary relationships are easily expressed and a fairly rich set of constraints can be modelled. ORM produces a very flat conceptual model; there is no encapsulation mechanism and the distinction between entities and attributes is postponed to the logical model. We have used ORM to produce a reference data model for the Hypermedata domain; this is presented in the Appendix. It has already proved useful for validating HIS schemas because its similarity to E-R models makes comparison easy but its greater precision and expressiveness allows new insights into the models' semantics. The reference model draws on various sources of information, including the analysis of the participating systems, general knowledge of healthcare organisation and general principles of best practice data modelling.

4.2 Object-Z Modelling

Nevertheless, ORM does not remedy all the deficiencies of the E-R family of formalisms. In particular, we can note

- the lack of an encapsulation mechanism for large scale modelling;
- the absence of life history and other temporal constructs;
- the separation of state from process;
- the conflation of structural (is-a, is-part-of) relationships with contingent relationships (e.g. lives-at, works-for);

Some of the limitations of E-R modelling and its derivatives are addressed by the various flavours of object-oriented data modelling. We have conducted experiments using Object-Z [Ros92], which adds object-oriented and temporal features to the Z [Spi89] formalism. They demonstrate the ability of Object-Z to express the strong structure of objects and the internal constraints on their states and operations. Conversely, inter-object relationships and constraints are expressed more awkwardly, often requiring the invention of pseudo-objects to carry them. Z schemas are not directly executable but recent developments in deductive object-oriented databases (DOOD) have demonstrated executable systems which have close equivalences to most of the features (except temporal predicates) in Object-Z. We have examined the feasibility of using the ROCK & ROLL [FPWB92, BPF+94] DOOD as a modelling tool for Hypermedata at the logical level.

5 Conclusions

5.1 Systems are 'dirty'

Our analysis has confirmed the conventional wisdom that currently operational systems are very difficult to document for the following reasons:

1. The data structures are not always in third normal form.
2. The semantics are concealed in code and usage.
3. Constraints are not clearly identified.

In particular there are difficulties in establishing that data is semantically consistent especially for a range of novel datatypes.

5.2 Role of Reference Model in Handling Evolving Heterogeneity

The use of a reference model provides a framework for discussion, investigation and design. It avoids problems of code and usage semantics of an individual system being propagated into the federation of systems. Such a reference model must be stable against:

- re-analysis

- evolution of local systems
- additional systems joining the interoperation
- the users finding additional purposes for interoperation within the current system
- requirements extension for the whole of the federated system

We conclude that defining a domain ontology [VH94] provides the best way of ameliorating the problems of heterogeneity that arise as local and federation requirements change.

6 Acknowledgements

The authors acknowledge the work of project colleagues on analysis of schemas and basic data models (Frantisek Krejci (LOGIS), Milan Sarek (AMIS), Rudolf Lukacka (ELAS), Michal Javornik (MU-ICS)) and the work on the definition of the languages to support the representation of graphs (Miroslav Benesovsky, Karel Skoupy (MU-ICS)). The prototype which proved feasibility was constructed by Karel Skoupy.

References

[BG92] D. Bell and J. Grimson. *Distributed Database Systems*. Addison-Wesley, 1992.

[BPF+94] Maria L. Barja, Norman W. Paton, Alvaro A. A. Fernandes, M. Howard Williams, and Andrew Dinn. An effective deductive object-oriented database through language integration. In *Proceedings of the 20th VLDB Conference, Santiago, Chile*, 1994.

[DKJ95] Maurice Dixon, John Kalmus, and Keith Jeffery. A 5-layer information expressivity model applied to semantic heterogeneity in a decentralised organisation. In *Proceedings of ERCIM Database Research Group Workshop 8 on Database Issues in Co-operative Information Systems, Trondheim, 1995*, 1995.

[FPWB92] Alvaro A. A. Fernandes, Norman W. Paton, M. Howard Williams, and Andrew Bowles. Approaches to deductive object-oriented databases. *Information and Software Technology*, 34(12):787–803, December 1992.

[GGJ93] Carol Goble, Andrzej Glowinski, and Keith Jeffery. Semantic constraints in a medical information system. In M. Worboys and A. F. Grundy, editors, *Proceedings BNCOD-11: Advances in Databases*, pages 40–57. Springer-Verlag, 1993. Lecture Notes in Computer Science 696.

[HBP94] A. R. Hurson, M. W. Bright, and S. Pakzad. *Multidatabase Systems: An Advanced Solution for Global Information Sharing.* IEEE Computer Society Press, 1994.

[HDR96] S. Hamill, M. Dixon, and B. J. Read. Classifying schematic and semantic heterogeneities in interoperating database systems. RAL Report, Rutherford Appleton Laboratory, 1996. In preparation.

[Hei95] S. Heiler. Semantic interoperability. *ACM Computing Surveys*, 27(2), 1995.

[JHK+94] K. G. Jeffery, L. Hutchinson, J. R. Kalmus, M. Wilson, W. Behrendt, and C. Macnee. A model for heterogeneous distributed database systems. In *Proceedings of British National Conference on Databases 12*, page 221. Springer Verlag, 1994. Lecture Notes in Computer Science 826.

[KKE94] E. Kuhnel, G. Klepser, and R. Englelbrecht. Smart cards and their opportunities for controlling health information systems. *International Journal of Bio-Medical Computing*, 35:153–157, 1994.

[Kla94] W. Klas. Overview on interoperable database systems. In *Proceedings of ERCIM Database Research Group Workshop 6 on Deductive and Interoperable Databases, Barcelona, 1994*, page 1, 1994.

[KS91] Won Kim and Jungyun Seo. Classifying schematic and data heterogeneity in multidatabase systems. *IEEE Computer*, page 12, December 1991.

[Lav94] S. M. Lavelle. The intelligent decision mapping patient record (IDMR). *Computer Methods and Programs in Biomedicine*, 45:97–99, 1994.

[NH89] G. M. Nijssen and Terry J. Halpin. *Conceptual Schemas and Relational Database Design: A Fact-Oriented Approach.* Prentice Hall, 1989.

[Ros92] G. A. Rose. Object-Z. In Susan Stepney, Rosalind Barden, and David Cooper, editors, *Object Orientation in Z*, pages 59–77. Springer-Verlag, 1992.

[SCGS91] F. Saltor, M. Castellanos, and M Garcia-Solaco. Suitability of data models as canonical models for federated databases. *SIGMOD Record*, 20:44, 1991.

[SL90] Amit P. Sheth and James A. Larson. Federated database systems for managing distributed, heterogeneous, and autonomous databases. *ACM Computing Surveys*, 22(3):183–236, September 1990.

[Spi89] J. M. Spivey. *The Z Notation: A Reference Manual.* Prentice-Hall, 1989.

[VH94] V. Ventrone and S. Heiler. Some practical advice for dealing with semantic heterogeneity in federated database systems. *International Journal of Computer Aided Design*, 1994.

A Object-Role Model

This appendix presents extracts from a draft reference model for the HIS domain within the scope of Hypermedata. It is derived from various sources including the analaysis of participating systems. The model uses the Object-Role Modelling [NH89] formalism, a derivative of NIAM.

A.1 General Observations

The material presented here was generated using the Asymetrix Infomodeler CASE tool. The model focuses on entities and relationships; attributes are hidden. The suggested implementations of value types (longinteger, varchar etc) are not prescriptive but merely defaults of the modelling tool.

A.2 Object Model

A.2.1 Model definitions

Objects:

```
1. Clinic: Simple object-type identified by 'name' and  implemented as 'varchar 20'
   Clinic is a subtype of Treatment Centre
   subtype description: "A Treatment Centre where Patients receive out-patient
   treatment."

2. Examination Event: Simple object-type identified by 'id' and  implemented
   as 'smallint'

3. Examination Request: Simple object-type identified by '#' and  implemented
   as 'longinteger'
   description: "Examination Request is ..."
```

(items not shown)

Facts:

```
1. Examination Event produces Examination_resultset
   Each Examination Event produces at most one Examination_resultset and
   Every Examination_resultset has exactly one Examination Event that produces it

2. Examination Request issued by Staff / Staff issues Examination Request
   Every Examination Request issued by exactly one Staff

3. Examination Request specifies Examination Event
   Every Examination Event has exactly one Examination Request that specifies it

4. Examination Request uses Laboratory_capability
   Every Examination Request uses exactly one Laboratory_capability

5. Hospital contains Treatment Centre / Treatment Centre belongs to Hospital
   Each Treatment Centre belongs to at most one Hospital

6. Hospital employs Staff / Staff works for Hospital
   Each Hospital employs zero or more Staff and
   Each Staff works for zero or more Hospital
```

(items not shown)

A.2.2 NIAM Notation

Symbol Interpretation

ellipse : entity

rectangle : role in relationship

round-cornered rectangle : nested object

double-headed arrow : unique role(s) in relationship

bullet : mandatory role in relationship

single-headed arrow : subtype — supertype relationship

A.3 Diagram

See figure on next page.