

Logical Foundations for the Infrastructure of the Information Market

Michael Heather¹, David Livingstone² and Nick Rossiter²

¹ Ambrose Solicitors, St Bede's Chambers, Jarrow NE32 5JB, UK,
michael.heather@cantab.net

² School of Computing, Engineering and Information Sciences, Northumbria
University, NE1 8ST, UK

{david.livingstone,nick.rossiter}@unn.ac.uk

WWW home page: <http://computing.unn.ac.uk/staff/CGNR1/>

Abstract. The European knowledge-based economy has a complex product to market to the rest of the world. Techniques of the past, based on the closed-world assumption, have proved useful in many types of local information systems. However, theory and practice suggest that this approach may be inadequate for the infrastructure required. In databases the relational model through SQL has maintained wide dominance in business data processing. However, interoperability between different databases, even when based on the relational model, is proving a major problem. Predicate logic (consistent and complete to first-order) has many advantages for practical application. Interoperability however requires higher order logic as the arguments themselves are relations and functions. Higher order logic in the context of set theory behaves less satisfactorily according to Gödel's theorems as such logic cannot satisfy all three of soundness, completeness and effectiveness. This may be a fundamental reason why interoperability is proving to be so difficult. This paper looks at underlying problems and suggests that they may be avoided by the use of categorial higher order logic. Cartesian categories are complete, consistent and decidable. They can be employed as an engineering technique to construct a general architecture of interoperability.

1 Introduction

Information systems are basic building materials for the knowledge-based economy of the 2000 Lisbon strategy. The Lisbon agenda of March 2000 set out a strategy resulting in the eEurope 2005 Action Plan to build the knowledge-based economy of the single market [17]. The fundamental basis of knowledge is information which has to be handled appropriately by both the technology and by the law. What was appropriate for the physical media of the last millennium cannot be just carried over to the new digital media.

Modern information systems operate at every level: from data held in a single purpose fixed device, through common PCs at home or mobile computing with business systems of an SME, to databases, intra-acting locally and at national

level, inter-acting between nation states and even then open to wider global systems outside of Europe. This interoperability requires global coherence, which mirrors the interoperation of the EU itself. The reported slow progress with the Lisbon agenda is reflected in a similar tardy development of interoperable information systems. The latter is perhaps one of the causes of the former. The report of Wim Kok, quoted in [5], recommended that the agenda be re-focused on growth and employment to remedy the small progress over the first five years in member states. Like national employment the focus needs to be on the details of operation on local information systems.

Particular attention needs to be paid to the delivery of the Lisbon agenda. In order to achieve these objectives, the Union must do more to mobilise all the resources at national and Community levels so that their synergies ... can be put to more effective use ([6] at p.9)

There is a problem with logic. Successful local systems are first order but they need to participate in higher-order activity. The quest for synergy between levels runs into problems arising from semantic interoperability. There are two fundamental types of data: images and text. Images can be structured, as in graphs, or unstructured, as in photographs. Text can be structured, as in relational databases, or 'unstructured', as in natural language [16]. There is also process data, represented as transactions in many current systems. Music as performed is an interesting example of process data whose permanent form is expressed as graphics with notes on musical staves and as text with expressions in italics such as *andante*.

Syntactical interoperability is already achieved for digital and analogue data, either where the data is unstructured or where the data has a natural structure representable in a form that holds the informational content, for instance as ordered pixels or natural language text. These may be universally recognisable as in pictures or other interpretable materials where a common understanding readily exists. For instance an English text is interoperable throughout the world only in its native state or by translation into some other natural language.

Consequently there are not too many problems with interoperability of the natural data. We simply need operations at the syntactical level, which Google does well for text and for well-established image formats like jpeg, tiff and gif. However, these forms of information are essentially raw data. When value is added to the data through the application of analytical methods, we obtain structured data, for which inter-communication is problematical. In the past these problems were minimised because the data was used mainly in a local setting, that is intra-communication was needed where any standard was common to the locality. Interoperability itself is concerned with the inter-communication of data at different and therefore usually heterogeneous localities. Figure 1 summarises those definitions in the form of a table. The importance of exactitude in transmission of data around the single market and commercial inter-communication with the rest of the globe needs to be emphasised.

Type	Structure	Examples	Applications
Images	Structured	Graphics	Business graphics
	Natural ('Unstructured')	Photographs	Publishing
Text	Structured	Relational databases	Business data
	Natural ('Unstructured')	The web	Google
Intermediate Data	Meta-structure added to unstructured data	Semi-structured data (as in XML), Semantic Web, CAD/CAM Engineering Drawing with instructions	Fitting spare parts
Process Data	Dynamic	Banking Transactions	ATM accounting
		Biometric identity	Iris data

Fig. 1. Natural and Structured Data Types

2 Exactness

The concept of exactness is important in commerce. Customers always want what is ordered, whether it is good for their services or not, and not some approximation of the order. Of course in the real world there are always experimental errors but these can normally be controlled in a local environment and in practical terms can be minimised according to how much the customer wants to pay. Commerce has always had an international dimension but until very recently it has tended to be locally based. Communications between local bases have been point to point, with transportation by land, sea or air and conversations by mail, telephone or wireless broadcasting. The world consisted of a large number of mainly autonomous locally-based entities with simple inter-connection where the main effort is intra-activity at the local level, fit for the purpose with merchantable quality. A satisfactory theory can rely on linear models, linear logic, etc [8]. Local equates with classical and of the same type.

Exceptionally large systems, even if of the same type, may lie outside local operating conditions. For example the large database cannot be maintained by one person and a very large database cannot be copied because in the meantime it is changed. However, distinction is usually qualitative rather than quantitative.

Non-locality on the other hand may still be composed of what approximates to a set of localities of the same type. It then operates as a classical organisation. The US legal system is able to operate this way with a federal law coordinating different local state laws. The early days of the EU imposed the same laws on the member states which at that time were few in number. But the character of Europe is diversity and variety, in language, culture, customs and style with quite different ways of working in manufacturing, commodities and providing services. When the number of member states of Europe became enlarged, imposing the same laws became impossible and there was a move to harmonisation. It is a comparable position with information systems in different states. Co-ordinating

systems of the same type is only a first-order activity but for heterogeneous systems higher-order operations are needed.

Legal systems are archetypal general information systems. Intensionally the legal systems of the member states of Europe are identical and feed into the over-arching European law in the European Court of Justice. But each local legal system is extensionally different. We shall see here how the theory of inter-operating business information follows the same practice. The topos in Figure 8 could represent the legal systems of Europe just as easily as information systems.

Exact operations with systems of the same type mean it is easy to assume that the same conditions will apply for different type systems but in reality the results can be radically different and unpredictable even leading to dangerous results and therefore a subject for rigorous risk management. These may be derived from theory or experimental results, preferably both. The theory in this respect is shown by the Austrian mathematician Kurt Gödel. Gödel was able to show in his doctoral thesis of 1929 on the completeness theorem [9] that first-order predicate logic is complete.

The significance is that consistent logical propositions may be applied as a model of first-order systems. Thus a digital computer operating with a von Neumann architecture gives satisfactory results with first-order limits. Most of the work in applied mathematics in the last two centuries has been gauged on clever theories to keep within the first-order limits. However, there are applications which have defied such analytical methods like turbulence where it has been necessary to resort to more qualitative techniques as found in chaos theory. These do not provide exact results. This can also be explained by Gödel's theory of undecidability which is perhaps even better known [10]. Gödel's theorem shows that both intensional and extensional systems which rely on axiom and number are undecidable ³.

Traditional mathematical modelling, which relies on set theory cannot therefore be applied directly to higher order behaviour. An example of undecidability in the case of a computational machine relying on the Church-Turing thesis is the halting problem. A commercial example of this can be found in the implementation of Codd's relational model [1] as it is utilised in modified form in much of current data processing. As we shall see in the next section, in its pure form the relational model works well for atomic data because it is within Gödel's principle of first-order predicate logic and is therefore complete. That is the relational model and its corresponding calculus give exact results for atomic data. While the commercial version of the relational model SQL is a vast improvement on earlier data models, it has compromised some of the relational model features and is neither complete, nor decidable. The relational model is sometimes trumpeted as an example of the effectiveness of logic in computer science [11]. This

³ There is the question of even how to define completeness, consistency, decidability, soundness and effectiveness. The literature itself is not consistent and we will therefore leave aside what these words mean in a set theoretic context and rely below on corresponding categorial concepts as definitive.

is only for theoretical computer science. Implementations of the model show a divergence between theory and practice.

Of course real-world data does not consist of homogeneous independent items making up atoms. This has resulted in various techniques such as normalisation with a series of normal forms: first normal form, second, third, etc, which attempt to squeeze real-world phenomena into a collection of first-order relations, to behave optimally with regard to update and search operations. Hence data, lacking any naturally regular atomic form, may be squeezed by normalisation into such a structure in a consistent manner.

3 Practical Examples of Interoperability Problems

The relational model predominates in much of commerce today as the format for structured information. Yet there are very significant interoperability problems between one relational database and another. Some of these can be attributed to problems with the underlying SQL standard as described below.

3.1 Variants of SQL

Vendors of SQL DBMS support different variants of SQL, all in varying degrees differing from the versions of the SQL International Standard ⁴, either having additional features and/or omitting features.

Features	Achievements	Problem in interoperability
Full facilities	Not achieved by MySQL	Very difficult between MySQL and other DBMS
Hierarchies, manipulation	Peculiar to Oracle	Difficult between DB2 and Oracle in network/hierarchical structures
Recursive union, assembling networks	Peculiar to DB2	
Implementation of integer type	Oracle treats as numeric(38)	Difficult between Oracle and other DBMS in formatting and rounding numbers
Dates	Different logical formats	Difficult between all systems in reliable data format recognition

Fig. 2. Effects of Variants of SQL on Interoperability

Some features described in the standard are labelled implementation-dependent, meaning they are independent of the standard. Others are implementation-defined, meaning the manner in which the feature is achieved is at the discretion of the implementer. Therefore it is not always possible to guarantee a semantically valid transfer of data from one SQL DBMS to another, since the recipient DBMS may treat the received data in a different way to that which the sending DBMS would have treated it, had it carried out ostensibly the

⁴ Information Technology – Database Languages - SQL, ISO/IEC 9075:2003 (2003).

same operations. This situation arises because the standards are not based completely on scientific or mathematical principles. Standards are also influenced by the software vendors, who are looking for pragmatic and strategic ways in which their products can be promoted. Examples of problems at the data level are shown in Figure 2.

Therefore if two databases are to be interoperable, it is much simpler if they are both managed by the same DBMS package, because then the only problems in this context are those arising from the consistent application of one variant of SQL. For this and other reasons, such as reduced DBMS maintenance and licence fees, in practice multiple-vendor SQL DBMS are unheard of, except where they arise due to force of circumstance, for example the merger of two previously independent companies.

3.2 SQL versus the Relational Model

The relational model is based on two mathematical theories: first-order predicate calculus and relations ([1] at p.v). Interestingly the relations permitted are not completely general (at p.467-477). In particular a collection of n-ary relations of assorted degrees is strongly encouraged where n is an integer, giving the degree of a particular relation. Both a single universal relation and a collection of binary relations are strongly discouraged, as the former loses flexibility in logical navigation and the latter is cumbersome and unnatural. Further if the collection of n-ary relations is constrained to be in first normal form with all values atomic, then the predictable regular form greatly simplifies the query language.

No version of SQL implements the full relational model, either that specified by Codd or evolved from Codd's model by others, for example see [3]. Some differences between SQL and the relational model are summarised in Figure 3. The different structures of set and bag respectively for the relational model and SQL are of particular interest. Sets, bags and other container-types such as sequences of tuples are specifically defined and cannot be used interchangeably. However, there are means of carrying out conversions from one kind of container type to either of the other two, in an attempt to make it trivial to achieve mathematical exactitude in defining and manipulating the different kinds of tuple containers. Current work at Northumbria University in the Open Database Project shows the need for rigorous definition at the local level in prototyping languages like Raquel [14] to produce an open source implementation of the features specified in the Third Manifesto [3]. The aim is to keep as close as possible to the philosophy of the pure relational database model including object classes as data types orthogonal to relations, an open architecture satisfying this philosophy, a design for the architecture and an implementation of that design. Interoperability is facilitated by the use of pure relational languages such as Raquel, together with conversion techniques for mapping between different container types.

Feature in SQL	Feature in relational model	Consequence for SQL
Default structure is bag	Structure is set	Duplicate rows permitted, inconsistency in updates
Row identifier	No identifiers	Physical bias to extension
Rows may be sequenced	No sequencing of rows	Data is apparently ordered
Set operations such as union based on column position	Set operations such as union based on column name	Set operations are based on physical, not logical, ordering of columns
Duplicate column names allowed in output	No duplicates allowed	Confusing output

Fig. 3. Differences between SQL and the Relational Model

3.3 Closed World Assumption

The definition of a relation should be a logical predicate such that each tuple in a relation corresponds to a logical proposition that is true for that predicate. By the Closed World Assumption (that is the CWA) any tuple not in the relation represents a false proposition. This is an attempt to satisfy Gödel's decidability principle: that tuples in the relation are true and tuples outside the relation are false [2]. However it is not possible for a relational DBMS to guarantee that all the tuples in a database represent true propositions, only that they all consistently meet a set of integrity constraints that partially represent the real-world logical predicate. This is because the typing system is based on set inclusion principles rather than constructive ones. Therefore decidability may be a problem with all relational systems.

Nulls are another example of where database approaches based on CWA run into difficulties. ([1] at p.383-387) suggests that the relational model should permit nulls as markers, with two interpretations: missing-but-applicable and missing-and-inapplicable. Nulls are not data values. A four-valued logic is then employed to manipulate such data with the outcome of true, false or two types of maybe. SQL claims to have a three-valued logic since logical variables may take the values true, false or null. However, it is not clear that null can be safely equated with maybe and a number of problems arise as shown in Figure 4. From the Gödel perspective, nulls make a system undecidable. It is therefore not surprising that practical implementations such as SQL have many problems in handling nulls. Some more recent versions of the relational model do not permit nulls, for example [3]. It is likely that the handling of nulls will be facilitated by the use of metadata to describe the reason for the null. The questionable CWA assumption also raises problems in proving that the result of a query is logically valid. This gives rise to undecidability and a lack of completeness. Compounding the problem is that many users look for plausible results, often on small volumes of data.

Case	Result	Problem
Creation of nulls whether value is missing-but-applicable or missing-and-inapplicable	No distinction	Semantic simplification
Use of null to represent maybe in the Boolean type	Three values for Boolean logic	Contrary to normal view of Boolean logic as binary valued
Comparing a null value with a null value	maybe with join/restrict, true in set operations	Difference in outcome between set operations and other operations such as join
Split table into a set of sub-tables using restrict; union resulting sub-tables	No guarantee that this will be the original table	Restrict only returns rows where the comparison returns true; hence those returning null are ignored and lost
Aggregation operators applied to columns containing some nulls	Count includes them; others ignore them	Arbitrary application
Aggregation operators applied to columns containing all nulls	Count returns zero; others return null	Arbitrary application
Second order distributivity (e.g. fuzzy sets)	Logical equivalences are not true	Inconsistent treatment of nulls

Fig. 4. Problems with Handling of Nulls by SQL

It is perhaps worth raising the question as to whether object-oriented databases would overcome some of the disadvantages above. The answer is no, as objects will suffer from all the problems of sets and methods are implemented through an enriched type system, similar to sketches or perhaps 3-categories. The object-oriented approach needs to be founded in category theory to be complete and decidable. Codd said (at p.22): “One of the main reasons that object-oriented DBMS and prototype products are not going to replace the relational model and associated DBMS products is their systems appear to omit support for predicate logic. It will take brilliant logicians to invent a tool as powerful as predicate logic. Even then such an invention is not an overnight task. Once invented it may take more than a decade to be accepted by logicians. Thus features that capture more of the meaning of data, which is important, should be added to the relational model, instead of being proposed as replacements” [1]. Intuitionistic logic offers a more convincing way forward. Compared to first-order Boolean predicate logic the intuitionistic is more naturally applicable to open systems, is constructive and avoids the problem of impredication.

4 Interoperability and Categories

In category theory [15] alternative meanings of decidability, completeness, satisfiability, soundness and consistency, all used by Gödel, converge. They come

together in the composition diagram in Figure 5(a). The negation of these terms or where they fail wholly or in part are all subsumed in the diagram of punctured composition [7] in Figure 5(b). At one level the composition diagram of Figure 5(a) is a formal categorical representation of Gödel's result that first-order predicate logic is complete. This diagram therefore satisfies a local intra-operability of a single system and first-order interoperability between simple systems. The difference is that moving to higher orders such as axiomatic number systems is undecidable. However, that limitation does not apply to a process view of the arrow. Composition is still satisfied by the diagram but where the arrows can be a different type or from different levels. We have shown in [19] that free interchange between four levels can satisfy any realisable system. The conditions for interoperability come from adjointness [18] between the two composition triangles in Figure 6.




Fig. 5. Commuting Diagrams for a) Composition, b) Punctured Composition: $(g \circ f) \neq (g \circ f)'$

Critical details of these two triangles are the values η , ϵ , respectively the unit and counit of adjunction [13] in Figure 7. F is the functor that carries the data across from the left system to the right system. G is the underlying functor giving the rules for that transmission. f, g respectively represent dynamic data in the left and right systems. L is an object in category \mathbf{L} and R an object in category \mathbf{R} . Note that this only defines information in one direction. For two-way communication there has to be a self-adjointness of both left and right systems. As explained above in the example of legal systems this is a process of harmonisation. The systems do not need to be identical, that is η may be other than \perp and \top other than ϵ .

4.1 Architecture for Interoperability

The architecture for interoperability between more than two systems is then a composition diagram of the form shown in Figure 8. Fundamental category




Fig. 6. Cartesian closed adjointness. Each circle superimposes clockwise and anticlockwise closed arrows as identity functors indicated by the contravariant arrow heads (top and bottom representing initial and terminal objects respectively). The left (L) and right (R) categories are themselves opposites as shown by the arrowhead directions.




Fig. 7. Roles in Adjointness of a) η , b) ϵ

theory shows that for physical existence the real world operates as a cartesian closed category. All the categories drawn above are therefore cartesian closed. The theory also shows that any such operation involves only two categories (L, R) and a context category C , that is a left system communicates with a right system in the context of all other systems. All other systems are therefore a single context category. More precisely this is a topos T as in Figure 8. Remember that interoperability is really a global character where everything is connected to everything else. We are not dealing with discrete systems. The context category described above is only a view and is really the limit of the topos. So $C \rightarrow T$. And of course L and R are themselves subobjects of T . This is the essence of interoperability where category theory, or nothing less than category theory, can give the required insight to construct exact and decidable interoperating information systems. Category theory has long been associated

with graph theory. A graph-based approach to interoperability such as that by [20] using metamodels is broadly in the same direction as our work.

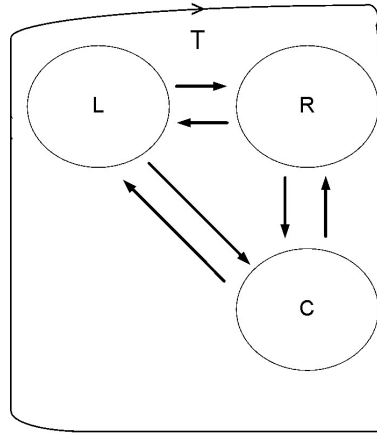


Fig. 8. Architecture for Interoperability: Topos T involving categories L , R and Context Category C

5 Conclusions

From the work of Gödel, first-order predicate systems are complete, consistent and decidable. Much of the attention in defining a relational data model has focused on keeping to a strict first-order system. The treatment of issues such as normalisation, nulls and recursion by workers developing a pure relational model [3, 1] is designed to avoid the need to handle higher-order logic in set theory. Indeed the relational model in its proper form is classified as one of the outstanding successes of logic in computer science [11]. The more casual treatment of such factors in SQL has led to systems, which are no longer consistent and decidable, giving many problems in interoperability. The kinds of problems with SQL standardisation are covered in [4] where underlying weaknesses in standards are thought to occur early on in the standardisation chain, such as through a weakness in the standards idea or the standards process.

Interoperability is essentially a higher-order problem, For higher order systems we need composability to achieve the same rigour as found in first-order predicate systems. Composability is a cornerstone of category theory and an architecture has been proposed, based on the topos, for achieving interoperability while meeting Gödel's requirements.

6 Acknowledgements

We are grateful to Hugh Darwen for his comments on the SQL standard and the relational model and to the contributors at the AREIN workshop, Madeira, March 2007 where a preliminary version of this work was first presented [12].

References

1. Codd, E F, *The Relational Model for Database Management* Addison-Wesley (1990).
2. Date, C J, *Gödel, Russell, Codd: A Recursive Golden Crowd* <http://www.dcs.warwick.ac.uk/~hugh/TTM/goedel.pdf> 6pp July 17th (2006).
3. Date, C J, & Darwen, Hugh, *Databases, Types and the Relational Model: The Third Manifesto* 3rd ed, Addison Wesley (2006).
4. Egyedi, T, Experts on Causes of Incompatibility between Standard-Compliant Products, in: *Enterprise Interoperability: New Challenges and Approaches* Doumeingts, G, Müller, J, Morel, G, & Vallespir, B, (edd) Springer 553-563 (2007).
5. Euractiv Network, *Relaunch of the Lisbon Strategy* <http://www.euractiv.com/en/innovation/growth-jobs-relaunch-lisbon-strategy/article-131891> (2005).
6. European Commission, Internal Guidelines in: *Integrated Guidelines for Growth and Jobs 2005-2008* 2005/0057, http://ec.europa.eu/growthandjobs/pdf/COM2005_141_en.pdf (2005).
7. Freyd, P, & Scedrov, A, *Categories, Allegories*, North-Holland (1990).
8. Girard, Jean-Yves, Une extension de l'interprétation de Gödel, à l'analyse, et son application l'élimination des coupures dans l'analyse et la théorie des types, *Studies in Logic and the Foundations of Mathematics* North-Holland 63-92 (1971).
9. Gödel, Kurt, *Über die Vollständigkeit des Logikkalküls* Doctoral Thesis, D1.736 33pp, University of Vienna (1929). Reprinted in Feferman, S, ed. *Gödel Collected Works*, volume 1 (1986).
10. Gödel, Kurt, Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme, I. *Monatshefte für Mathematik und Physik* **38** 173-198 (1931): translated in Jean van Heijenoort, *From Frege to Gödel: A Source Book on Mathematical Logic* Harvard 596-616 (1967).
11. Halpern, Joseph Y, Harper, Robert, Immerman, Neil, Kolaitis, Phokion G, Vardi, Moshe Y, Vianu, Victor, On the unusual effectiveness of logic in computer science *Bulletin of Symbolic Logic* **7**(2) 213-236 (2001).
12. Heather, Michael, Livingstone, David, & Rossiter, Nick, Higher Order Logic and Interoperability in Global Enterprise, *AREIN I-ESA* 26-30 March 12pp (2007).
13. Lawvere, F W, Adjointness in Foundations, *Dialectica* **23** 281-296 (1969).
14. Livingstone, David, *Open Database Project*, CEIS, Northumbria University, <http://computing.unn.ac.uk/openDBproject/> (2007).
15. Mac Lane, *Categories for the Working Mathematician*, 2nd ed, Springer (1998).
16. Probst, G, Raub, S, & Romhardt, K, *Managing Knowledge Building-Blocks for Success* Wiley (2000).
17. Rodrigues, Maria Joao, *The Debate Over Europe and the Lisbon Strategy for Growth and Jobs*, 2005.10.21, http://www.mariajoarodrigues.eu/files/The_Debate_over_Europe_and_LS.doc (2005).
18. Rossiter, Nick, & Heather, Michael, Conditions for Interoperability, *7th ICEIS* Florida, USA, 25-28 May 2005, 92-99 (2005).
19. Rossiter, Nick, Heather, Michael, & Nelson, David, A Natural Basis for Interoperability, in: *Enterprise Interoperability: New Challenges and Approaches* Doumeingts, G, Müller, J, Morel, G, & Vallespir, B, (edd) Springer 417-426 (2007).
20. Ziemann, J, Ohren, O, Jaëkel, F-W, Kahl, T, & Knothe, T, Achieving Enterprise Model Interoperability Applying a Common Enterprise Metamodel, in: *Enterprise Interoperability: New Challenges and Approaches* Doumeingts, G, Müller, J, Morel, G, & Vallespir, B, (edd), Springer p.199-208 (2007).