

Using Information-Flow Theory to Enable Semantic Interoperability

Marco Schorlemmer¹ Yannis Kalfoglou²

¹*Centre for Intelligent Systems and their Applications
School of Informatics
The University of Edinburgh*

¹*Escola Universitària de Tecnologies d'Informació i Comunicació
Universitat Internacional de Catalunya*

²*Department of Electronics and Computer Science
University of Southampton*

Abstract. We observe an ever growing need for integration in today's research agendas across a variety of organisations. The proliferation of ontologies and other similar knowledge-rich and labour-intensive structures as well as their exposure to a distributed environment like the Web, and eventually its successor, the Semantic Web, justifies the need. Although a plethora of solutions have been proposed and used, there are many issues which remain unclear. The most striking one is the antithesis in the availability of solutions for semantic integration as opposed to the abundance of techniques and methods for syntactic integration. In this paper we make the first step towards semantic integration by proposing a mathematically sound application of channel theory to enable semantic interoperability of separate ontologies representing similar domains.

1 Introduction

In the context of a distributed environment like the Web, Uschold and Grüninger point out that, “two agents are semantically integrated if they can successfully communicate with each other” [12], and successful communication means that they understand each other and there is guaranteed accuracy. This is a requirement for complete semantic integration in which the intended models of both agents are the same, that is, all the inferences that hold for one agent, should also hold when translated into the other agent's ontology, the authors continue. This is proposed as the golden standard of semantic integration, but we are skeptical about how or whether it can be achieved in computationally tractable manners. As it has been shown in a recent case study by Corrêa da Silva et al. [3], ontologies, which are naturally believed to be the right vehicle for this task, “fall short in providing adequate solutions in certain knowledge sharing scenarios.” These are mostly concerned with problem solving knowledge, where inferential knowledge needs to be made explicit when shared. As the authors state, “there ought to be, beyond the usual ontological correspondence between the communicating systems, a correspondence between the inference engines, in terms of their operators and deduction rules.”

Although the debate on adequacy of ontologies is interesting, it is out of the scope of this paper. We tackle the problem of semantic heterogeneity from a theoretical standpoint with attainable practical applications in a variety of knowledge sharing structures, including ontologies. One way to achieve the ambitious goal of semantic integration is to proceed in a step-wise fashion. In our view, to be semantically integrated presupposes to be semantically inter-operable. That's the focus of this paper. Semantic interoperability as a prerequisite for semantic integration. Our aim is to capture semantic interoperability between separate systems and to represent and model it in formal structures in order to reason over those in subsequent integration steps. Having achieved that, we will then be able to establish semantic-preserving exchange of information between the communicating systems, which is the first, and arguably, the most crucial step in achieving the kind of inferential knowledge sharing Uschold and Grüninger [12] and Corrêa da Silva et al. [3] are calling for.

2 The Role of Information Flow

A satisfactory way then, to approach semantic interoperability is via a formal notion of *information flow*. For that reason we will use *channel theory*, a modern theory of semantic information and information flow put forward by Barwise and Seligman in [1]. This theory underlies also Kent's Information Flow Framework [8], which attempts to accomplish this goal of interoperability as well. In Appendix A we list the main definitions we are using in this paper in order to explore how channel theory can help us to put the task of semantic interoperability on a firm theoretical ground. We have been putting the prefix 'IF' in front of channel-theoretic constructions to distinguish them from their standard meaning. For a more in-depth understanding of channel theory we point the interested reader to [1].


The key channel-theoretic construct we are going to exploit is that of a *distributed IF logic*. This is the IF logic that represents the information flow occurring in a distributed IF system. In particular we will be interested in a restriction of this IF logic to the language of those communities we are attempting to integrate. The basic idea is the following.

Suppose two communities **A** and **B** need to inter-operate, but are using different ontologies in different contexts. We use an *IF classification* as a very simple mathematical structure that effectively captures the local syntax and semantics of a community for the purpose of semantic interoperability. The syntactic expressions that a community uses will constitute the *types* of the IF classification. Depending on the kind of semantic interoperation we want to achieve, *types* can be concept or class symbols, relation names, complex queries or logical expressions, or even sets of expressions. The meaning that these expressions take within the context of the community will be represented by the way *tokens* are classified to *types*. Hence, the semantics is characterised by what we choose to be the tokens of the IF classification for a particular community; therefore, these will vary depending on the particularities of a semantic interoperability scenario. *Tokens* may, for example, be particular instances of classes or abstract first-order structures. The crucial point is that the semantics of the interoperability scenario crucially depends on our choice of types, tokens and their IF classification for each community. The example in Section 3 will make this point clearer.

To have communities **A** and **B** semantically inter-operating will mean to *know the semantic relationship in which they stand to each other*. In terms of the channel-theoretic context, this means to know an IF theory that describes how the different types from **A** and **B** are logically related to each other, i.e., an IF theory on the union of types $typ(\mathbf{A}) \cup typ(\mathbf{B})$

that respects the local IF classification systems of each community—the meaning each community attaches to its expressions—but also interrelates types whenever there is a similar semantic pattern, i.e., a similar way communities classify related tokens. In such an IF theory a sequent like $\alpha \vdash \beta$, with $\alpha \in \text{typ}(\mathbf{A})$ and $\beta \in \text{typ}(\mathbf{B})$, would represent an implication of types among communities that is in accordance to how the tokens of different communities are connected between each other.

This IF theory is the IF theory of the distributed IF logic of an IF channel



that represent the information flow between \mathbf{A} and \mathbf{B} . This channel can either be stated directly, or indirectly by some sort of partial alignment of \mathbf{A} and \mathbf{B} . The logic we are after is the one we get from *moving* a logic on the core \mathbf{C} of the channel to the sum of components $\mathbf{A} + \mathbf{B}$.

- Its set of types is the disjoint union of all the types of the component IF classifications: That is the language we speak in a semantic interoperability scenario, because we want to know when type α of one component corresponds to a type β of another component.
- Its IF theory will be over this set of types, hence a constraint $\alpha \vdash \beta$ will represent that every α is a β , together with a constraint $\beta \vdash \alpha$ we obtain type equivalence.
- The IF theory will be induced at the core of the channel; this is crucial. The distributed IF logic is the inverse image of the IF logic at the core; therefore the type and tokens system at the core and the IF classification of tokens to types will determine the IF logic at this core. We usually take the natural IF logic as the IF logic of the core. This seems natural, and is also what happens in the various interoperability scenarios we have been investigating.
- It is interesting though, that since the distributed IF logic is an inverse image, soundness is not guaranteed, which means that the semantic interoperability is not reliable in general. Even if $\alpha \dashv\vdash \beta$ in the IF logic, there might be tokens (instances, situations, models, possible worlds) of the respective components for which this is not the case. Reliable information flow is only achieved for tokens that are connected through the core. The way in which infomorphisms from components to the core are defined in an interoperability scenario is crucial. If these infomorphisms are token-surjective, then the distributed IF logic will preserve the soundness of the IF logic of the core. Proving the token-surjectiveness is hence a necessary task in order to guarantee reliable semantic interoperability.

In the following section we develop the above key ideas using an hypothetical, but realistic example.

3 Interoperability via IF Channels

We elaborate on an imaginative scenario to demonstrate the strengths of channel theory in capturing semantically rich information for alignment purposes. We are dealing with a situation where an agent or a group of agents (human or artificial) are faced with the task of

aligning organisational structures and responsibilities of ministries across different governments. This is a realistic scenario set out in the domain of e-governments, and despite its imaginative nature, its complexity and importance differentiates it from mapping ontologies of real world academic departments described in [6], where similar technology was used.

Our agents have to align UK and US governments, by focusing on governmental organisations, like ministries. The focal point of this alignment, is not only the structural and taxonomic differences of these ministries but the way in which responsibilities are allocated in different departments and offices within these ministries.

For the sake of brevity and space reasons, we only describe here four ministries: The UK Foreign and Commonwealth Office, the UK Home Office, the US Department of State, the US Department of Justice (hereafter, FCO, HO, DoS and DoJ, respectively). We gathered information related to these ministries from their web sites¹ where we focused on their organisational structures, assuming that the meaning of these structures is in accordance to the separation of responsibilities. These structures were trivial to extract, either from the hierarchical lists of departments, agencies, bureau, directorates, divisions, offices (which we shall commonly refer to as *units*) within these ministries, or organisational charts and organograms publicly available on the Web. The extraction of responsibilities and their units though, requires an intensive manual knowledge acquisition exercise (typically, a mission statement under a *what we do* hyperlink).

The ministries' taxonomies range from 38 units comprising the US DoJ to 109 units for the UK HO. In this example we focus on the alignment of 3 common responsibilities between these ministries:

- *passport services*, responsibility of HO and DoS;
- *promote productive relations*, responsibility of FCO and DoS;
- *immigration control*, responsibility of HO and DoJ.

Four steps towards semantic interoperability

In order to achieve the semantic interoperability we desire, we will go through the following four steps:

1. We define the various contexts of each community by means of a distributed IF system of IF classifications.
2. We define an IF channel—its core and infomorphisms—connecting the IF classifications of the various communities.
3. We define an IF logic on the core IF classification of the IF channel that represents the information flow between communities.
4. We distribute the IF logic to the sum of community IF classifications to obtain the IF theory that describes the desired semantic interoperability.

These steps illustrate a theoretical framework and need not to correspond to actual engineering steps; but we claim that a sensible implementation of semantic interoperability can be

¹ Accessible from www.homeoffice.gov.uk, www.fco.gov.uk, www.state.gov and www.usdoj.gov.

achieved following this framework, as it constitutes the theoretical foundation of a semantic interoperability scenario. In fact, Kalfoglou and Schorlemmer use similar techniques to assist in ontology mapping [6]. In the remainder of this section we apply the above four steps to our hypothetical interoperability scenario.

3.1 Community IF Classifications

UK and US governments use different ontologies to represent their respective ministries; we shall be dealing, therefore, with two separate sets of types:

$$\begin{aligned} \text{typ}(\mathbf{UK}) &= \{\mathbf{FCO}, \mathbf{HO}\} \\ \text{typ}(\mathbf{US}) &= \{\mathbf{DoS}, \mathbf{DoJ}\} \end{aligned}$$

We model the interoperability scenario using a separate IF classification for each government, **UK** and **US**, whose types are ministries.

To have UK and US ministries semantically inter-operable will mean to know the semantic relationship in which they stand to each other, which we take to be their set of responsibilities. It is sensible to assume that there will be no obvious one-to-one correspondence between ministries of two governments because responsibilities of a ministry in one government may be spread across many ministries of the other, and vice versa. But we can attempt to derive an IF theory that describes how the different ministry types are logically related to each other—an IF theory on the union of ministry types $\text{typ}(\mathbf{UK}) \cup \text{typ}(\mathbf{US})$ in which a constraint like $\mathbf{FCO} \vdash \mathbf{DoS}$ would represent the fact that a responsibility of the UK Foreign and Commonwealth Office is also a responsibility of the US Department of State.

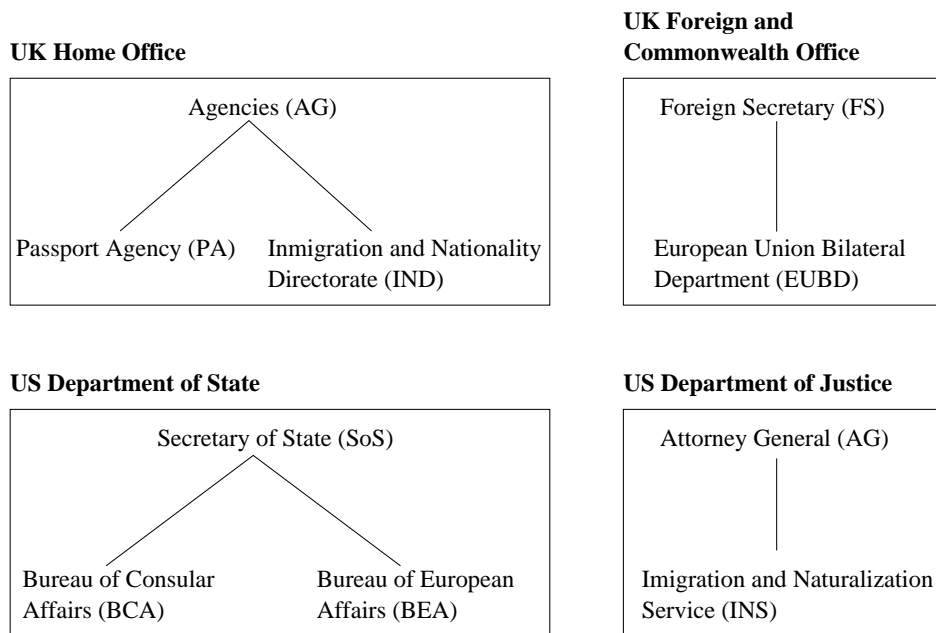


Figure 1: Hierarchical structures of government ministries

We shall construct the IF channel that will allow us to derive the desired IF theory using the hierarchical structure of units shown in Figure 1. Within the context of one government, different ministries represent already the top-level separation of responsibilities.

From the hierarchical structures we extract an IF theory on unit types for each government. Following are the two IF theories of UK and US units, respectively:

$\vdash \text{AG,FS}$ $\text{AG,FS} \vdash$ $\text{PA} \vdash \text{AG}$ $\text{IND} \vdash \text{AG}$ $\text{PA,IND} \vdash$ $\text{EUBD} \vdash \text{FS}$	$\vdash \text{SoS,AGe}$ $\text{SoS,AGe} \vdash$ $\text{BCA} \vdash \text{SoS}$ $\text{BEA} \vdash \text{SoS}$ $\text{BCA,BEA} \vdash$ $\text{INS} \vdash \text{AGe}$
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

By extracting responsibilities from the units' web sites we are able to define an IF classification for each government whose tokens are responsibilities and whose types are ministry units, and then classify responsibilities to their respective units. These IF classifications will have to be in accordance to the hierarchy as represented in the IF theories. That is, if a responsibility is classified to a unit, it shall also be classified to all its supra-units. This can be done automatically. In the case of UK units, the IF classification \mathbf{A}_{UK} will be the following:

	AG	PA	IND	FS	EUBD
r_1	1	1	0	0	0
r_2	1	0	1	0	0
r_3	1	0	0	0	0
r_4	0	0	0	1	1
r_5	0	0	0	1	0

Here tokens r_1 to r_5 represent responsibilities extracted from the units' web sites. So, token r_2 stands for the responsibility *immigration control* of the Immigration and Nationality Directorate, and hence also for the Agencies, while token r_3 stands for a responsibility of the Agencies only. For the US units we proceed in the same way:

	SoS	BCA	BEA	AGe	INS
s_1	1	1	0	0	0
s_2	1	0	1	0	0
s_3	1	0	0	0	0
s_4	0	0	0	1	1
s_5	0	0	0	1	0

However, the phrasing of responsibilities in the US web sites might differ from that in the UK web sites, which will result in a separate set of tokens s_1, \dots, s_5 for IF classification \mathbf{A}_{US} .

To represent how ministry types (like FCO,HO, etc.) from the IF classification UK relates to the IF classification \mathbf{A}_{UK} of ministerial units, we will use the *flip* \mathbf{A}_{UK}^\perp of the IF classification table and its *disjunctive power* $\vee \mathbf{A}_{UK}^\perp$. The flip classifies ministerial units to responsibilities, and for the UK case it is:

	r_1	r_2	r_3	r_4	r_5
AG	1	1	1	0	0
PA	1	0	0	0	0
IND	0	1	0	0	0
FS	0	0	0	1	1
EUBD	0	0	0	1	0


The disjunctive power of this flip classifies ministerial units to sets of responsibilities, whenever some of its responsibilities is among those in the set. Here is a fragment of this IF classification:

	$\{r_1, r_2, r_3, r_4, r_5\}$	\dots	$\{r_1, r_2, r_3\}$	\dots	$\{r_4, r_5\}$
AG	1		1		0
PA	1		1		0
IND	1		1		0
FS	1		0		1
EUBD	1		0		1

The way ministries relate to these sets of responsibilities can then be represented with an infomorphism $h_{UK} : \mathbf{UK} \rightleftarrows \vee \mathbf{A}_{UK}^\perp$:

$$\begin{aligned} h_{UK}(\mathbf{HO}) &= \{r_1, r_2, r_3\} \\ h_{UK}(\mathbf{FCO}) &= \{r_4, r_5\} \end{aligned}$$

Each context for a government, with its ministries, their respective units, and hierarchy captured by an IF theory, is then represented as a distributed IF system of IF classifications. For the UK government this distributed system is the following:



In the next step we use the flips \mathbf{A}_{UK}^\perp and \mathbf{A}_{US}^\perp to align responsibilities in order to achieve the desired semantic interoperability.

3.2 The IF Channel

We construct an IF channel from a partial alignment of some of the responsibilities extracted from the ministerial units' web sites. This is the crucial aspect of the semantic interoperability, since it is the point where relations in meaning are established. We assume a partial alignment, that is, one where not all responsibilities r_1 to r_5 are related to responsibilities s_1 to s_5 . In particular we shall assume the alignment of UK responsibilities r_1 , r_2 and r_4 with US responsibilities s_1 , s_4 and s_2 :

- passport services: $r_1 \longleftrightarrow s_1$
- immigration control: $r_2 \longleftrightarrow s_4$
- promote productive relations: $r_4 \longleftrightarrow s_2$

The focus of this paper is not how this partial alignment is established; various heuristic mechanisms have been proposed in the literature (see e.g., [10]). We assume that we have already applied one of those heuristics. Our purpose here is to provide a framework that shows how a partial alignment of a few responsibilities fits into the larger picture of an alignment scenario as the one described here, and represented as a distributed IF system, and how a

global IF theory of semantic interoperability on the level of government ministries is derived from this partial alignment.

The above partial alignment is a binary relation between $typ(\mathbf{A}_{UK}^\perp)$ and $typ(\mathbf{A}_{US}^\perp)$. In order to represent this alignment as a distributed IF system in channel theory, we decompose the binary relation into a couple of total functions $\hat{g}_{UK}, \hat{g}_{US}$ from a common domain $typ(\mathbf{A}) = \{\alpha, \beta, \gamma\}$. (For example $\hat{g}_{UK}(\beta) = r_2$ and $\hat{g}_{US}(\beta) = s_4$.) This will constitute the type-level of a couple of infomorphisms. We complete the alignment to a system of IF classifications

$$\mathbf{A}_{UK}^\perp \xleftarrow{g_{UK}} \mathbf{A} \xrightarrow{g_{US}} \mathbf{A}_{US}^\perp$$


by generating the IF classification on $typ(\mathbf{A})$ with all possible tokens, which we generate formally, and their classification:

	α	β	γ
n_0	0	0	0
n_1	0	0	1
n_2	0	1	0
n_3	0	1	1
n_4	1	0	0
n_5	1	0	1
n_6	1	1	0
n_7	1	1	1

To satisfy the fundamental property of infomorphisms, the token-level of g_{UK}, g_{US} must be as follows:

$$\begin{array}{ll} \check{g}_{UK}(\mathbf{AG}) = n_6 & \check{g}_{US}(\mathbf{SoS}) = n_5 \\ \check{g}_{UK}(\mathbf{PA}) = n_4 & \check{g}_{US}(\mathbf{BCA}) = n_4 \\ \check{g}_{UK}(\mathbf{IND}) = n_2 & \check{g}_{US}(\mathbf{BEA}) = n_1 \\ \check{g}_{UK}(\mathbf{FS}) = n_1 & \check{g}_{US}(\mathbf{AGe}) = n_2 \\ \check{g}_{UK}(\mathbf{EUBD}) = n_1 & \check{g}_{US}(\mathbf{INS}) = n_2 \end{array}$$

This alignment allows us to generate the desired channel between **UK** and **US** that captures the information flow according to the aligned responsibilities. This is done by constructing a classification **C** and a couple of infomorphisms $f_{UK} : \forall \mathbf{A}_{UK}^\perp \rightleftarrows \mathbf{C}$ and $f_{US} : \forall \mathbf{A}_{US}^\perp \rightleftarrows \mathbf{C}$ that correspond to a category-theoretic colimit [9] of the following distributed IF system, which includes the alignment and the contexts of each government:



This is a cover of the distributed IF system.

3.3 The IF Logic on the Core

This is how colimit \mathbf{C} is constructed: Its set of types $typ(\mathbf{C})$ is the disjoint union of types of $\forall \mathbf{A}_{UK}^\perp$ and $\forall \mathbf{A}_{US}^\perp$; its tokens are connections—pairs of tokens—that connect a token a of $\forall \mathbf{A}_{UK}^\perp$ with a token b of $\forall \mathbf{A}_{US}^\perp$ only when a and b are sent by the alignment infomorphisms g_{UK} and g_{US} to tokens of the alignment IF classification \mathbf{A} that are classified as of the same type. For example, the core \mathbf{C} will have a token $\langle \mathbf{AG}, \mathbf{SoS} \rangle$ connecting $\forall \mathbf{A}_{UK}^\perp$ -token \mathbf{AG} with $\forall \mathbf{A}_{US}^\perp$ -token \mathbf{SoS} , because $\check{g}_{UK}(\mathbf{AG}) = n_6$ and $\check{g}_{US}(\mathbf{SoS}) = n_5$, and both n_5 and n_6 are of type α in \mathbf{A} .

The following is a fragment of the IF classification on the core (not all types are listed, but all tokens are):

	$\{r_1, r_2, r_3\}$	$\{r_4, r_5\}$	$\{s_1, s_2, s_3\}$	$\{s_4, s_5\}$
$\langle \mathbf{FS}, \mathbf{BEA} \rangle$	0	1	1	0
$\langle \mathbf{EUBD}, \mathbf{BEA} \rangle$	0	1	1	0
$\langle \mathbf{FS}, \mathbf{SoS} \rangle$	0	1	1	0
$\langle \mathbf{EUBD}, \mathbf{SoS} \rangle$	0	1	1	0
$\langle \mathbf{IND}, \mathbf{AGe} \rangle$	1	0	0	1
$\langle \mathbf{IND}, \mathbf{INS} \rangle$	1	0	0	1
$\langle \mathbf{PA}, \mathbf{BCA} \rangle$	1	0	1	0
$\langle \mathbf{PA}, \mathbf{SoS} \rangle$	1	0	1	0
$\langle \mathbf{AG}, \mathbf{BCA} \rangle$	1	0	1	0
$\langle \mathbf{AG}, \mathbf{SoS} \rangle$	1	0	1	0

It shows the IF classification of all connections to those types of the core that are in the image of $f_{UK} \circ h_{UK}$ and $f_{US} \circ h_{US}$, which are the infomorphisms we will use in the next step to distribute the IF logic on the core to the IF classifications \mathbf{UK} and \mathbf{US} .

As the IF logic on the core we will take the natural IF logic of the IF classification \mathbf{C} , whose constraints are:


$$\begin{aligned}
 \{r_4, r_5\} &\vdash \{s_1, s_2, s_3\} \\
 \{s_4, s_5\} &\vdash \{r_1, r_2, r_3\} \\
 \{r_1, r_2, r_3\}, \{r_4, r_5\} &\vdash \\
 &\vdash \{r_1, r_2, r_3\}, \{r_4, r_5\} \\
 \{s_1, s_2, s_3\}, \{s_4, s_5\} &\vdash \\
 &\vdash \{s_1, s_2, s_3\}, \{s_4, s_5\}
 \end{aligned}$$

The natural IF logic is the one that captures in its constraints a complete knowledge of the IF classification. Since we have constructed the IF classification from those in the distributed system—which captured the contexts of governments together with the alignment of certain responsibilities—the natural IF logic will have as its IF theory all those sequents that conform to the government’s contexts as well as to the alignment, which is what we desire for semantic interoperability.

3.4 The Distributed IF Logic

The natural IF logic has an IF theory whose types are sets of responsibilities taken from UK or US web sites, but we want to know how this theory translates to government ministries, by

virtue of what responsibilities each ministry has. For that reason we take the IF theory of the distributed IF logic of the IF channel:



which is the inverse image along $(f_{UK} \circ h_{UK}) + (f_{US} \circ h_{US})$ of the natural IF logic $Log(\mathbf{C})$ generated from the core IF classification. Its theory has the following constraints:

$$\begin{array}{ll}
 \text{FCO} \vdash \text{DoS} & \text{DoJ} \vdash \text{HO} \\
 \text{FCO,HO} \vdash & \vdash \text{FCO,HO} \\
 \text{DoS,DoJ} \vdash & \vdash \text{DoS,DoJ}
 \end{array}$$

These constraints capture the semantic interoperability between all ministries, UK and US.

4 Related Work

Previously, Kalfoglou and Schorlemmer have shown that information flow can be used to assist in ontology mapping [6]. Their work demonstrates a practical application of information-flow theory in the area of ontology mapping where the ontologies used were representing academic departments from different universities, which were eventually mapped onto each other.

A complementary agenda is that currently pursued by Kent with his Information Flow Framework [8], which aims at providing a meta-level framework for a standard that will specify an upper ontology, enabling computers to interoperate. Targeted to upper ontologies, his effort is focused on providing the necessary framework for specifying ontologies that define concepts that are meta, generic, abstract and philosophical, and therefore are general enough to address (at a high level) a broad range of domain areas.

Similar work on using the notion of classification of tokens is demonstrated in the FCA-MERGE system [11], where the underpinning theory is that of formal concept analysis [5]. Last, but not least, there is a plethora of less formal approaches for semantic integration, notably the work on using communities of practice and learning algorithms [4], and on constraint-satisfaction-based systems [2]. We refer to [7] for an extensive survey.

5 Conclusions

In this paper we presented a practical application of channel theory to capture and model semantic interoperability in terms of information flow between different systems that needed to be integrated. The strong mathematical foundations of channel theory and their potential transformation to logic programs enabled us to work out a real world integration scenario with semantic-preserving exchange of information. These could provide a better understanding of the foundations for building and deploying semantically integrated systems in distributed environments.

Acknowledgements

This work is supported under the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University.

The first author is also supported by a ‘Ramón y Cajal’ Fellowship from the Spanish Ministry of Science and Technology.

References

- [1] J. Barwise and J. Seligman. *Information Flow*. Cambridge University Press, 1997.
- [2] S. Bressan and C. Goh. Semantic integration of disparate information sources over the internet using constraint propagation. In *CP'97 Workshop on Constraints and the Internet*, 1996.
- [3] F. Corrêa da Sliva, W. Vasconcelos, D. Robertson, V. Brilhante, A. de Melo, M. Finger, and J. Agustí. On the insufficiency of ontologies: problems in knowledge sharing and alternative solutions. *Knowledge-Based Systems*, 15(3):147–167, 2002.
- [4] N. Friesen. Semantic interoperability, communities of practice and the CanCore learning object metadata profile. In *11th International World Wide Web Conference*, 2002.
- [5] B. Ganter and R. Wille. *Formal Concept Analysis*. Springer, 1999.
- [6] Y. Kalfoglou and M. Schorlemmer. Information-flow-based ontology mapping. In *On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE*, LNCS 2519, pages 1132–1151. Springer, 2002.
- [7] Y. Kalfoglou and M. Schorlemmer. Ontology mapping: the state of the art. *Knowledge Engineering Review*, 2003. In press.
- [8] R. Kent. The information flow foundation for conceptual knowledge organization. In *Sixth International Conference of the International Society for Knowledge Organization*, 2000.
- [9] S. Mac Lane. *Categories for the Working Mathematician*. Springer, second edition, 1998.
- [10] P. Mitra and G. Wiederhold. Resolving terminological heterogeneity in ontologies. In *Ontologies and Semantics Interoperability*, CEUR-WS 64, 2002.
- [11] G. Stumme and A. Maedche. FCA-Merge: Bottom-up merging of ontologies. In *17th International Joint Conference on Artificial Intelligence*, 2001.
- [12] M. Uschold and M. Grüninger. Creating semantically integrated communities on the World Wide Web. In *WWW 2002 Semantic Web Workshop*, 2002.

A Channel Theory

IF classification: $\mathbf{A} = \langle tok(\mathbf{A}), typ(\mathbf{A}), \models_{\mathbf{A}} \rangle$ consists of a set $tok(\mathbf{A})$ of *tokens*, a set $typ(\mathbf{A})$ of *types*, and a binary relation $\models_{\mathbf{A}}$ between $tok(\mathbf{A})$ and $typ(\mathbf{A})$.

Infomorphism: $f : \mathbf{A} \rightleftarrows \mathbf{B}$ from classifications \mathbf{A} to \mathbf{B} is a contra-variant pair of functions $f = \langle f^{\sim}, f^{\wedge} \rangle$ satisfying the Fundamental Property $f^{\sim}(b) \models_{\mathbf{A}} \alpha$ iff $b \models f^{\wedge}(\alpha)$, for each token $b \in tok(\mathbf{B})$ and each type $\alpha \in typ(\mathbf{A})$; f is *token-surjective* if f^{\sim} is surjective.

Flip: \mathbf{A}^{\perp} is the classification whose tokens are $typ(\mathbf{A})$ and types are $tok(\mathbf{A})$, such that $\alpha \models_{\mathbf{A}^{\perp}} a$ iff $a \models_{\mathbf{A}} \alpha$.

IF channel: \mathcal{C} is an indexed family $\{f_i : \mathbf{A}_i \rightleftarrows \mathbf{C}\}_{i \in I}$ of infomorphisms with a common codomain \mathbf{C} , the *core* of \mathcal{C} . The tokens of \mathbf{C} are called *connections*.

Sum: $\mathbf{A} + \mathbf{B}$ of classifications has as set of tokens the Cartesian product of $tok(\mathbf{A})$ and $tok(\mathbf{B})$ and as set of types the disjoint union of $typ(\mathbf{A})$ and $typ(\mathbf{B})$, such that for $\alpha \in typ(\mathbf{A})$ and $\beta \in typ(\mathbf{B})$, $\langle a, b \rangle \models_{\mathbf{A} + \mathbf{B}} \alpha$ iff $a \models_{\mathbf{A}} \alpha$, and $\langle a, b \rangle \models_{\mathbf{A} + \mathbf{B}} \beta$ iff $b \models_{\mathbf{B}} \beta$. Given two infomorphisms $f_{1,2} : \mathbf{A}_{1,2} \rightleftarrows \mathbf{C}$, the sum $f_1 + f_2 : \mathbf{A}_1 + \mathbf{A}_2 \rightleftarrows \mathbf{C}$ is defined by $(f_1 + f_2)^\wedge(\alpha) = f_i(\alpha)$ if $\alpha \in \mathbf{A}_i$ and $(f_1 + f_2)^\vee(c) = \langle f_1^\vee(c), f_2^\vee(c) \rangle$, for $c \in tok(\mathbf{C})$.

Distributed IF system: \mathcal{A} consists of an indexed family $cla(\mathcal{A}) = \{\mathbf{A}_i\}_{i \in I}$ of classifications together with a set $inf(\mathcal{A})$ of infomorphisms all having both domain and codomain in $cla(\mathcal{A})$.

Cover: An IF channel $\mathcal{C} = \{h_i : \mathbf{A}_i \rightleftarrows \mathbf{C}\}_{i \in I}$ covers a distributed IF system \mathcal{A} if $cla(\mathcal{A}) = \{\mathbf{A}_i\}_{i \in I}$ and for each $i, j \in I$ and each infomorphism $f : \mathbf{A}_i \rightleftarrows \mathbf{A}_j$ in $inf(\mathcal{A})$, $h_i = h_j \circ f$.

Disjunctive power: $\vee \mathbf{A}$ of an IF classification \mathbf{A} is the classification whose tokens are the same as \mathbf{A} , whose types are subsets of $typ(\mathbf{A})$, and given $a \in tok(\mathbf{A})$ and $\Phi \subseteq typ(\mathbf{A})$, $a \models_{\vee \mathbf{A}} \Phi$ iff $a \models_{\mathbf{A}} \sigma$ for some $\sigma \in \Phi$. There exists a natural embedding $\eta_{\mathbf{A}} : \mathbf{A} \rightleftarrows \vee \mathbf{A}$ defined by $\eta_{\mathbf{A}}^\wedge(\alpha) = \{\alpha\}$ and $\eta_{\mathbf{A}}^\vee(a) = a$, for each $\alpha \in typ(\mathbf{A})$ and $a \in tok(\vee \mathbf{A})$.

IF theory: $T = \langle typ(T), \vdash \rangle$ consists of a set $typ(T)$ of types, and a binary relation \vdash between subsets of $typ(T)$. Pairs $\langle \Gamma, \Delta \rangle$ of subsets of $typ(T)$ are called *sequents*. If $\Gamma \vdash \Delta$, for $\Gamma, \Delta \subseteq typ(T)$, then the sequent $\Gamma \vdash \Delta$ is a *constraint*. T is regular if for all $\alpha \in typ(T)$ and all sets $\Gamma, \Gamma', \Delta, \Delta', \Sigma_0, \Sigma_1$ of types:

1. *Identity:* $\alpha \vdash \alpha$
2. *Weakening:* If $\Gamma \vdash \Delta$, then $\Gamma, \Gamma' \vdash \Delta, \Delta'$
3. *Global Cut:* If $\Gamma, \Sigma_0 \vdash \Gamma, \Sigma_1$ for each partition $\langle \Sigma_0, \Sigma_1 \rangle$ ($\Sigma_0 \cup \Sigma_1 = typ(T)$ and $\Sigma_0 \cap \Sigma_1 = \emptyset$), then $\Gamma \vdash \Delta$.

IF classification generated by an IF theory: Given a regular IF theory T , the classification $Cl_a(T)$ generated by T is the classification whose tokens are partitions $\langle \Gamma, \Delta \rangle$ of $typ(T)$ that are not constraints of T , and types are the types of T , such that $\langle \Gamma, \Delta \rangle \models_{Cl_a(T)} \alpha$ iff $\alpha \in \Gamma$.

IF logic: $\mathcal{L} = \langle tok(\mathcal{L}), typ(\mathcal{L}), \models_{\mathcal{L}}, \vdash_{\mathcal{L}}, N_{\mathcal{L}} \rangle$ consists of a classification $cla(\mathcal{L}) = \langle tok(\mathcal{L}), typ(\mathcal{L}), \models_{\mathcal{L}} \rangle$, a regular theory $th(\mathcal{L}) = \langle typ(\mathcal{L}), \vdash_{\mathcal{L}} \rangle$ and a subset of $N_{\mathcal{L}} \subseteq tok(\mathcal{L})$ of *normal tokens*, which satisfy all the constraints of $th(\mathcal{L})$; a token $a \in tok(\mathcal{L})$ satisfies a constraint $\Gamma \vdash \Delta$ of $th(\mathcal{L})$ if, when a is of all types in Γ , a is of some type in Δ . An IF logic \mathcal{L} is *sound* if $N_{\mathcal{L}} = tok(\mathcal{L})$.

Natural IF logic: It is the IF logic $Log(\mathbf{A})$ generated from an IF classification \mathbf{A} , and has as classification \mathbf{A} , as regular theory the theory whose constraints are the sequents satisfied by all tokens, and whose tokens are all normal.

Inverse image: Given an infomorphism $f : \mathbf{A} \rightleftarrows \mathbf{B}$ and an IF logic \mathcal{L} on B , the inverse image $f^{-1}[\mathcal{L}]$ of \mathcal{L} under f is the local logic on \mathbf{A} , whose theory is such that $\Gamma \vdash \Delta$ is a constraint of $th(f^{-1}[\mathcal{L}])$ iff $f^\wedge[\Gamma] \vdash f^\wedge[\Delta]$ is a constraint of $th(\mathcal{L})$, and whose normal tokens are $N_{f^{-1}[\mathcal{L}]} = \{a \in tok(\mathbf{A}) \mid a = f^\vee(b) \text{ for some } b \in N_{\mathcal{L}}\}$. If f is token-surjective and \mathcal{L} is sound, then $f^{-1}[\mathcal{L}]$ is sound.

Distributed IF logic: Given a binary IF channel $\mathcal{C} = \{f_{1,2} : \mathbf{A}_{1,2} \rightleftarrows \mathbf{C}\}$ and an IF logic \mathcal{L} on its core \mathbf{C} , the distributed IF logic $DLog_{\mathcal{C}}(\mathcal{L})$ is the inverse image of \mathcal{L} under the sum $f_1 + f_2$.