Institution:	**University of Oxford**
Unit of Assessment:	**11 Computer Science and Informatics**
Title of case study:	**Reasoning Systems (HermiT) (8)**

1. Summary of the impact (indicative maximum 100 words)

State-of-the-art reasoning systems developed in the UoA have underpinned the standardisation of ontology languages, and play a critical role in numerous applications. For example, HermiT, software developed in the UoA, is being used by Électricité de France (EDF) to provide bespoke energy saving advice to 265,000 customers in France, and a roll out of the use of the system to all of their 17 million customers is planned.

2. Underpinning research (indicative maximum 500 words)

Ontologies are formalised vocabularies of terms. The meaning of each term in the ontology is defined by its relationships with other terms in the ontology. Ontologies thus capture the semantics of the vocabulary in a way that can be exploited in software systems; they can be thought of as logical theories in a carefully designed decidable subset of First Order Logic. Ontologies have long since been important as a means to develop and maintain large structured vocabularies, and for capturing knowledge about an application domain. They are important for capturing domain knowledge in areas as diverse as healthcare and energy.

Ontologies are an important tool in efforts to improve access to and sharing of data on the web, where they are used to structure data and to provide structured vocabularies for annotations. In order to ensure that ontologies can be shared across all web-based applications, the World Wide Web Consortium (W3C) has standardised an ontology language called OWL that builds on existing web standards such as XML and RDF. The OWL standard, and the resulting improvement in the range and quality of ontology tools, has given an added impetus to applications of ontologies in industry, where they are increasingly used to address information integration and access problems.

Effective reasoning systems are essential both to support ontology development, and for exploiting ontologies in applications. HermiT was originally developed in the EPSRC funded HermiT project (EP/F065841/1), which ran from August 2008 to February 2012, and continues to be developed as an Oxford-led open source project (see http://www.hermit-reasoner.com/). Ian Horrocks was the principal investigator on the project and Boris Motik the co-investigator. They led a team of post-doctoral researchers and PhD students including Birte Glimm, Giorgos Stoilos and Rob Shearer.

HermiT is based on a novel hypertableau algorithm [1]. This algorithm is inherently less nondeterministic than the tableau algorithms used in earlier reasoners, and HermiT also employs a wide range of innovative optimisation techniques that further improve performance on typical ontologies. These include individual reuse [4] and core blocking [2], as well as a novel classification algorithm that can significantly reduce the number of reasoning tasks that need to be performed in order to compute all atomic subclass entailments [3].

HermiT is the only reasoner that fully supports the OWL 2 ontology language standard, which requires the basic calculus to be extended with reasoning over datatypes [5] and classification of properties (binary predicates) as well as classes [3]. It has also been extended to support several

language features that go beyond what can be expressed in OWL 2, in particular Description Graphs and Rules, which together support modelling of complex structures such as those found in human anatomy [6].

3. References to the research (indicative maximum of six references)

The three asterisked outputs best indicate the quality of the underpinning research.

***[1] Boris Motik, Rob Shearer, and Ian Horrocks. Hypertableau Reasoning for Description Logics. J. of Artificial Intelligence Research, 36:165-228, 2009.**
https://www.jair.org/media/2811/live-2811-4689-jair.pdf
Presents the Hypertableau algorithm used in HermiT and proves various computational properties including soundness, completeness, termination and complexity.
[2] Birte Glimm, Ian Horrocks, and Boris Motik. Optimized Description Logic Reasoning via Core Blocking. In Jürgen Giesl and Reiner Hähnle, editors, Proc. of the Int. Joint Conf. on Automated Reasoning (IJCAR 2010), volume 6173 of Lecture Notes in Artificial Intelligence, pages 457-471. Springer, 2010.
http://dx.doi.org/10.1007/978-3-642-14203-1_39
Presents a novel blocking optimisation, proves that it does not adversely affect decidability, and demonstrates empirically the resulting improvement in HermiT's performance.
***[3] Birte Glimm, Ian Horrocks, Boris Motik, Rob Shearer, and Giorgos Stoilos. A Novel Approach to Ontology Classification. J. of Web Semantics, 10(1), 2011.**
http://dx.doi.org/10.1016/j.websem.2011.12.007
Presents a novel classification procedure, proves that it is worst case optimal, and demonstrates empirically the resulting improvement in HermiT's performance.
***[4] Motik, B., Horrocks, I., 2008. Individual reuse in description logic reasoning. In: Proceedings of the International Joint Conference on Automated Reasoning (IJCAR 2008). Vol. 5195 of Lecture Notes in Computer Science. Springer, pp. 242–258.**
http://dx.doi.org/10.1007/978-3-540-71070-7_20
Presents a novel individual reuse optimisation, proves that it does not adversely affect decidability, and demonstrates empirically the resulting improvement in HermiT's performance.
[5] Motik, B., Horrocks, I., 2008. OWL datatypes: Design and implementation. In: Proceedings of the 7th International Semantic Web Conference (ISWC 2008). Vol. 5318 of Lecture Notes in Computer Science. Springer, pp. 307– 322.
http://dx.doi.org/10.1007/978-3-540-88564-1_20
Describes how reasoning with datatypes, such as integers and strings, can be integrated into a logical ontology language such as OWL, and evaluates an implementation in HermiT.
[6] Motik, B., Cuenca Grau, B., Horrocks, I., Sattler, U., 2009. Representing ontologies using description logics, description graphs, and rules. Artificial Intelligence Journal 173 (14), 1275–1309.
http://dx.doi.org/10.1016/j.artint.2009.06.003
Presents an extension to the OWL ontology language designed to more faithfully represent complex structures, proves various computational properties, and evaluates an implementation that extends HermiT.

4. Details of the impact (indicative maximum 750 words)

The HermiT reasoner has made an important contribution to the development of ontology language standards and the uptake of ontology based systems: it was the first reasoner to support the OWL 2 standard (a revision of OWL published by the W3C in 2009) [A], it is used in an increasing number and range of applications, and it is distributed as the standard reasoner with the Protégé OWL ontology editor, a popular ontology design tool developed at Stanford University that has

more than 230,000 registered users drawn from every sector of the economy [B].

An example application of HermiT is in the Energy Management Adviser (EMA), a system developed by EDF Energy in France. EMA is designed to produce personalised energy saving advice for EDF's customers. In order to determine what kind of advice is appropriate to a given customer, each customer is classified based on a profile that includes information such as the kind of property they live in, their location, and their past energy usage. This classification uses an ontology that captures a wide range of information relevant to energy consumption. An important advantage of this design is that EDF's home management specialists can update and extend the ontology to reflect changing environmental circumstances and new energy saving ideas, and the reasoner will automatically adjust the nature of the advice given to customers to reflect these changes. Although hard to quantify, EMA has the potential to save substantial amounts of energy.

The letters are generated from the EMA based on information from Hermit which enables tailored advice to be provided based on knowledge about the customer. Hermit provides essential reasoning capability to make sense of this information and for EDF to be able to easily amend the reasoning techniques and advice based on emerging information about customer profiles and energy usage e.g. weather forecasts or new product information [H].

Typical advice would include "Electricity consumption between December and February is higher than during the same period last year. This increase may be caused by a lower out-door average temperature (2 degrees lower), which could have led to increased use of your central-heating system."
This would be followed by some general advice on how to make more efficient use of the heating system, including making a small reduction in the thermostat temperature. An interesting feature of the system is that it will omit this last piece of advice (reducing the temperature) for "vulnerable" customers, e.g. pensioners.

EDF uses Protégé and HermiT to develop and maintain their ontology, and they use HermiT in the EMA itself. The EMA is currently part of a premium service package called Suivi Conso that is an extra-cost option subscribed to by 265,000 EDF customers in France [C]. The EMA is used to provide advice to these customers at least twice per year, with the advice being distributed along with energy bills.

It is planned to extend this service to all 17 million EDF customers (which will require HermiT to classify more than 1 million customers per week), and also to use EMA to provide advice to customers via the EDF web site [D]. Horrocks and Motik are acting as paid consultants to EDF, helping them with the design of the EMA and of the associated ontology; some of this work has been documented in a joint EDF-Oxford paper to be presented at the 2013 International Semantic Web Conference [E].

Another example of HermiT's use in products and applications is the "ontology based automotive off-board diagnostic system" developed by Gigatronik Stuttgart GmbH. In this case the ontology is used to capture knowledge about relevant automotive systems and fault diagnoses. Reasoning (via HermiT) is used to support both ontology development (i.e., adding knowledge about faults and their possible causes), and fault diagnosis (i.e., by identifying different kinds of fault and suggesting possible causes) [F,G].

5. Sources to corroborate the impact (indicative maximum of 10 references)

[A] http://en.wikipedia.org/wiki/Semantic_reasoner
The Wikipedia entry that explains what an OWL reasoner is and presents a comparison of various reasoners, including HermiT.

[B] http://protege.stanford.edu
The web page of Protégé, a widely used ontology design tool that bundles HermiT as a reasoner.

[C] http://particuliers.edf.com/gestion-de-mon-contrat/ma-facture/suivi-conso-48007.html
The EDF web page that describes their "Suivi Conso" service that includes advice produced by the Energy Management Adviser driven by HermiT.

[D] Letter from Professor Stéphane Andrieux, Scientific Director of EDF R&D.
Letter from the Scientific Director of EDF confirming the role played by HermiT in EMA, and their ongoing relationship with the Oxford KRR group.

[E] P. Chaussecourte, B. Glimm, I. Horrocks, B. Motik, and L. Pierre. The energy management adviser at EDF. In Proc. of the 12th International Semantic Web Conference (ISWC 2013), Lecture Notes in Computer Science, 2013.
A paper describing the Energy Management Adviser that was co-authored with EDF researchers and published in the "in use" track at ISWC, the premier conference for Semantic Web research and applications.

[F] Gunther Sudra, Dominic Lyons, and Juergen Schwarz. Ontology-based Knowledge Representation for an Automotive Off-board Diagnostic System. In Proc. of the Tagung der Gesellschaft für Informatik, Workshop “Applications of Semantic Technologies” (http://ast2013.fzi.de/?page_id=7).
A paper describing the use of reasoning in the Gigatronik diagnostic system.

[G] Dominic Lyons. Konzeption und Entwicklung eines wissensbasierten Visualisierungsframeworks fu r eine interaktive Nachstellung von Diagnosefa llen. Studiengang Sofwaretechnik Fakulta t Informationstechnik Hochschule Esslingen.
An MSc thesis describing the use of HermiT in the Gigatronik diagnostic system.

[H] An example letter from EDF which is sent to customers tailored to themselves and their circumstances is held on record.