Distributed Knowledge Technologies for Intelligent Web Services

Alun Preece

Department of Computing Science

University of Aberdeen

The need for Intelligent Web Services

The growing web overwhelms users' abilities to locate and gather together disparate pieces of information.

Maintaining the web's usefulness as a resource requires us to develop a new generation of intelligent web services.

Aim: to turn the vast array of inhomogeneous sources into a coherent information base.

Approach: use distributed knowledge technologies.

Distributed knowledge technologies for intelligent web services

Knowledge fusion services: combine datasource querying & constraint solving to automate aspects of ecommerce.

Ontology-driven knowledge management: use ontologies/schemas to provide better knowledge management services.

Distributed knowledge systems architectures: build on web standards to allow integration of - and uniform access to - intelligent systems components.

Knowledge fusion services: the KRAFT approach

Designed to support distributed configuration design, within a multi-partner virtual organisation:

- manufacturing (PCs, network systems, ...)
- services (travel planning, meeting scheduling, ...)

Knowledge is exchanged between agents in the form of constraints expressed against an object data model.

Objects are defined in a shared ontology.

[KRAFT was funded by EPSRC & BT]

Knowledge fusion services: KRAFT architecture

The core KRAFT architecture provides services to:

- locate appropriate on-line sources of knowledge
- transform heterogeneous knowledge to a homogeneous constraint interchange format (CIF)
- fuse the constraints with associated data to form a dynamically-composed constraint satisfaction problem (CSP)
- harness existing constraint solver engines to compute CSP solutions

Knowledge fusion services: agent roles in KRAFT

The KRAFT architecture is agent-based:

Wrapper agents act as proxies for all participants in the virtual organisation, including vendors and customers.

Facilitator agents provide directory and matchmaking services for the marketplace.

Mediator agents:

- provide "value adding" services (e.g. reselling)
- enforce rules of the marketplace

Knowledge fusion services: "small print" constraints

Instances from vendors' product catalogues often have associated constraints (the "small print").

In natural language:

all of Storage Inc's Zip disk drives require a PC with a USB-type port

In KRAFT-CIF:

```
constrain each d in disk_drive
  such that name(vendor(d))
      = "Storage Inc" and type(d) = "Zip"
  at least 1 p in ports(host_pc(d))
      to have type(p) = "USB"
```


Knowledge fusion services: a generic KRAFT application

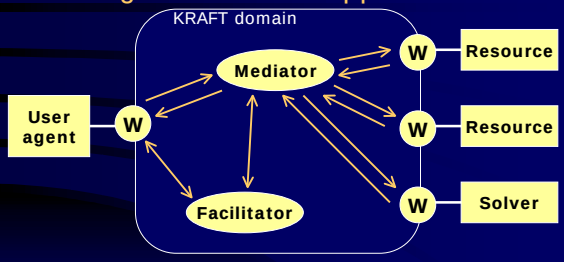

[See Preece et al, *Int J Cooperative Information Systems*, 2001]

Knowledge fusion services: implementation & evolution

Prototype uses KQML-97 for communication:

original version used Prolog term structures

new version uses XML/RDF (including CIF)*

[See Gray et al, *IJCAI-01 EBIWEB workshop*, 2001]

Hub-based physical network design:

original version used Linda servers as hubs

new version uses XML messaging (AKT-bus)*

Mediators implemented in Prolog, Java, CLIPS

*Ongoing work in EPSRC-funded AKT project

Knowledge fusion services: future work

Extend the simple interaction protocols to incorporate more sophisticated negotiation.

Take account of quality-of-service metrics in forming virtual organisations.

Allow for dynamic restructuring of virtual organisations at operation-time.

Test for scalability (no. of agents; amount of info).

Collaborators:

Southampton Uni (Jennings)
Cardiff University (Gray)
BT

Ontology-driven knowledge management

Ontologies are conceptual models of information domains – similar in form & intent to schemas.

Every knowledge or information resource has an ontology, even if it's implicitly defined.

By creating explicitly-defined ontologies, we aim to:

provide metalevel descriptions of what our resources contain – to assist resource location

enable the principled integration of resources

Ontology-driven KM: recent work (1)

Baker Hughes OASIS TCS Project

Built a large ontology (1200 concepts, 240 relations) using LOOM, in the domain of drilling process improvement.

Captured case bases of valuable experts' experience using the ontology.

Subsequently reimplemented & deployed in (less expressive) Lotus Notes / Domino.

[See Preece et al, *IEEE Intelligent Systems*, 2001]

Ontology-driven KM: recent work (2)

DISCOVER

- A tool to help integrate disparate knowledge bases under a common ontology.
- Verifies whether a given knowledge base & set of mapping rules is consistent with a given common ontology.

Reconciliation of Experts' Ontologies

- Goal is to articulate & help resolve mismatches between ontologies acquired from domain experts.
- As a precursor to integrating their knowledge.

Ontology-driven KM: future work (1)

Baker Hughes – Phase II

- Integrate an expressive knowledge representation & reasoning system with Lotus Notes / Domino.
- Use this as the basis for ontology-driven services:
  - Just-in-time task-based knowledge delivery
  - On-the-job knowledge capture
  - Quality assurance of knowledge assets

Ontology-driven KM: future work (2)

Empirical Study of Ontology Added Value

Web visionaries are promoting a semantic web, in which ontologies improve information retrieval (IR) and processing.

We're seeking to:

1) test the extent to which ontologies actually improve IR in the semantic web context
2) explore the costs involved in using ontologies during IR

Summary: Distributed knowledge technologies for intelligent web services

Knowledge fusion services: combine datasource querying & constraint solving to automate aspects of ecommerce.

Ontology-driven knowledge management: use ontologies/schemas to provide better knowledge management services.

Distributed knowledge systems architectures: build on web standards to allow integration of - and uniform access to - intelligent systems components.

Credits

Work done in the context of the following dept'l research groups:

- Knowledge-Based Systems Group
- Databases & Distributed Information Systems Group

Main collaborators: (Apologies to anyone I've missed!)

- (Staff) Pete Edwards, Peter Gray, Graham Kemp, Kit Hui, Tim Norman, Derek Sleeman
- (Students) Alan Flett, Gunnar Grimnes, Adil Hameed, Andrew Waterson
- (Other universities) Cardiff University, Liverpool University, AKT Consortium: Edinburgh, OU, Sheffield, Southampton
- (Companies) Baker Hughes, BT

Questions & comments?