Autonomy, the Web and Knowledge-based Services



Dr Terry R. Payne Department of Computer Science





Autonomy, the Web and Knowledge-based Services

- Perspectives on Agents & Knowledge based services
 - 1. Interface Agents & the Web/Internet
 - 2. Multi-Agent Systems & the need for Semantics
 - 3. Services and the need for Semantics
 - 4. Autonomous Agreements on Knowledge



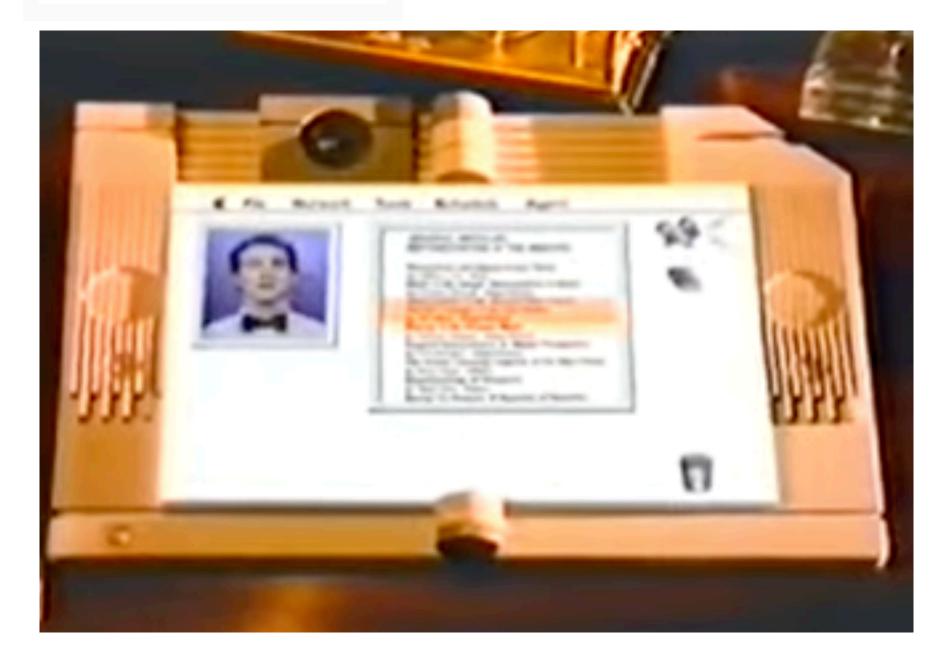
Agents and Nicholas Negroponte

"...The 'agent' answers the phone, recognises the callers, disturbs you when appropriate, and may even tell a white lie on your behalf. The same agent is well trained in timing, versed in finding opportune moments, and respectful of idiosyncrasies..." (p | 50)

"... If you have somebody who knows you well and shares much of your information, that person can act on your behalf very effectively. If your secretary falls ill, it would make no difference if the temping agency could send you Albert Einstein. This issue is not about IQ. It is shared knowledge and the practice of using it in your best interests..." (p I 5 I)

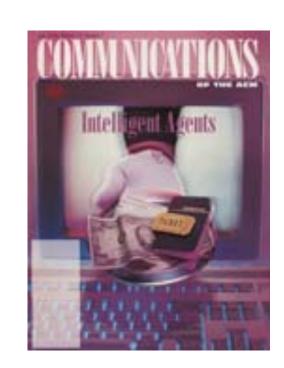
"...Like an army commander sending a scout ahead ... you will dispatch agents to collect information on your behalf. Agents will dispatch agents. The process multiplies. But [this process] started at the interface where you delegated your desires..." (\$\rm 158\$)

Apple Knowledge Navigator (1987)

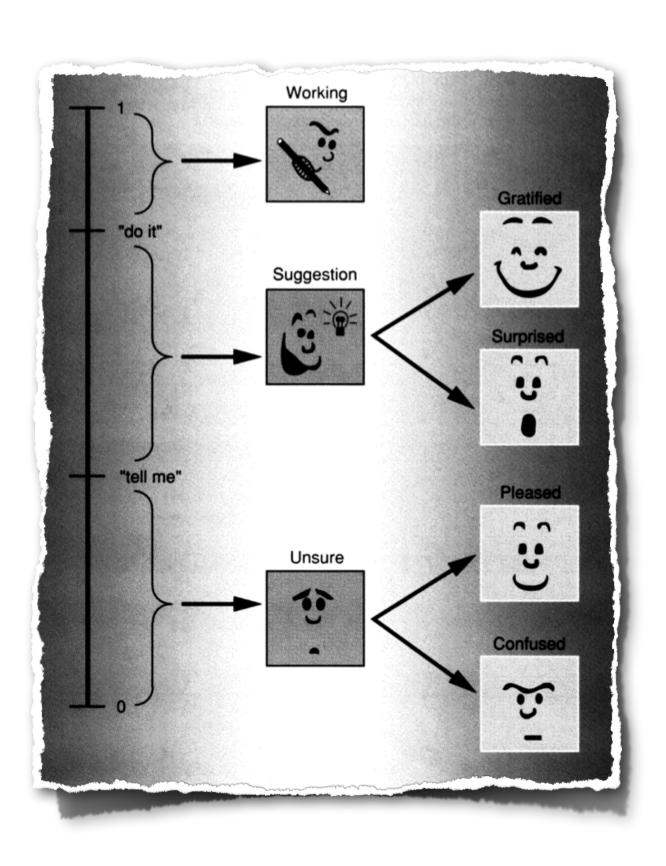


(From **Being Digital**, 1985)

Interface Agents (Comms ACM 1994)



- Agents sat 'over' applications, watching, learning, and eventually doing things without being told — taking the initiative.
 - Pioneering work at MIT Media Lab (Pattie Maes):
 - USENET news reader & mail readers
 - Web browsers
 - Early recommender systems
 - Other work from Katia Sycara, Oren Etzioni, Daniel Weld & Tom Mitchel
- Utilised ML to adapt to user preferences in doing and automating tasks
 - No MAS "Agency"



Distributed Systems & Multiple Agents

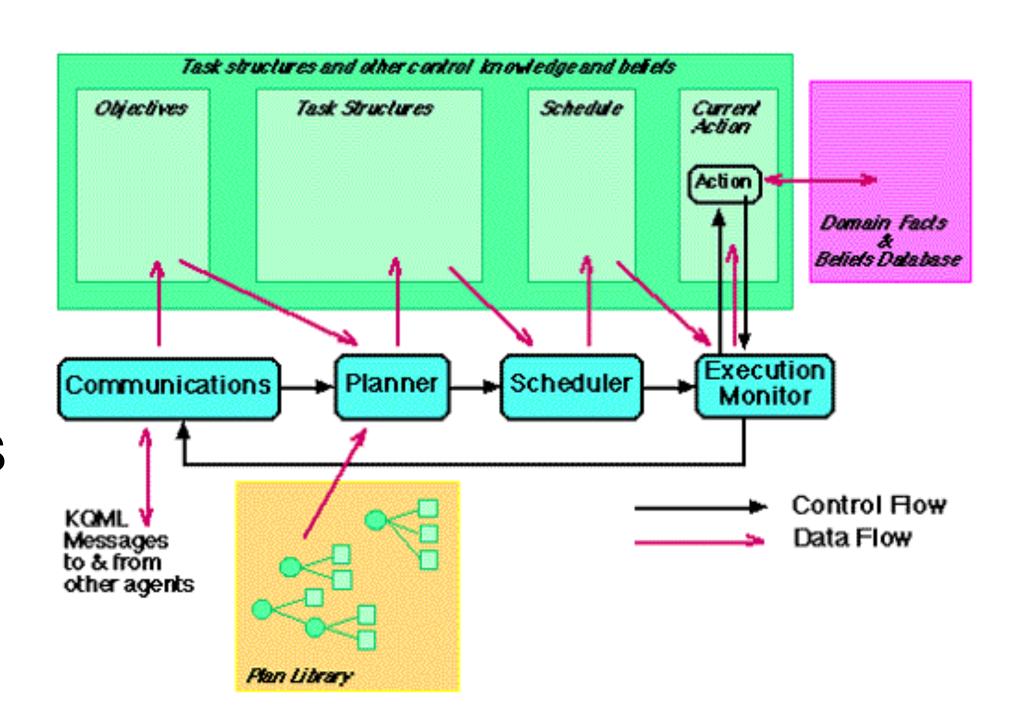
- Multi-Agent Systems emerged as a way of solving complex distributed problem
 - Agency focused on the *autonomous* coordination of decision making and control, between different agents (without human intervention)
- Need to model knowledge of environment and knowledge of other agents
 - Not just of capability, but of reliability, veracity etc
 - Typically making syntactic/semantic interoperability assumptions



A Theory of Agency

- MAS was Classic Al... but distributed!
 - Often built using a modular BDI structure
- Theoretical models often focussed on
 - Modal Logics for distributed beliefs
 - Game Theoretic Models for coordination and negotiation
- Implemented Models considered practicalities
 - Modelling Environment
 - Discovery Services
 - Managing Communication

Retsina Agent Architecture (Katia Sycara) https://www.cs.cmu.edu/~softagents



Curse of Heterogeneity

- Many such MAS infrastructures emerged
 - Centralised communication vs Peer-based communication
 - e.g. **OAA** Martin et. al., **Retsina** Sycara et. al.
 - *Team-based* organisation vs *Institutions*
 - e.g. **STEAM teamwork model** Tambe et. al., **Electronic Institutions** Sierra et. al.
- Communication through ACLs
 - KQML / KIF, developed by the DARPA-funded Knowledge Sharing Effort (KSE)
 - Helped resolve challenges with syntactic interoperability (somewhat...!)
- Content/knowledge suffered from a lack of semantic interoperability

Aligning Multi Agent Systems

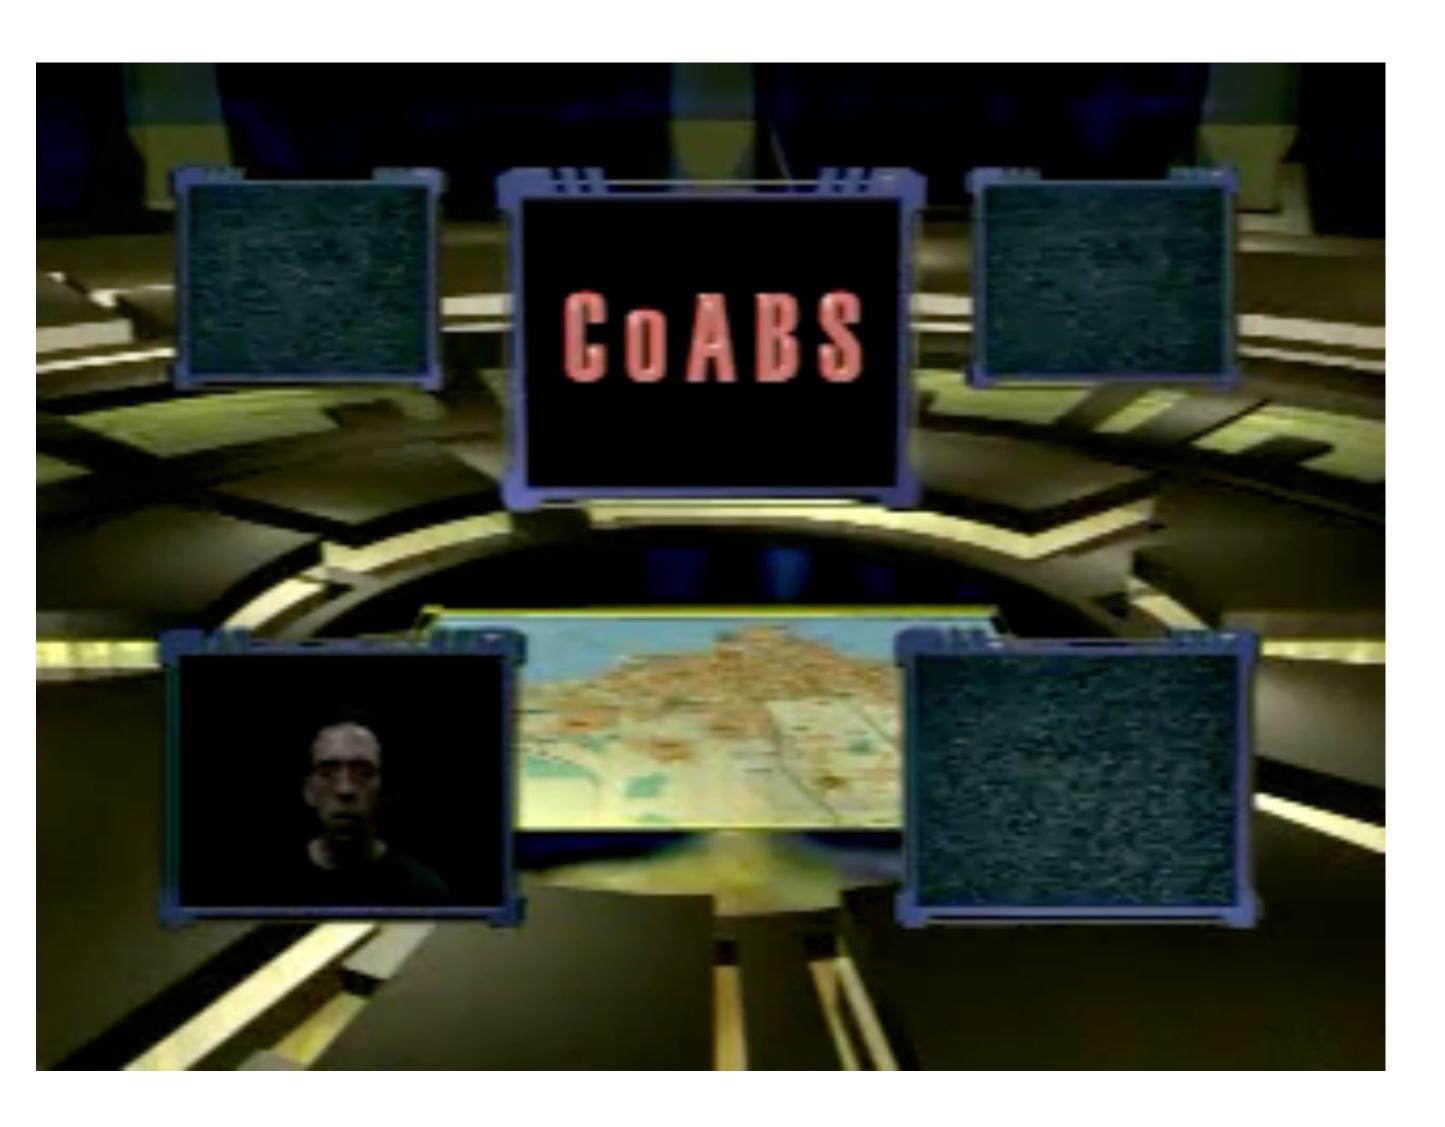
Control of Agent Based Systems (CoABS)

- DARPA Vision:
 - to bring together multiple, heterogenous technologies to autonomously respond to dynamic and evolving scenarios.
- Developed middlewear for architecture heterogeny, security, logging, & discovery
 - But it was difficult to debug, and didn't address the semantics of knowledge



Jim Hendler - CoABS & DAML Program Lead

CoABS TIE - NEO (1999)



- Facilitated Interoperation through Middlewear
 - CoABS Grid
 - Interoperator Agents
- In reality...
 - ... only if the agents pre-agreed:
 - the *ontologies*
 - and their semantics!

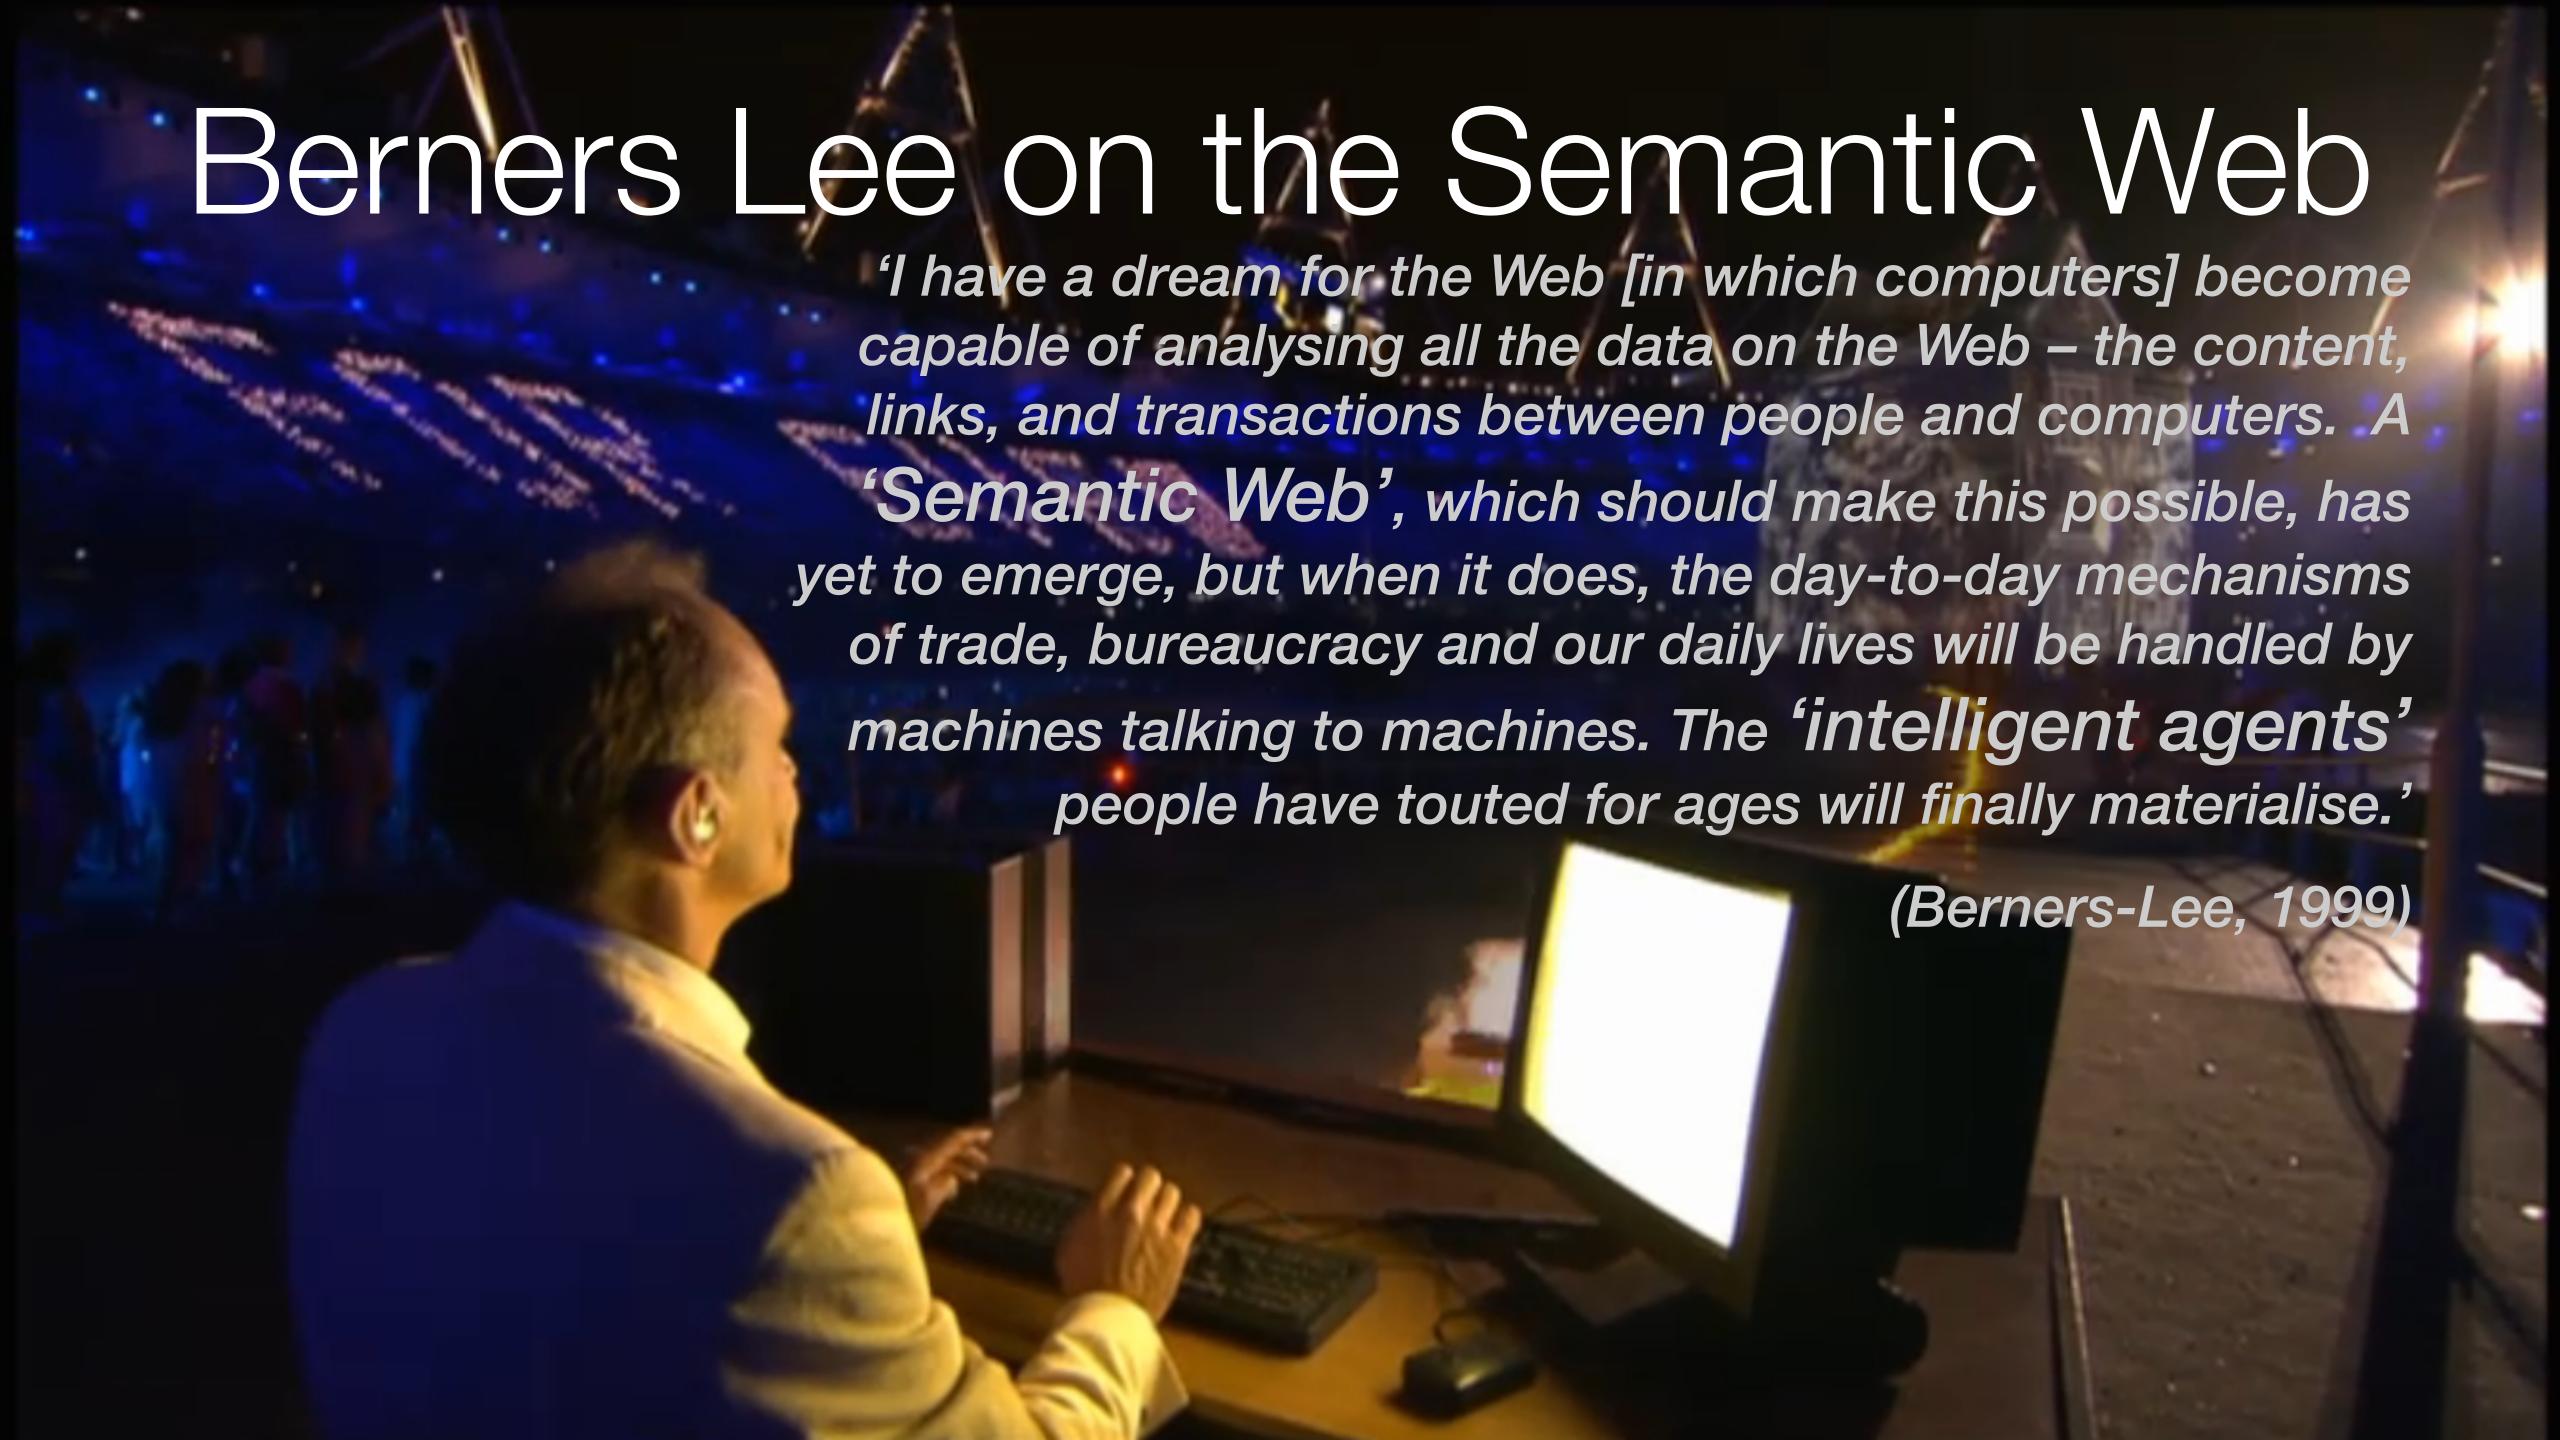
Tackling Semantic Interoperability...

- Semantic Interoperability was a major hurdle for
 - Locating Services
 - Different terms used for advertisements and requests
 - Negotiating contracts & communications
 - Different protocols used by different communities when agreeing whether to transact
 - Invoking
 - Constructing valid messages based on the published signature/interface of a service
 - Understanding
 - Interpreting the results of invoking a service
 - Composing Services
 - Constructing plans to achieve meta-goals based on available Services/Agents

DARPA Agent Markup Language

- Proposed by J. Hendler to address the knowledge problem for agents
 - Distributed ontologies using Web-based links
 - Based on DLs, but introducing new challenges
 - Helped kick-start the Semantic Web





DARPA Agent Markup Language

- Focussed on the challenge of representing *knowledge* for the *web*
 - DAML-ONT (2000)
 - Built upon RDF(S) Ora Lassila & Ralph Swick
 - DAML+OIL (2001)
 - Joint EU/US Committee brought together DAML-ONT and OIL from the EU funded Onto-Knowledge project
 - OWL (2004)
 - Three flavours: Lite, DL and Full

Date: Tue, 10 Oct 2000 12:28:34 -0400
To: www-rdf-logic@w3.org
From: Jim Hendler < jhendler@darpa.mil>
Subject: 1st public release - DAML-ONT

As program manager of the "DARPA Agent Markup Language" initiative, I am pleased to announce the first release of a significant piece of the DAML language. DAML is an effort to help bring the "semantic web" into being, focusing on the eventual creation of a web logic language. It represents joint work between DoD, industry and academia in both the US and the European Community and we hope it will lead to the eventual web standard in this area.

This first release is a draft language for the "ontology" core of the language (roughly corresponding to a frame-based or description logic starting place) - this allows the definition of classes and subclasses, their properties, and a set of restrictions thereon. The language does not yet include any specification of explicit inference rules, which we hope will follow.

We believe this ontology core will be a useful starting place for extending the language, and for experiments in a web-based semantic language that is accessible to a wide audience.

The language is based on RDF (actually RDF Schema) and extends it to include new concepts.

A group of researchers from MIT's Laboratory for Computer Science had primary responsibility for developing this language core. This was then extended by a number of people including representives from the OIL effort, the SHOE project, the KIF work, and DAML contractors. Some of the primary contributors are listed below.

Details of the language release and pointers to the above named projects (and others) can be found at http://www.daml.org

- Jim Hendler DAML Program Manager

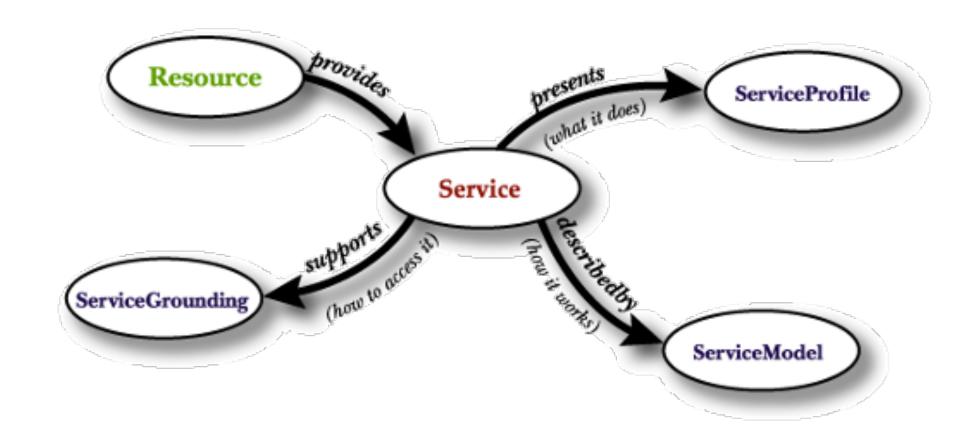
}-----

Preliminary release - DAML-ONT 0.5

DAML-S/OWL-S

- The DAML-S coalition tasked with developing ontologies for services
 - Adopted a mixed agent / planning / workflow perspective
 - Focus on the notion of:
 - Discovery
 - Composition
 - Execution

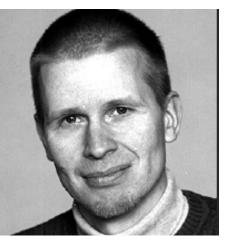
Anupriya Ankolekar, Mark Burstein, Jerry R. Hobbs, Ora Lassila, David Martin, Drew McDermott, Sheila A. McIlraith, Srini Narayanan, Massimo Paolucci, Terry Payne & Katia Sycara: "DAML-S: Web Service Description for the Semantic Web", ISWC 2002

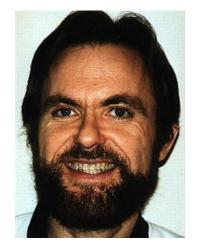
















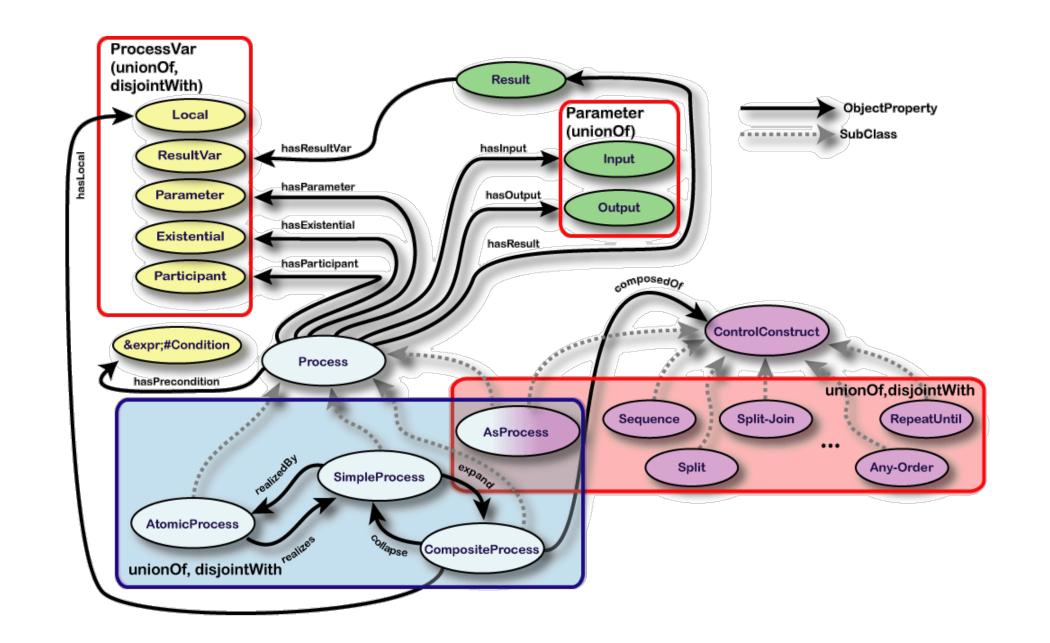


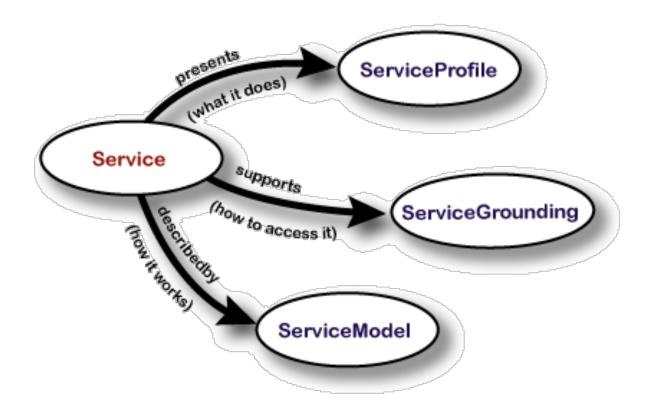




Workflows as Plans, Services as Agents

- Abstract (ontological) workflows
 - Each service element as a description of the runtime service to be provisioned
 - Model functional and QoS parameters
 - Knowledge-based descriptions of service interfaces and invocation mechanisms
 - Interoperability through *shim/mediator* services





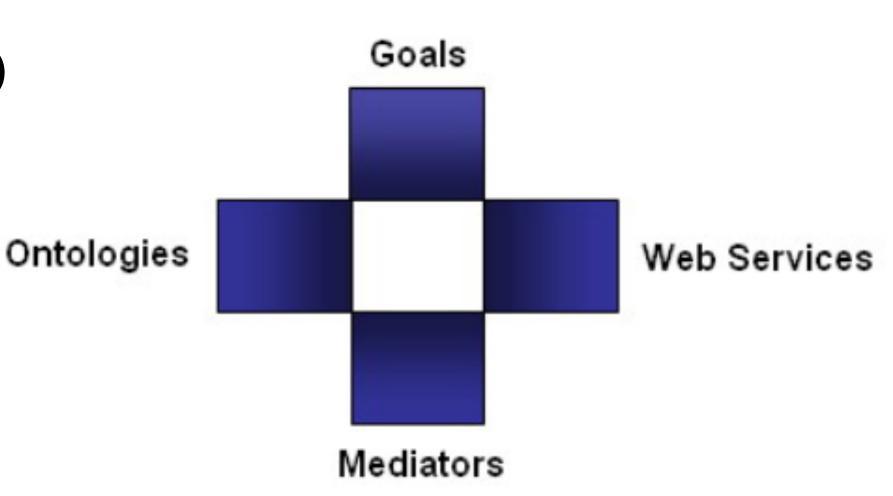
- Ontological Models to support Orchestration
 - Semantic Web Service Frameworks (OWL-S & WSMO)
- Run-time Dynamic Service *Discovery*
 - Supports re-planning given failure...!

Semantic Web Service Challenges

- DAML-S / OWL-S captured the moment
 - Excellent research outputs from other efforts:
 - iBrow and PSMs, WSMO, etc
 - Composition was captivating
 - Provided one of the de-facto standards for services in the semantic web
 - the other at the time being WSMO
- Distributed services and devices were emerging
 - eScience Grid with physical lab equipment and in-silico eScience services
 - proliferation of web services

Mediation and Alignment

- OWL-S had no explicit consideration for mediators
 - Left to other services or the agent
- Contrasts with the explicit modelling of WSMO
 - **ggMediators**: linking two goals
 - ooMediators: resolve representation mismatch between ontologies
 - wgMediators: map goals to web-services
 - wwMediators: linking two web services
- Semantic Interoperability in open systems still a fundamental challenge
 - reliability / serendipity of services



The top-level elements of WSMO https://www.w3.org/submissions/WSMO/

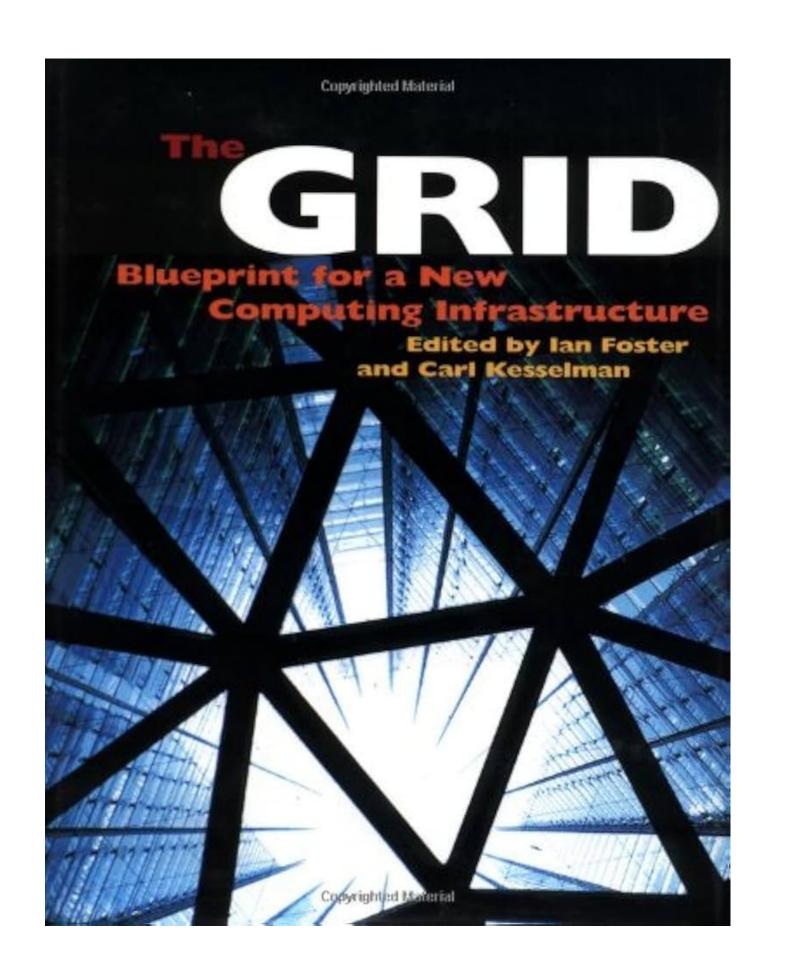
Back to Services... and agents

The Grid and Multi Agent Systems

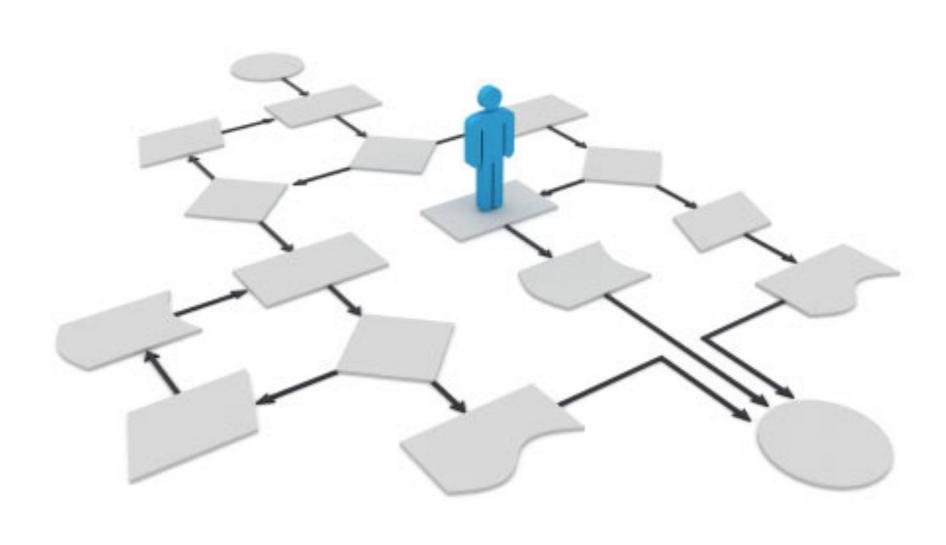
'The [eScience] Grid and agent communities are both pursuing the development of such open distributed systems, albeit from different perspectives. The Grid community has historically focussed on [. . .] "brawn": interoperable infrastructure and tools for secure and reliable resource sharing within dynamic and geographically distributed virtual organisations (VOs), and applications of the same to various resource federation scenarios.

In contrast, those working on *agents have focussed* on "*brains*", i.e., on the development of concepts, methodologies and algorithms for autonomous problem solvers that can act flexibly in **uncertain and dynamic environments** in order to achieve their objectives.'

(Foster et al, 2004)



Scientific workflows and services



- Processes, resources, tools as Services
 - Available *locally* as Virtual Organisations or *commercially* through 3rd parties
 - Encourages *reuse* of:
 - Large-scale *resources* (e.g. HPC) and *equipment* (e.g. IoT)
 - Terabytes of *Data*, both raw (warehoused) & inferred knowledge (Big Data)
- Frameworks exist to find & assemble workflows
 - Discovery using knowledge-based frameworks through workflow editors (myGrid Taverna)
 - Still the issue of:
 - availability, provisioning, scheduling, interoperability, unreliability!

Run-time provision of web services

- Based on Service Demand
 - Provisioning of services based on Abstract Workflows
 - Selection of services based on multiple factors such as cost, duration, reliability

Agent Task/ Service/ Capability

Task Modelling and Discovery

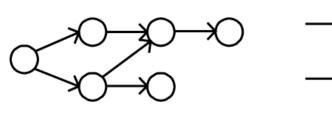
- Service Description / Modelling
- Discovery Architecture

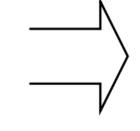
Task Facilitation

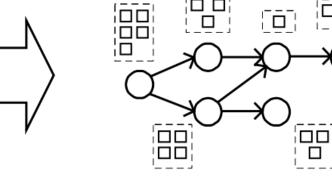
- Resolving Semantic Heterogeneity through Ontology Alignment using Dialogue, Negotiation and Argumentation
- Resolving Syntactic Heterogeneity through ACL and XML refactoring

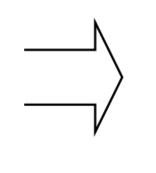
Task Provision

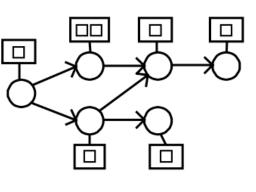
- Provisioning task supply based on dynamic demand
- Adaptive task use in noisy/unreliable environments.

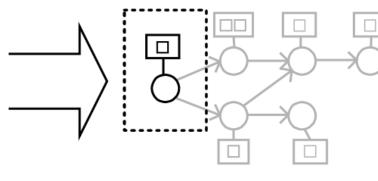










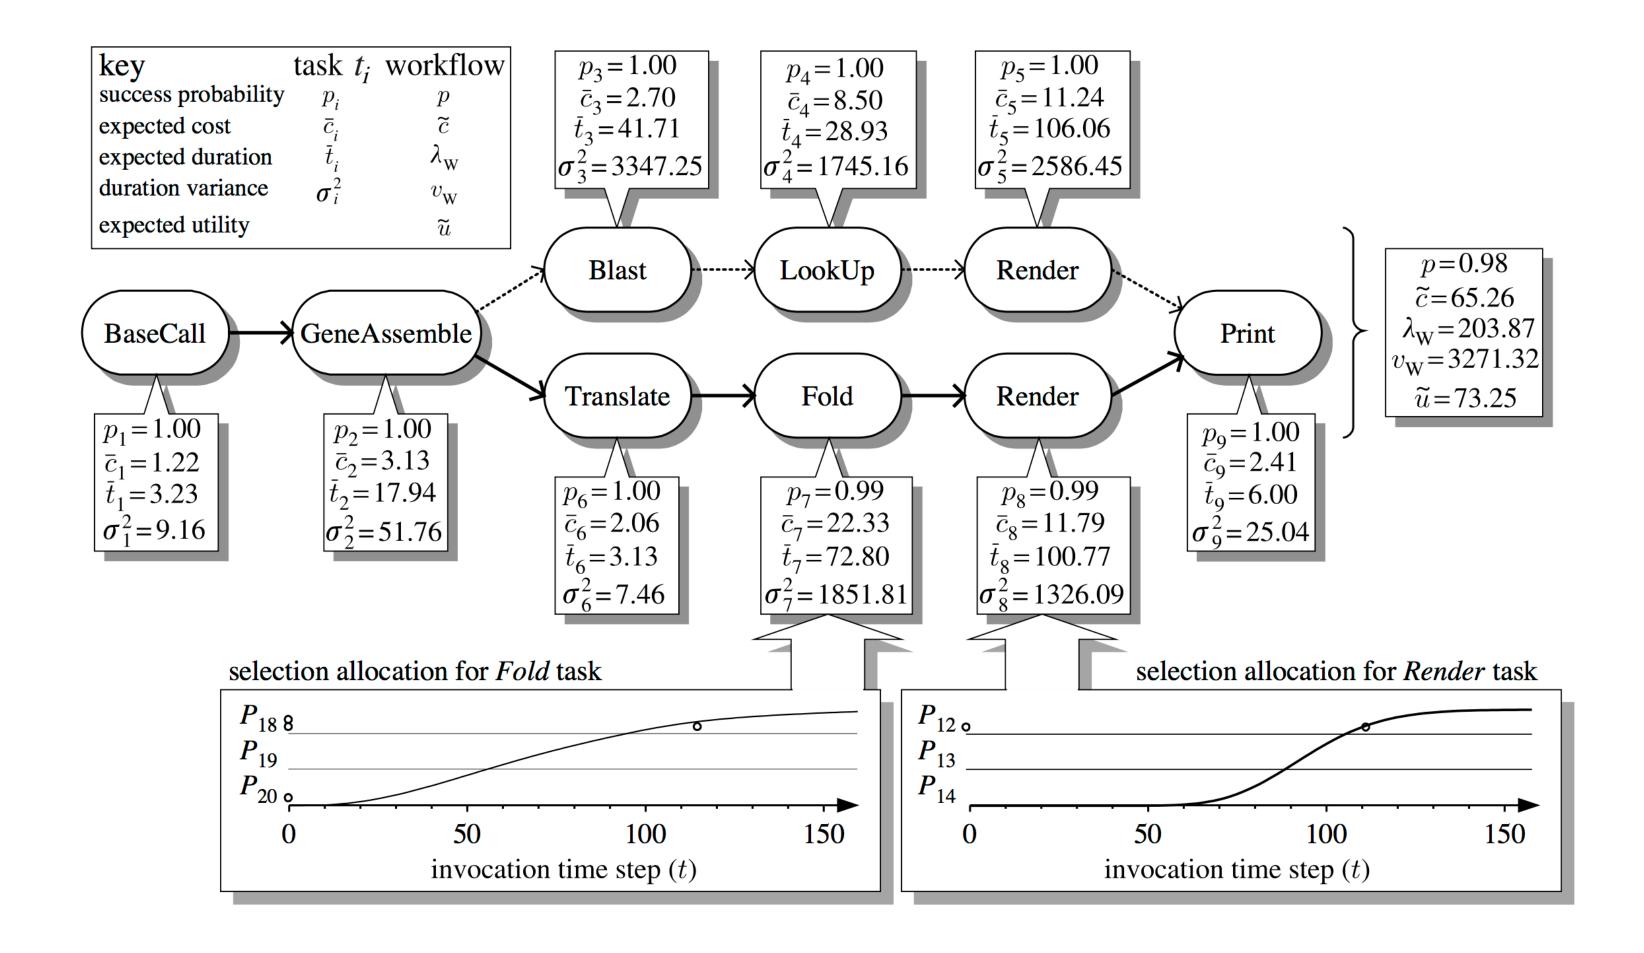


service discovery

service selection

invocation

Agent based Service Provisioning



Sebastian Stein, *Terry R. Payne*, and Nicholas R. Jennings. Flexible Selection of Heterogeneous and Unreliable Services in Large-Scale Grids. *Philosophical Transactions of the Royal Society A: Mathematical, Physical & Engineering Sciences*. 367 (1897). pp. 2483-2494, 2009.

Open (Web-scale) Environments

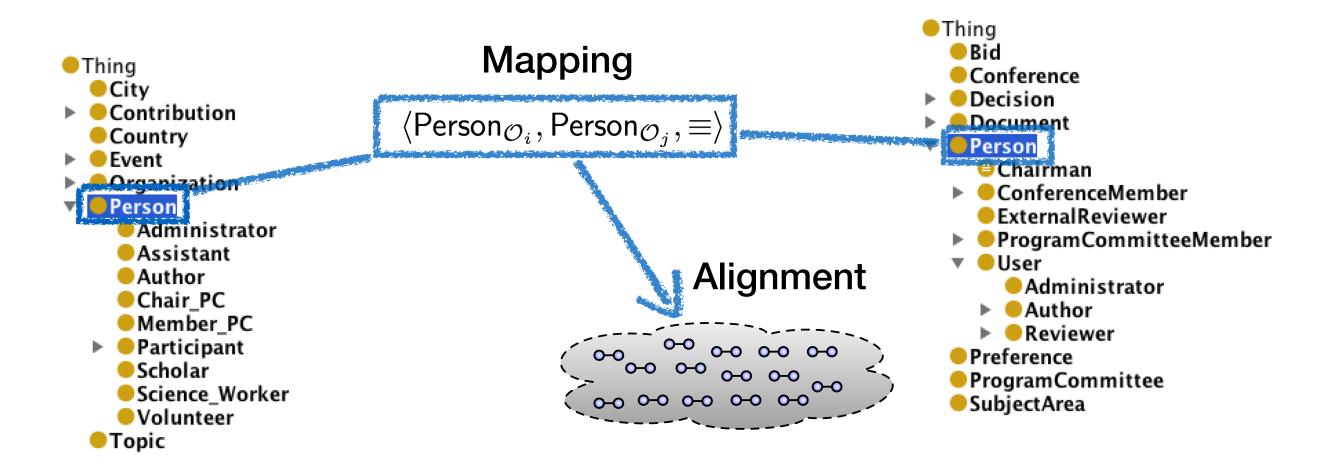
- Large, open environments can be characterised as...
 - Distributed Data Sources
 - Many different applications/services/devices/sensors
 - Owned by different stakeholders, and performing different tasks
 - Heterogeneous Ontologies
 - Multiple ontologies modelling data for similar domains
 - Transience, Autonomy and Opportunistic Behaviour (IoT)
 - Devices moving through an environment may need to provision services at runtime without human intervention
 - Danger of misinformation to the user, or exploitation of user data

A plethora of data models

- Different systems (sensors, services, applications, agents...) that generate or use knowledge usually make use of different models i.e. *ontologies*
 - Similar or overlapping information is modelled in diverse ways even inside organisations with strong governance and internal communication
- These differences in modelling become apparent when
 - These systems must be combined (*integration*)
 - Or be made to work together (interoperation)
- Ontology alignment is the process of determining mappings between semantically related entities
 - classes, relationships and instances

Open Systems, Ontologies and Alignment

- Systems can assume different ontological models
 - Modelled implicitly, or explicitly by defining entities (classes, roles etc), typically using some logical theory, i.e. an Ontology
- Alignment Systems align similar ontologies



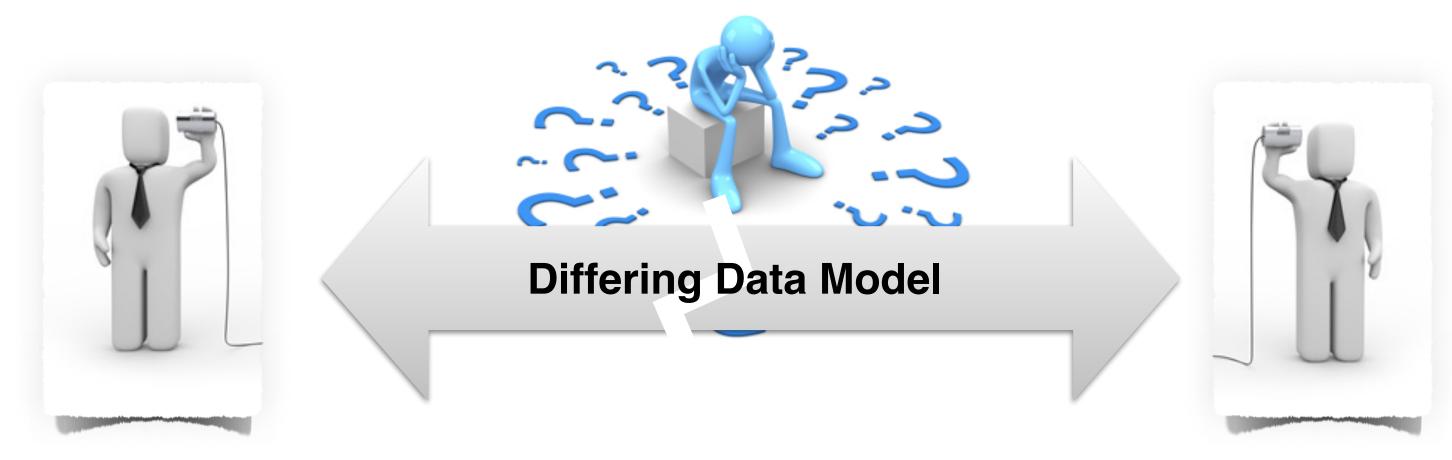
 If we assume that different alignments exist, how do agents choose which to use?

Alignments as shims



- For many workflows, shims are used to connect services
 - They map the data from one format to another
 - Syntactic mapping
 - Semantic mapping
- The notion of shims or mediators is central to Semantic Web Services
 - OWL-S assumes either that:
 - agents can manage the mediation themselves
 - mediators can be represented as services in workflows
 - WSML explicitly models mediators as part of the workflow model

Challenges in using alignments



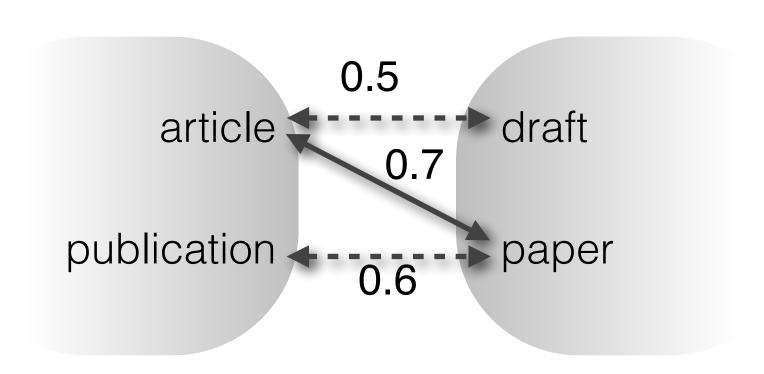
- Does the task affect the mappings used?
 - Will the inclusion of certain mappings introduce logical violations?
 - Different alignment mechanisms favour different mappings, even with the same ontologies!
- Fragments of the ontological space may be confidential, or commercially sensitive.
 - Axiomatic disclosure or exposure is potentially problematic

Just in time alignment

- Do we need everything to be aligned?
 - An agent may aggregate several ontologies for a variety of domains
 - A task may be relevant to only a single fragment within an ontology
 - Fragments of the ontological space may be confidential, or commercially sensitive.



Selecting the right mappings



- Quality vs Quantity
 - Do we maximise coverage
 - Preferable when merging the whole ontology
 - Do we find the "best" mappings
 - Preferable when aligning specific signatures



What alignments can we use?



- What does the agent know?
 - Pre-computed alignments exist, and can be shared and aggregated
 - But... different agents may possess different alignment fragments from different sources.
 - We need to choose the correspondences that are more useful
 - wrt to some criteria

Dialogues for collaborative decision making



• Why dialogues?

- Collaborative decision making mechanism
 - Public knowledge vs private strategy
 - Agents disclose *preferences* on correspondences, terms, axioms
 - Preferences are determined privately through strategic reasoning
 - Allows agents to reason with different types of information

Traceability

- The dialogue itself states publicly the commitments of each agent to the negotiation
- Explanations can be generated as to why the resulting mappings are determined

Approaches to negotiating mappings



- Aggregating different mappings
 - Agents selectively identify what mappings should be disclosed
 - Can choose what to share with the opponent, and what is private knowledge
 - Adapts mappings in an evolving environment

- Discovering novel mappings
 - Offer *limited ontological knowledge* of seed entities
 - Agents selectively share conceptual knowledge to identify localised structural similarity
 - Bootstraps the process of aligning different data systems

Dialogical approaches to ontology alignment @ Liverpool

- Dialogical approach for alignment aggregation
 - Construct a mutually acceptable alignment comprising selected correspondences from existing prior alignments.
- Dialogical approach for alignment repair
 - Cooperative decision making mechanism to repair an inconsistent alignment generated through aggregation.
- Dialogue protocol to support meaning negotiation
 - Repeated protocol for cooperatively agreeing on a correspondence with limited knowledge of the agents' ontologies.

Dialogues for finding Correspondences between partially disclosed Ontologies

- Allows two agents to exchange knowledge about *known mappings*
 - Each agent suggests preferred mappings
 - Agents negotiate over alternatives
 - Through argumentation / persuasion
 - Agents control what knowledge they disclose
 - Aligns only those entities in each agents' working fragments,
 without disclosing the ontologies, or all known mappings
 - Allows agents to suggest repairs if a mapping introduces a logical violation



Dialogue-Based Meaning Negotiation



- Explores a cognitive approach to reaching consensus over possible correspondences...
 - Agents identify possible concepts that may be ontologically equivalent in their respective ontologies
 - Each then seeks further evidence over the *locality* of each concept to verify if these are structurally similar.
 - Both agents have the opportunity to ask questions
 - Correspondences only accepted if both agents accept the same underlying support

Dynamic Alignment Open Problems

- Strategies can be adopted to determine what axioms to divulge
 - Requires quantifying or ordering the axioms in some way
- Utility gained from disclosure may vary as the negotiation progresses
 - Mechanism Design may discourage deception
 - Incentivise truth telling
 - However, such approaches can be NP-hard
- Evaluating the validity of the mappings
 - Strategies should evolve based on past performance
 - Thus the need for additional provenance models on alignments

Conclusions

- Semantic interoperability remains one of the key enablers for Hypermedia Multi-Agent Systems.
- Future systems must balance autonomy, privacy, and trust while operating at Web scale.
- The vision of agents acting flexibly through the Web is achievable but requires:
 - Robust alignment frameworks
 - Governance mechanisms for open hypermedia environments
 - Integration of Al reasoning with Web standards
- The Role of LLMs and GenAi?
 - That is still an open question:)

