Approaches to Resource Discovery in Computational Grids

Anurag Bhardwaj
CSE Deptt. University at Buffalo
Fall - 2006

Objectives

- Introduction
- MDS Architecture
- Other proposed approaches
- Scope / direction of future work
- Conclusion

What is Resource Discovery

- Involves manipulation of heterogeneous resources
- May require an indexing service
- Need to implement registry as well as user-query subsystems
- May be Centralized, Decentralized or a combination of both (Hierarchical).
- Should address performance,security,scalability and robustness requirements.

MDS @ Globus

- Contains Index Service as well as Trigger Service
- Index service collects data from various sources and provides a query/subscription interface to that data
- Trigger service collects data from various sources and can be configured to take action based on that data
- An *Archive Service*, which will provide access to historic data, is planned for a future release.

Architecture Overview

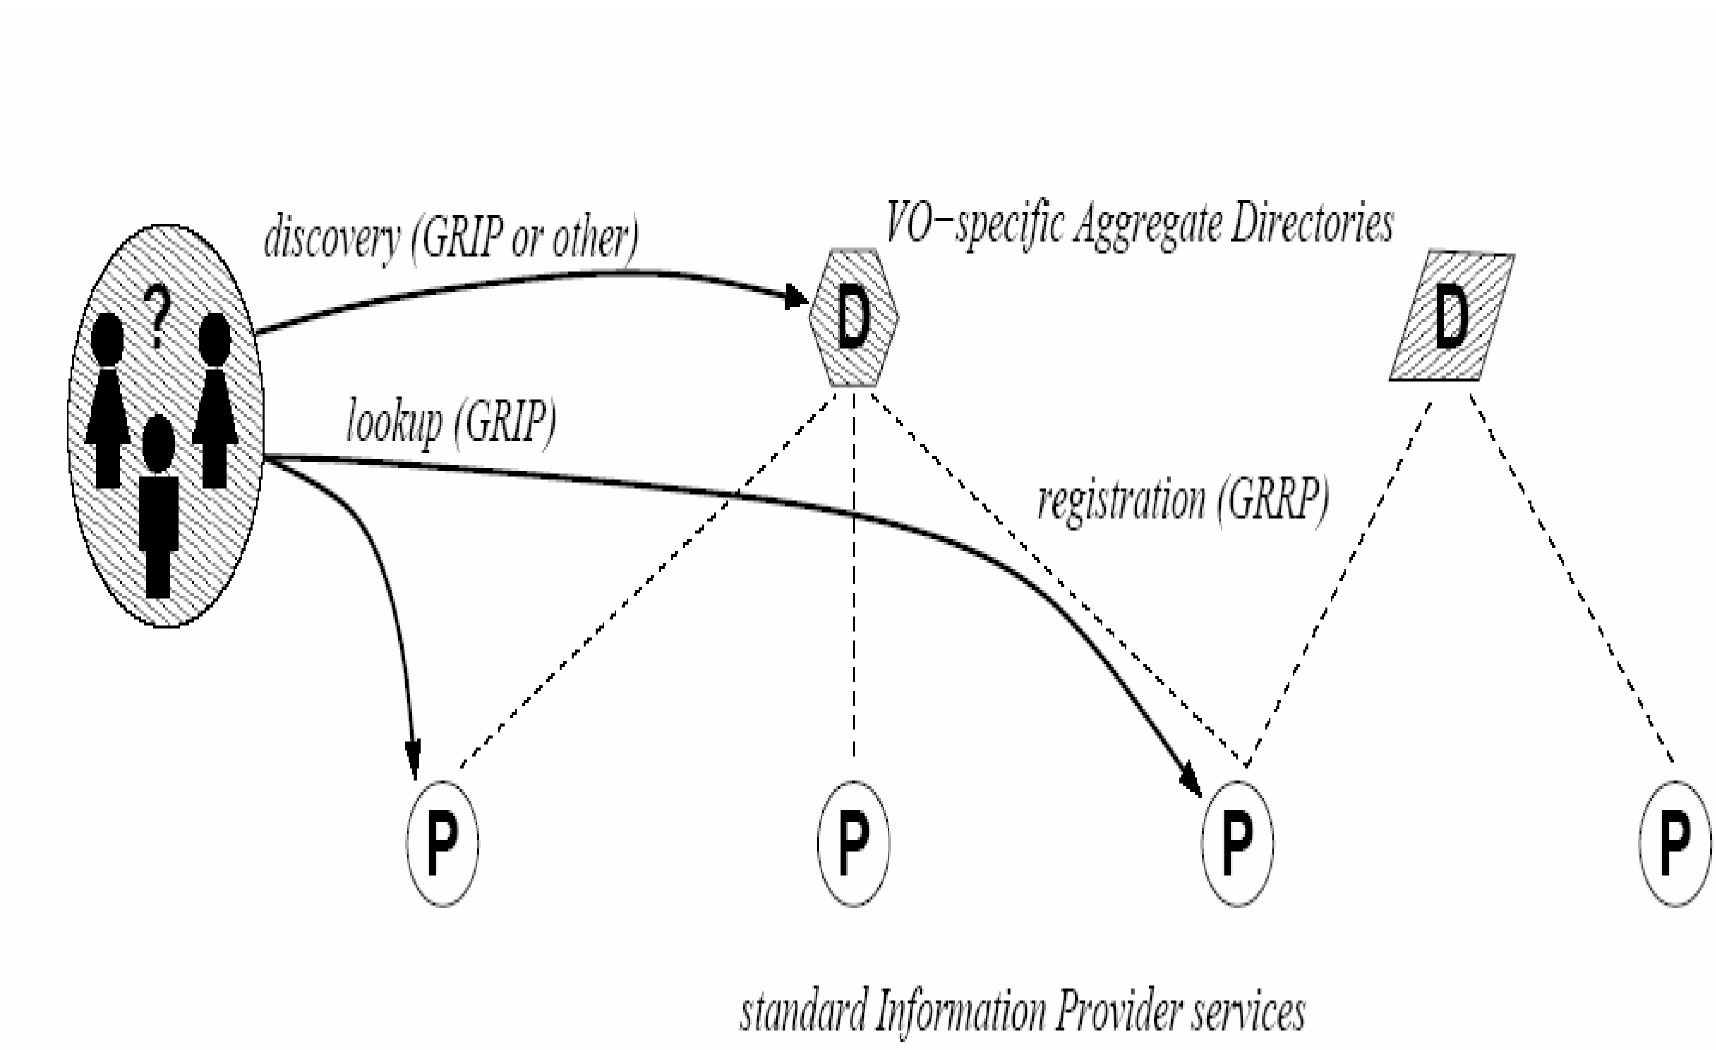

MDS Architecture

- Information Providers - Provide access to information about individual entities. information is structured in terms of a standard data model taken from a LDAP.
- Aggregate Directory Services - Facilitate resource discovery and monitoring for VO by implementing both generic and specialized views and search methods for a collection of resources.
- Grid Information Protocol - Used to access information about entities
- Grid Registration Protocol - Used to notify aggregate directory services of the availability of this information.

Information Providers

- Providers form a common, VO neutral infrastructure providing access to detailed, dynamic information about Grid entities.

- A provider for a compute resource might provide information about the number of nodes, amount of memory, operating system version number and load average; a provider for a running application might provide information about its configuration and current status.

- An information provider is defined as a service that speaks two basic protocols. The GRid Information Protocol (GRIP) is used to access information about entities, while the GRid Registration Protocol (GRRP) is used to notify aggregate directory services of the availability of this information.

Aggregate Directories

- Provide often specialized, VO-specific views of federated resources, services, and so on.
- For example, a directory intended for use by a broker might list the computers available to a VO organized by operating system type; another directory, intended to support application monitoring, might keep track of running applications.
- We define an aggregate directory as a service that uses GRRP and GRIP to obtain information (from a set of information providers) about a set of entities, and then replies to queries concerning those entities.

GRIP (Grid Information Protocol)

- Supports both discovery and enquiry.
- Discovery is supported via a search capability.
- Enquiry corresponds to a direct lookup of information.
- For example, consider an information provider that maintains information on a set of workstations. A broker might then perform a search on that provider to obtain a set of results that roughly match a given criteria. From the set of discovered resources, enquiry can be used to refine the set of resources upon which a broker may schedule.
- MDS adopts standard Lightweight Directory Access Protocol (LDAP) as the protocol for GRIP.
- Limitations – Cannot specify relational joins !

The LDAP Data Model

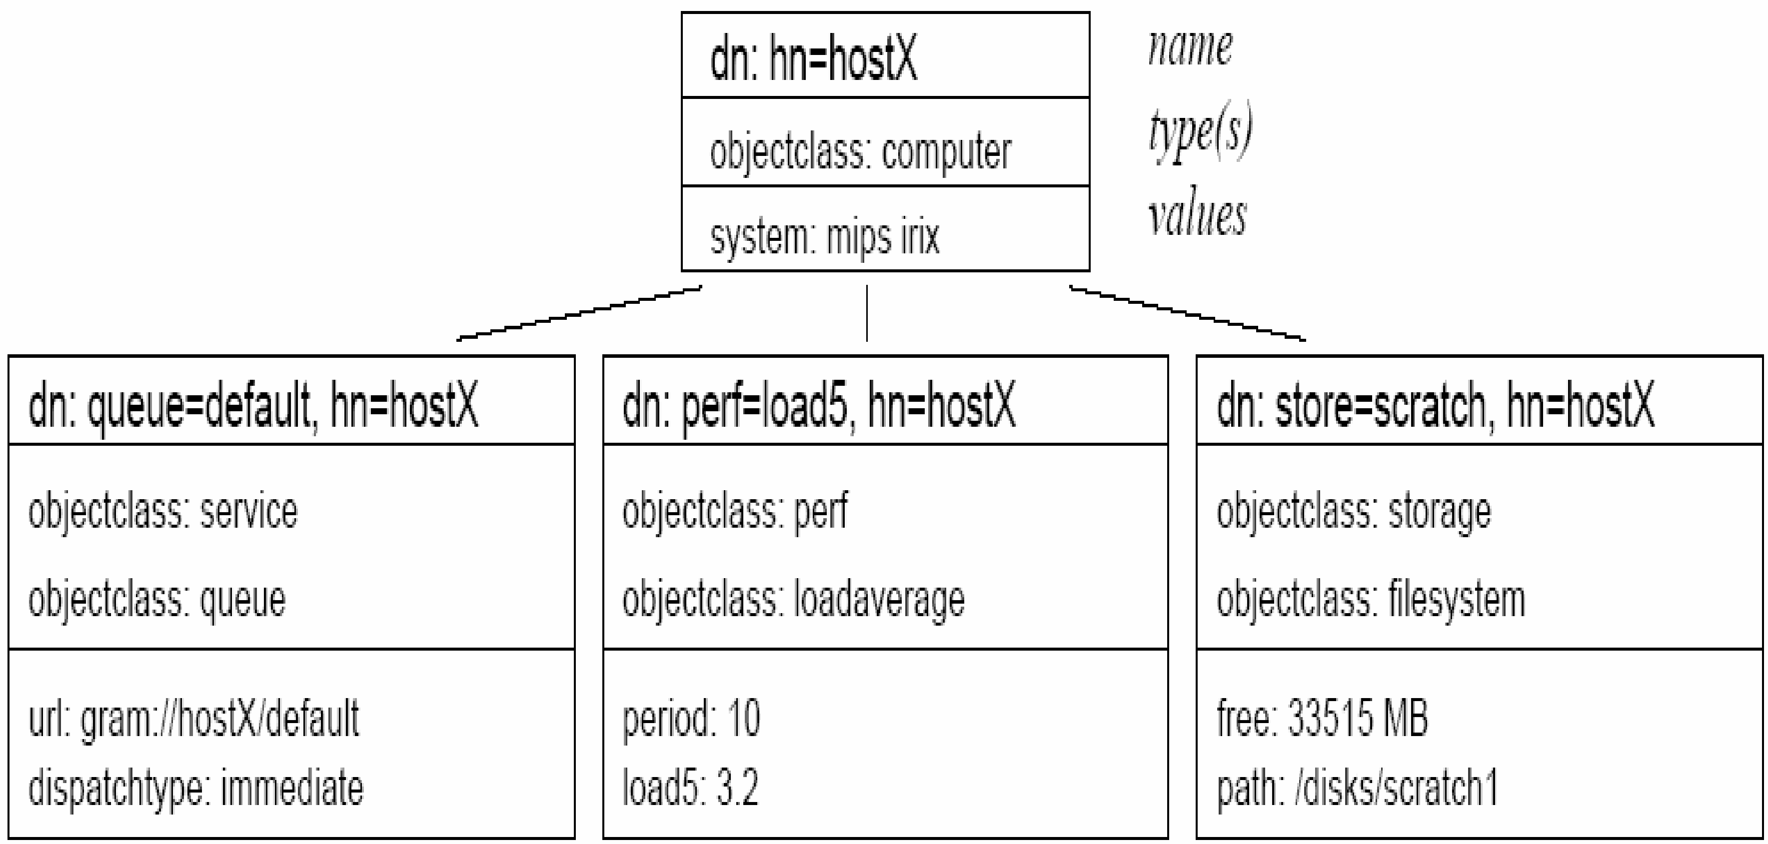

GRRP (Grid Registration protocol)

- Defines a notification mechanism that one service component can use to push simple information about its existence.

- For example, GRRP is used by an information provider to notify a aggregate directory of its availability for indexing, or by an aggregate directory to invite an information provider to join a VO.

- Limitations – Sometimes registration messages can get lost, which leads to incorrect removal of the Information Provider from the index of the Aggregate Directories.

A Decentralized Approach

- For large scale grids, centralized approaches may not scale well and prove to be performance bottlenecks even with replicated servers.

- Distributed approaches can remedy this problem.

- Use of an Overlay network has been proved to be quite successful in terms of resource discovery and sharing at global scales.

- Overlay networks implicitly balance loads, scale well to very large numbers , and can also meet the failure property of distributed operations.

A Decentralized Approach (Contd.)

- Build a distributed index of resources at each node.
- The search space is partitioned into various key spaces and each key space is assigned to a peer.
- Each peer builds a routing table which enables it to forward queries which it cannot answer locally.
- The common goal is to minimize the number of forwarding steps required to for successfully resolving queries in the presence of arbitrary peer and connection failures.

A Decentralized Approach (Contd.)

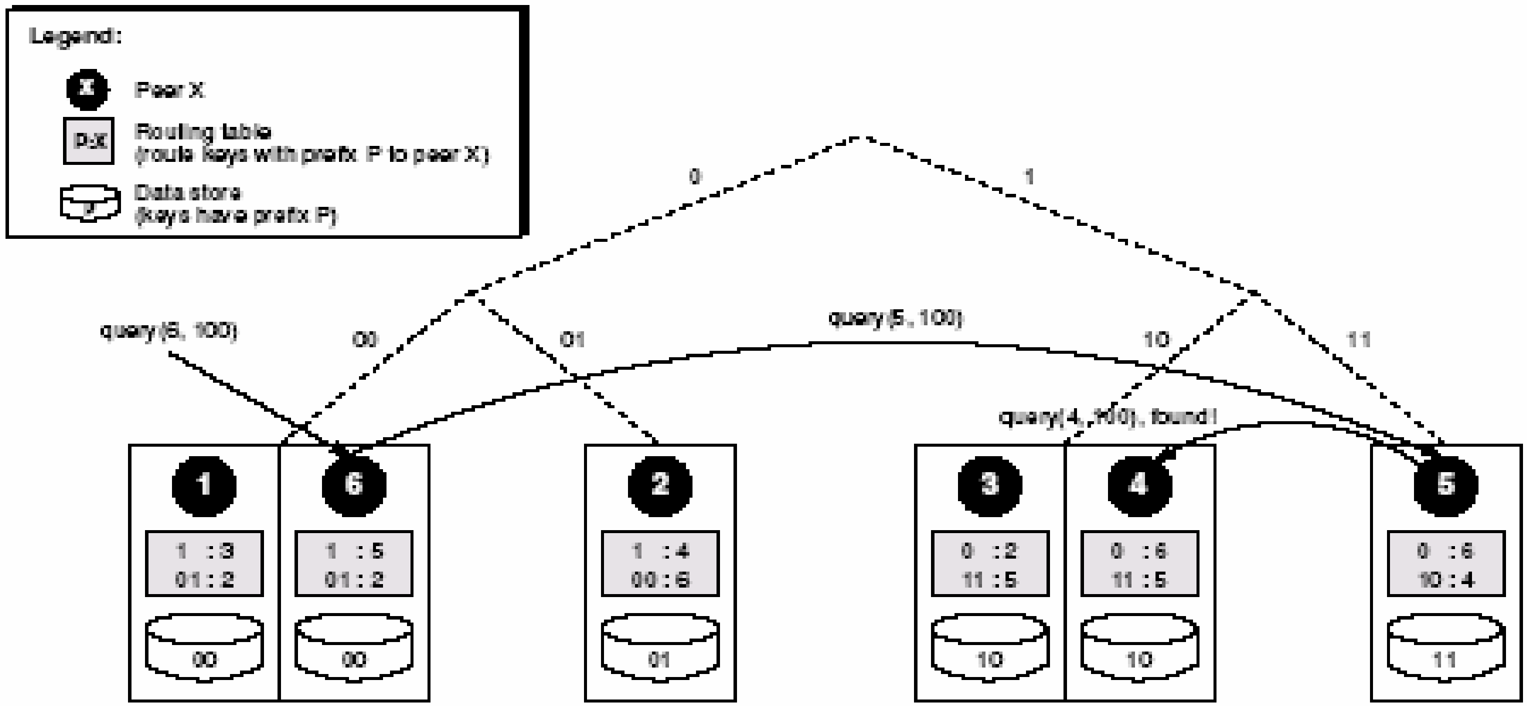

Figure 1. P-Grid overlay network

A Hierarchical Approach

- Unstructured Overlay networks often use broadcast and flooding mechanisms which are undesirable.
- Control resource heterogeneity to the extent that all the nodes in a group have some statistically-similar resource characteristic during a certain time period.
- Capture resource dynamics through grouping around dynamically elected leaders as well as publishing worker resource information.
- Enable the assemblage of large number of resources for applications within and across VO boundaries through a limited number of dynamic groups that are potentially searched faster than directly searching them individually.

Components

- Publish resource information;
- Handle group forming and maintenance;
- Receive and process query messages;

How does it work ?

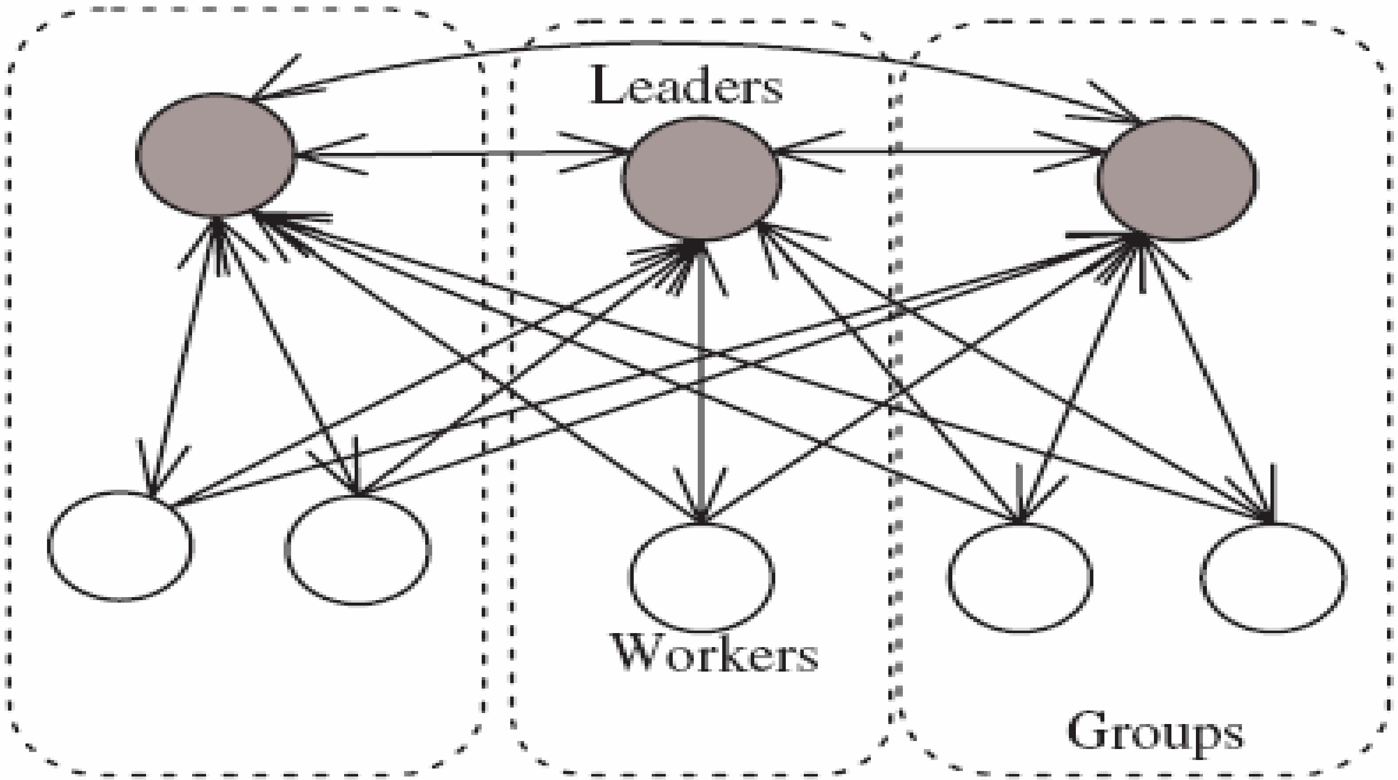

Fig. 1. The upper level graph maintained by SOG

A Semantic Approach

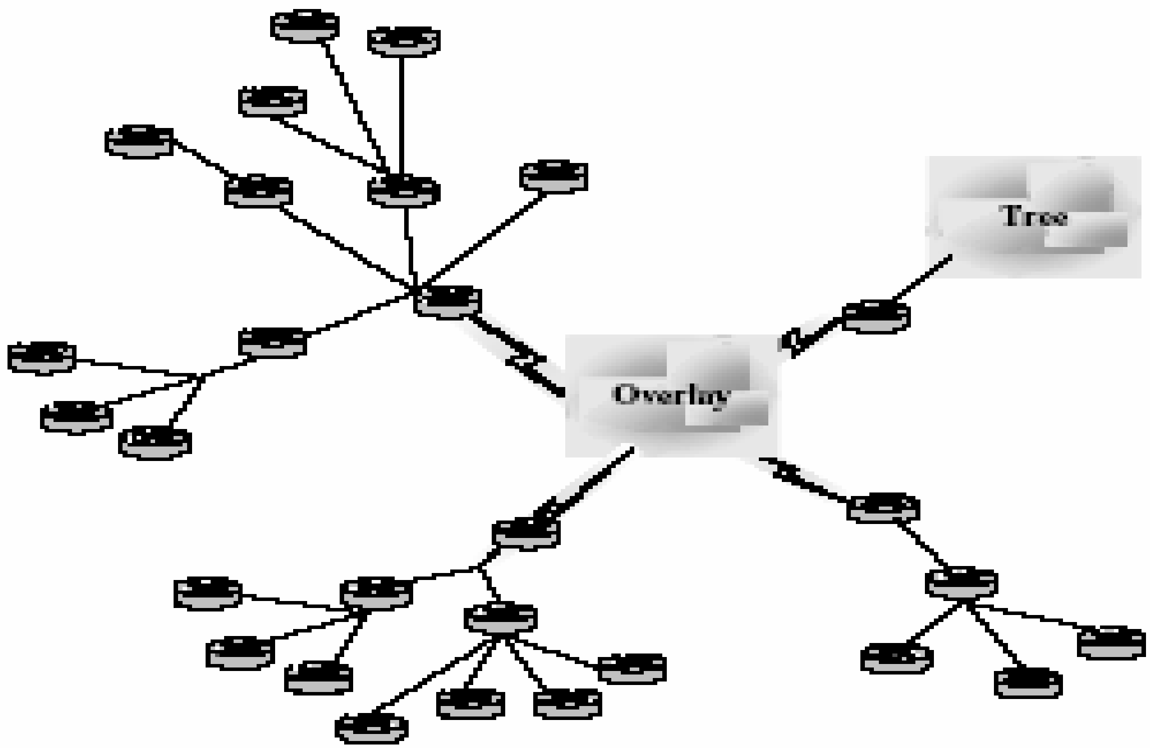

Figure 2. System architecture

Issues with some Approaches

- Centralized – Not globally scalable.
- Distributed – unstructured and use undesirable mechanisms like flooding and broadcasting.
- Hierarchical / semantic – Promising but lacks certain features.
- Lack of formalization and clustering ideas.
- Handing of complex, multi-attribute queries.
- Query expansion and relevance feedback may be useful !

Scope of future work

- Possibility of use of machine learning algorithms for unsupervised learning.
- A vector space model representation of resources and queries may scale very well for computing similarity measures.
- Query expansion and relevance feedback can exploit the inherent association of data.

A Sample Mixture Model

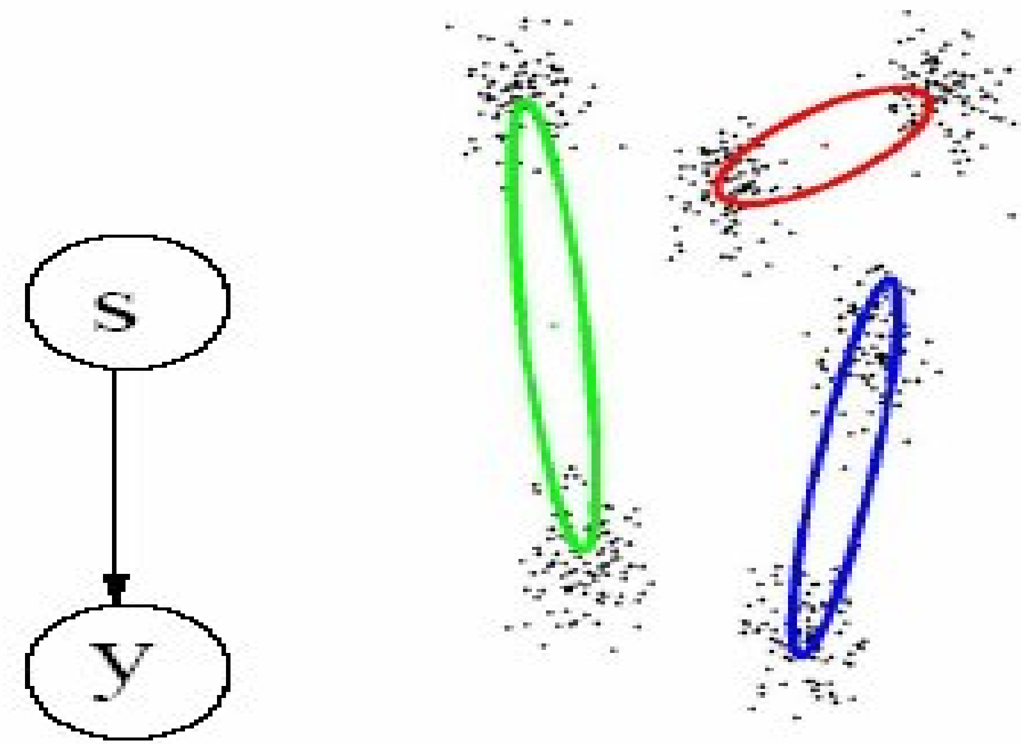

$$P(S^{(c)} = i|\pi) = \pi_i$$
$$p(y^{(c)}|S^{(c)} = i, \mu, \Sigma) = \mathcal{N}(\mu_i, \Sigma_i)$$

- Taken from Dr. Matthew J Beal’s lecture notes

VS Model: illustration

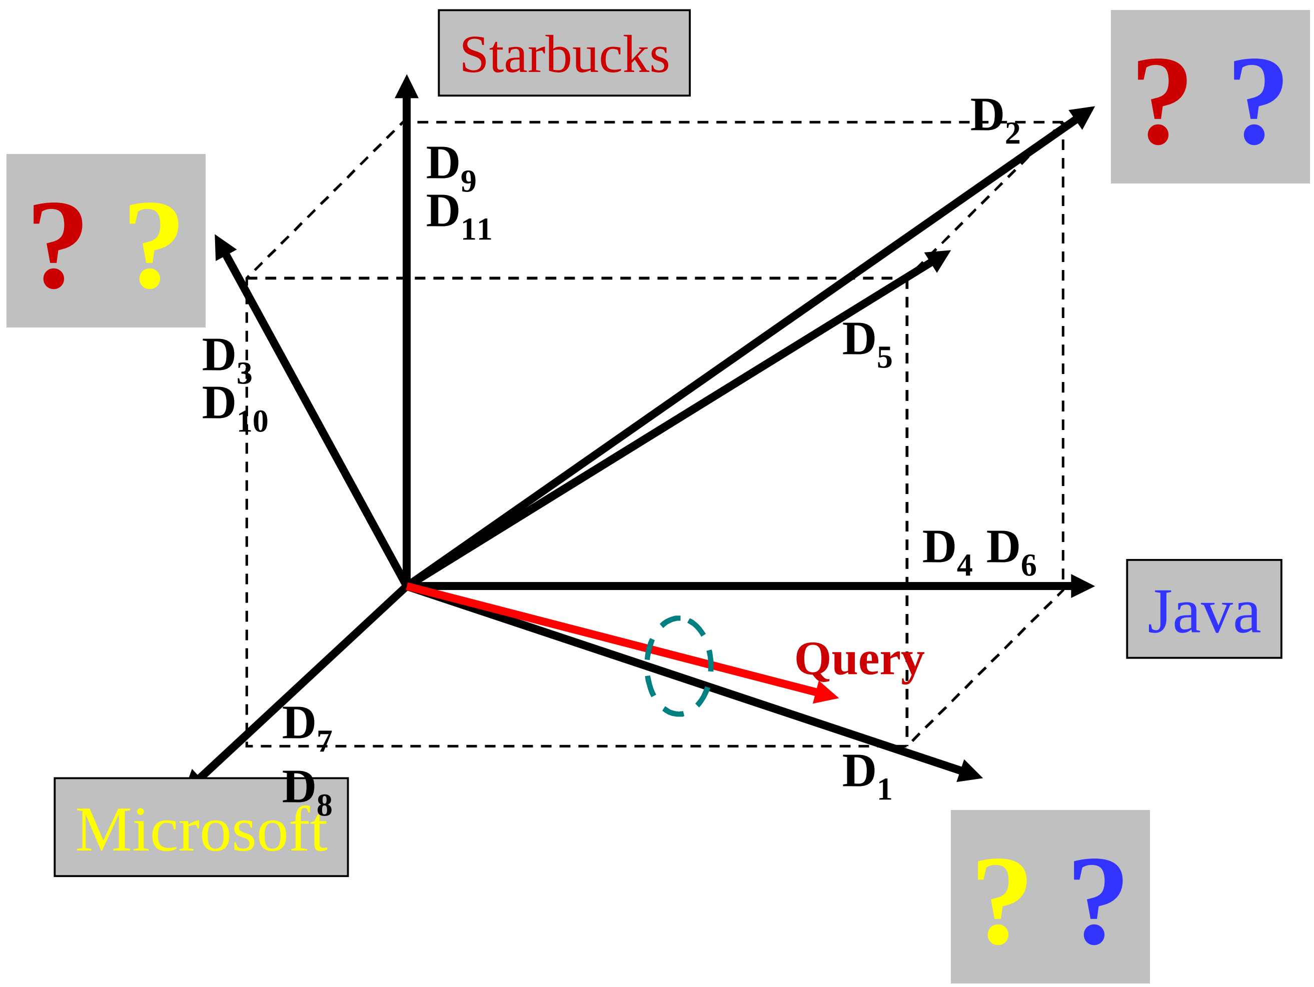

Query Re-formulation

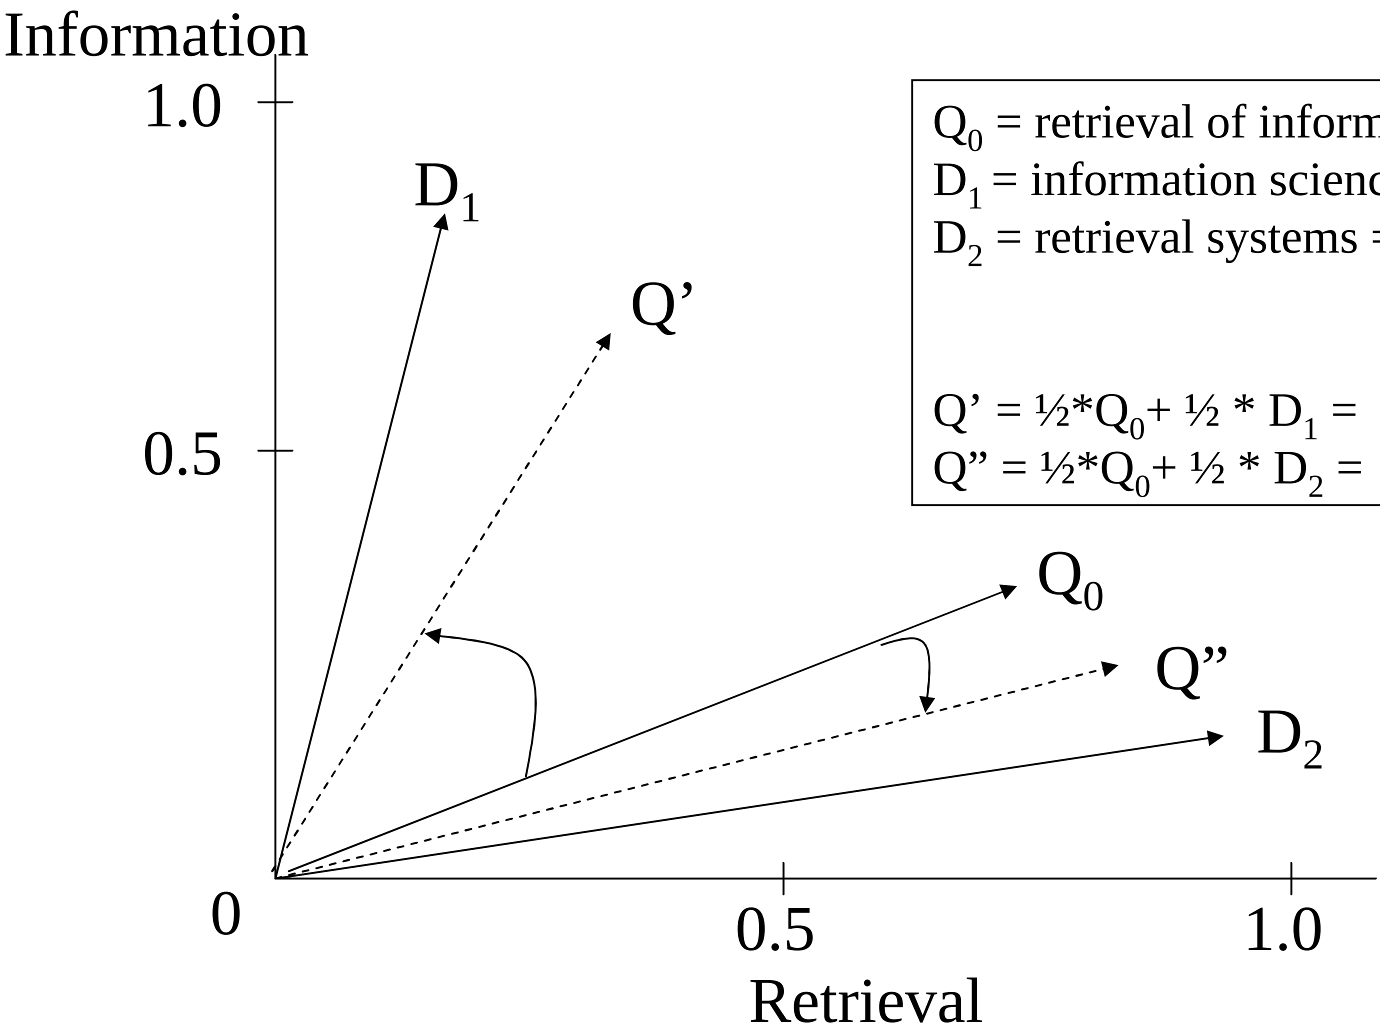

Q_0 = retrieval of information = (0.7,0.3)
D_1 = information science = (0.2,0.8)
D_2 = retrieval systems = (0.9,0.1)

Q’ = ½*Q_0+ ½ * D_1 = (0.45,0.55)
Q” = ½*Q_0+ ½ * D_2 = (0.80,0.20)

- Taken from Dr. Rohini Srihari’s Lecture notes

Conclusion

- Resource discovery can fundamentally be seen as a classic IR problem.
- Challenges : Heterogeneity and dynamic nature of data.
- Clustering ideas derive their origin from unsupervised learning in ML.
- Overall, a hybrid approach with some possible future work looks to be IN !!!

Q & A ??