



Simon Dobson and Paddy Nixon

Systems Research Group, School of
Computer Science and Informatics,
UCD Dublin, Dublin IE

<http://www.ucd.ie/csi>

simon.dobson@ucd.ie

paddy.nixon@ucd.ie

Adaptive middleware

Copyright © 2005, UCD Dublin

Why are we here?

The software environment is changing

- Component-based systems without centralised authorship
- Dynamic discovery and composition of components
- Mobility, ubiquity, autonomicity, self-* properties

The traditional approaches that served us well in the past need to evolve to meet these new challenges

- Middleware abstracts much of the complexity of interconnection
- ...but traditional middleware doesn't handle highly dynamic environments as well as we need it to

Our purpose in this tutorial

- To survey the current and emerging systems from the perspective of their support for developing adaptive systems
- To try to pick out some emerging trends

An executive summary

No current mainstream middleware provides good support for adaptive systems development

There are techniques that can be taken from a variety of systems and combined to construct adaptive applications

- Often mix paradigms

The research landscape is changing, and looks promising

- Several new systems that address some (but not all) issues

Who *are* these guys?

Academics in the Systems Research Group of UCD Dublin, Ireland

- Group focuses on pervasive and autonomic systems, system visualisation, semantics and software technology

Around 30 years' experience in designing and building large-scale distributed systems

- CORBA, event systems, programming languages, ...
- Ran a start-up company providing context-adaptive services to the travel industry – no, it's not still running... :-)
- Consulting for the likes of IBM, IONA, Siemens, ...

Structure

Adaptive 101

Middleware 101 (the 10 minute introduction)

A closer look at some systems

- Object-broker style systems
- Message- and event-oriented systems
- Tuple-space systems
- Peer-to-peer systems
- Knowledge-driven systems

What are the issues in supporting adaptive systems?

Some tentative conclusions

All references on website

Adaptive systems?

Aren't *all* systems adaptive to their inputs?

Well, maybe – but for the purposes of this tutorial

- A system whose *detailed responses and behaviour* change to deliver a service in the face of *changing context and constraints*

The system *changes* (at a system, configuration or communication level) in order to *stay the same* (at the application or user level)

- Somewhat paradoxically, an adaptive system aims to look *less* variable to its users than a “normal” system

Goals include predictability, reduced management complexity, better robustness, better tolerance of dynamic configuration


Places we encounter adaptation

Classic case is a location-based service

- Exactly what is served varies according to users' locations
- ...and possibly to the device they're using

Also appear in any sensorised environment, which these days includes networks and vehicles

- Adapt an entertainment system to the available content
- Adapt a network to the traffic load and the isochrony requirements of user tasks
- Adapt a business system to a changing population of end-point devices or applications



Different kinds of adaptation

Closed-adaptive systems

- Decide on the possible adaptations and how they will be selected
- Can be analysed to reduce impact on the user/programmer
- Typically will not be visible at all, e.g. TCP/IP congestion control

Open-adaptive systems

- Provide a framework for adaptivity
- Allow designers to download strategies as required
- Fewer guarantees possible, but probably better response to specific conditions
- Typically will result in at least some changes visible to users or programmers

Due to Rick Taylor and his colleagues

Starting points

The problems of complexity in current communication systems, even on the small scale, are identified by Bolosky:

"Users are subjected to random performance and service disruptions. Replacing or upgrading a personal computer, workstation, or server is very difficult. Even a moderate size computer network requires significant expertise to configure and maintain."

"However, over the next 15 years, we predict not just a quantitative expansion of computing, but qualitative change" –
Crowcroft, et al.

Adaptive systems compound these problems. They increase the level of dynamism, variability in infrastructure, and need for personalisation.

Eight fallacies of distributed computing

1. The network is reliable
2. Latency is zero
3. Bandwidth is infinite
4. The network is secure
5. Topology doesn't change
6. There is one administrator
7. Transport cost is zero
8. The network is homogeneous

Originally six, due to L. Peter Deutsch, updated with two additional ones (4 and 7) by James Gosling

The modern reality

1. Dynamism leads to network partitions and service loss
2. Communications time is unbounded
3. Bandwidth is generally small
4. Applications must deal with unknown parties
5. Topology and membership are unpredictable and variable
6. There is no (human) administrator
7. Wireless transport has a high power cost
8. Large degree of device heterogeneity

Our view of adaptation – 1

Two core features of the adaptive systems

- Physical integration – artefacts increasingly include information and communication
- Spontaneous interoperation – “encounter” services through mobility

From a systems perspective we would augment these two core observations with:

- Need to place local interactions in global setting
- Mobility of any part of the system is feasible (implied)
- Share data, services, ..., anything

Our view of adaptation – 2

We **cannot** assume everything we need is available

Nor can we assume that anything that is available **remains so**

We **cannot** assume everything that is available can be delivered when needed

We **cannot know** what is in the mind of the user

We can extract clues from the environment about **needs of the user**

We must use **available** resources to provide **best** services

We may use adaptation to change the costs of access of the availability or performance of resources

Many adaptations are provisional and error-prone, dependent on inference or invalidation over time

Our approach to, and resourcing of, service provision will change over time

Our view of adaptation – 3

To deal with adaptation these architectures and systems infrastructure must be able to

- Embrace contextual change
- Collect whatever meta-data is available
- Encourage *ad hoc* composition
- Facilitate sharing
- Support both local and the global computation
- Have multiple use view-points
(interactions designer/user/architect/programmer/system)
- Abstract away from the features that may adapt
- Allow well-founded adaptation to be specified

What we won't cover

Traditional middleware (other than in passing)

- A vague term that, when used in the context of Internet applications, means "software sold to people who don't know how to program by people who know how to program." [Philip Greenspun]

Shan't cover (but might touch on)

- Policy, security (in the broadest sense), quality of service, consistency, fault tolerance, low level detail (data formats etc.)

Will try to highlight – but avoid – the key overlaps with context awareness infrastructures

Will focus on interoperability, core principles, key paradigms, key challenges, what remains to be done

So what is middleware?

We used Google (like anybody else...) and:

- Middleware is commonly known as the **plumbing** of an information system as it routes data and information transparently between different back-end data sources and end-user applications.
- **The network-aware system software**, layered between an application, the operating system, and the network transport layers, whose purpose is to facilitate some aspect of cooperative processing.
- **Software that mediates** between an application program and a network. It manages the interaction between disparate applications across the heterogeneous computing platforms.

Or the favourite...

We like this one:

- “**Software that mediates** between different types of hardware and software on a network, so that they can function together. “


It places no constraints on what functioning together means

It captures (critically) heterogeneity of the software **and** hardware platforms

It highlights the importance of communication

It says nothing about transparency

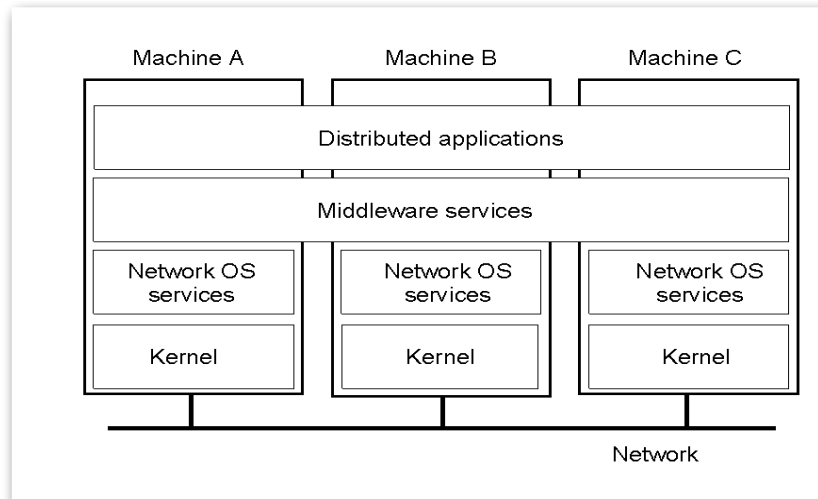
Distributed systems 101



The goal of location transparency has been assiduously pursued

- The web, CORBA, e-mail, ...
- Remove significance of – and usually any knowledge of – the (absolute or relative) locations of agents in a system
- Allow arbitrary interactions

Traditional middleware picture



In summary

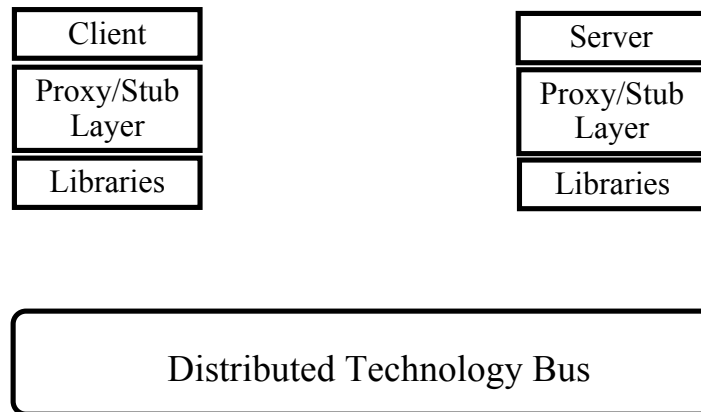
Middleware provides support for

- Naming, Location, Service discovery, Replication
- Protocol handling, Communication faults, QoS
- Synchronisation, Concurrency, Transactions, Storage
- Access control, Authentication


Middleware dimensions

- | | | |
|---------------------|-----|----------------------------|
| • Request/Reply | vs. | Asynchronous Messaging |
| • Language-specific | vs. | Language-independent |
| • Small-scale | vs. | Large-scale |
| • Tightly-coupled | vs. | Loosely-coupled components |

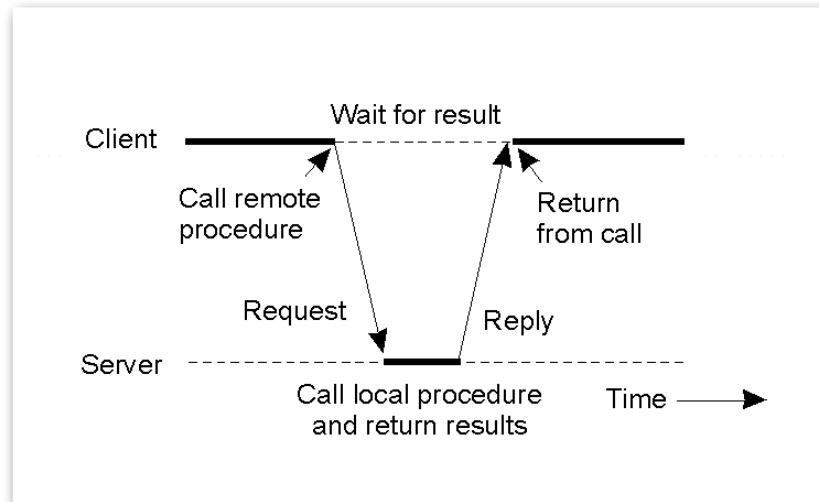
Traditional middleware picture



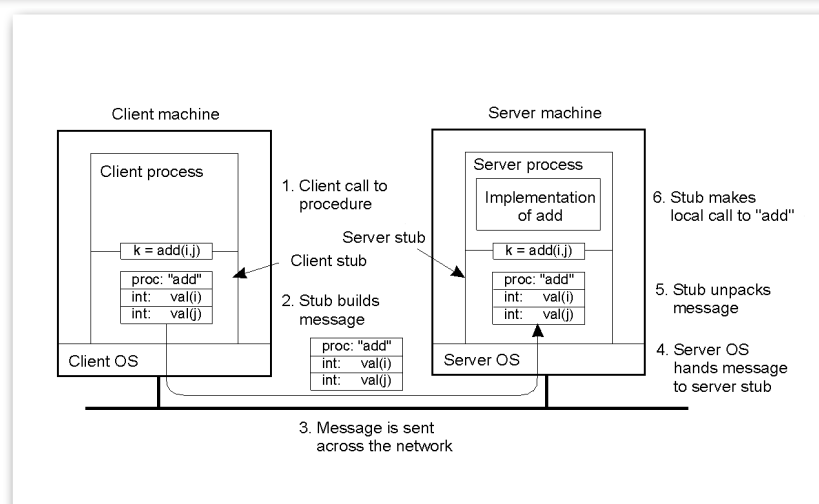
Traditional middleware picture




Basic synchronous invocation



Under the hood



CORBA




Common Object Request Broker Architecture

- Specification of a distributed middleware
- Specs drawn up by Object Management Group (OMG)
- <http://www.omg.org>

Goal: Interoperability with distributed applications on various platforms

CORBA overview




Object request broker (ORB)

- Core of the middleware platform
- Handles communication between objects and clients
- Handles distribution and heterogeneity issues
- May be implemented as libraries

Facilities: composition of CORBA services

Object model



Objects & services specified using an Interface Definition language (IDL)

- Used to specify interface of objects and/or services

ORB: run-time system that handles object-client communication

Dynamic invocation interface: allows object invocation at run-time

- Generic *invoke* operation: takes object reference as input
- Interface repository stores all interface definitions

CORBA services

Service	Description
Collection	Facilities for grouping objects into lists, queue, sets, etc.
Query	Facilities for querying collections of objects in a declarative manner
Concurrency	Facilities to allow concurrent access to shared objects
Transaction	Flat and nested transactions on method calls over multiple objects
Event	Facilities for asynchronous communication through events
Notification	Advanced facilities for event-based asynchronous communication
Externalization	Facilities for marshaling and unmarshaling of objects
Life cycle	Facilities for creation, deletion, copying, and moving of objects
Licensing	Facilities for attaching a license to an object
Naming	Facilities for systemwide name of objects
Property	Facilities for associating (attribute, value) pairs with objects
Trading	Facilities to publish and find the services on object has to offer
Persistence	Facilities for persistently storing objects
Relationship	Facilities for expressing relationships between objects
Security	Mechanisms for secure channels, authorization, and auditing
Time	Provides the current time within specified error margins

Object invocation models

Request type	Failure semantics	Description
Synchronous	At-most-once	Caller blocks until a response is returned or an exception is raised
One-way	Best effort delivery	Caller continues immediately without waiting for any response from the server
Deferred synchronous	At-most-once	Caller continues immediately and can later block until response is delivered

Invocation models supported in CORBA.

- Original model was RMI/RPC-like
- Current CORBA versions support additional semantics

Adapting object systems

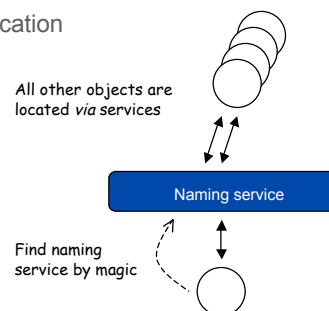
To interact with an object you need to know its identifier (IOR)

- By magic – hard-coded into the application
- Through some service(s)


Well-architected CORBA systems allow administrators to reconfigure applications *via* services

Still very much a manual task

- System can *be* adapted
- No infrastructure for it to *adapt itself*
- Deals well with slowly-changing systems with clear change boundaries



Event-based distributed programming




Event typically caused by state change in system

- Event changes state of object
- Multiple objects at different locations are informed of the occurrence of an **event** at a particular object, for example:
 - in a spontaneous computing environment, that a person's PDA has entered a hotel room
 - a client has entered participation in a collaborative work environment
 - an electronic document has been modified

Publish-and-subscribe events

Publish-and-subscribe (pub-sub) paradigm

- Object generating events **publishes** (producer) list of events for which other objects can receive notifications
- Object requiring notifications **subscribes** (consumer) to the notification service at an object offering notification for this particular event through its publication list
- Control how information propagates by controlling who registers for each service
- Typically have several different models for propagating events, with different performance characteristics



Notification – 1

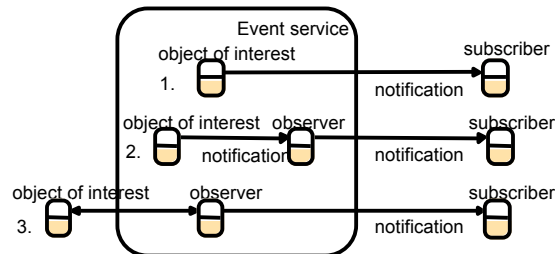
Architecture

- Event service

- maintains list of published events and registered subscriptions
- event service is notified of events at objects of interest
- subscribers subscribe for notifications at event service
- after occurrence of event, notification is sent to all subscribers

- Delivery semantics are implementation dependent

- e.g., notifications by IP multicast do not guarantee that any notification will ever get delivered to subscriber



Notification – 2

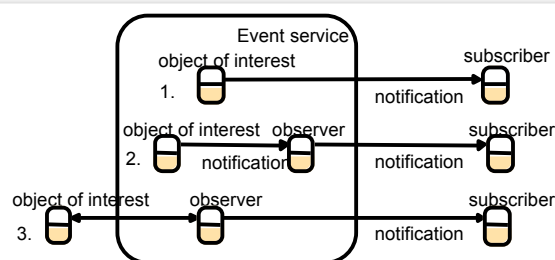
Architecture

- Observer objects

- decoupling of an object of interest from subscriber
 - responsible for all subscribers to events in some object
 - subscribers and the types of events they are interested in may be rather heterogeneous - hence, better to have observer deal with this (separation of concerns)

- Three cases

- in case 3, the object of interest is outside event service, hence, the observer needs to poll the object of interest for event occurrences




Message-oriented middleware

Communication using **messages**

Messages stored in **message queues**

Optional **message server** decouples client and server

Various assumptions about **message content**



Properties of MOM

Asynchronous interaction

- Client and server are only **loosely coupled**
- Messages are queued
- Good for application integration

Support for **reliable** delivery service

- Keep queues in persistent storage

Processing of messages by intermediate message server

- Filtering, transforming, logging, ...
- Networks of message servers
- Typically manipulate messages on the fly before forwarding

Natural for database integration

IBM's MQ-Series products are the canonical example of OM

Disadvantages of MOM

Poor programming abstraction

- Rather low-level (cf. packets)
- Results in multi-threaded code
- Request/reply more difficult to achieve

Message formats unknown to middleware

- No type checking

Queue abstraction only gives one-to-one communication

- Limits scalability

Adaptive messages/events

Most modern infrastructures do a good job of scaling services over the internet

- Intermediate relay servers
- Control information propagation – although this typically has to be done by hand

Most systems leave the location and configuration of the queues/servers to the designer

- Easy to get fossilised into a particular configuration
- No good ways of adapting automatically – services are typically too big and heavyweight
- Will provide good performance *as long as* the performance need is correctly anticipated

A hybrid: Akamai

The classic need to adaptation is content location in the web

- One place – get “Slashdotted” if you become popular
- Replicate – people may not find the replicas, there may not be one near many of the users

Akamai is an example of an adaptive web server

- Place content on a network of servers, managed as a whole
- Server network re-distributes cache of content depending on the observed patterns of requests
- Web sites point to “gateway” server which redirects to the best replica

Messages (HTTP requests) handed-off within the server network

Tuple-space systems

“Distributed workspace” research by David Gelernter and colleagues at Yale

Combines message passing and shared memory paradigms

- Nodes write arbitrary tuples (heterogeneous-type vectors) to shared memory
- Nodes read matching tuples from shared memory

Exact matching is required for extraction

Lookup calls always block until matching tuple exists


- Yuck!...

No guarantees about ordering, delay

Core API:

- out():
 - Writes tuples to shared space
 - Example: `out("abc", 1.5, 12)`
 - Result: Insert ("abc", 1.5, 12) into space
- read():
 - Retrieves tuple copy matching arg list (blocking)
 - Example: `read("abc", ? A, ? B)`
 - Result: Finds ("abc", 1.5, 12) and sets local variables `A = 1.5, B = 12`
 - Tuple ("abc", 1.5, 12) is still in the space.
- in():
 - Retrieves and deletes matching tuple from space (blocking)
 - Example: Same as above except ("abc", 1.5, 12) is deleted
- Eval()
 - Evaluates a tuple on the server

JavaSpaces – visual overview

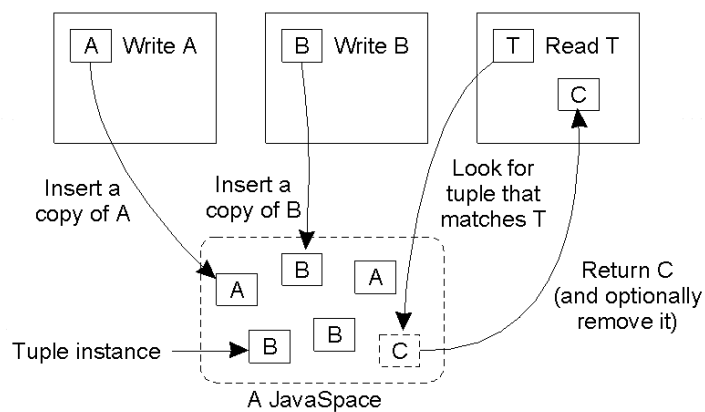


JavaSpaces overview


JavaSpace properties

- Store Java objects instead of tuples
- Spaces handle all details of sharing
- Objects are persistent (serializable) until removed or leases expire
- Object lookups are associative
- Transactionally secure (atomic)
- Objects may have executable content (e.g. methods)
- Security via identity servers

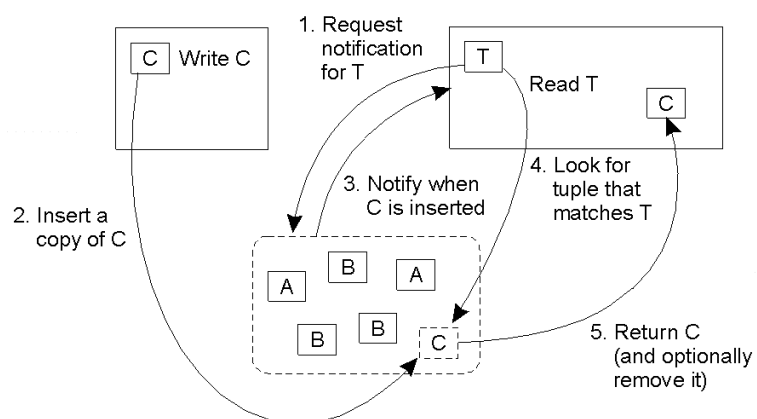
Overview of Jini



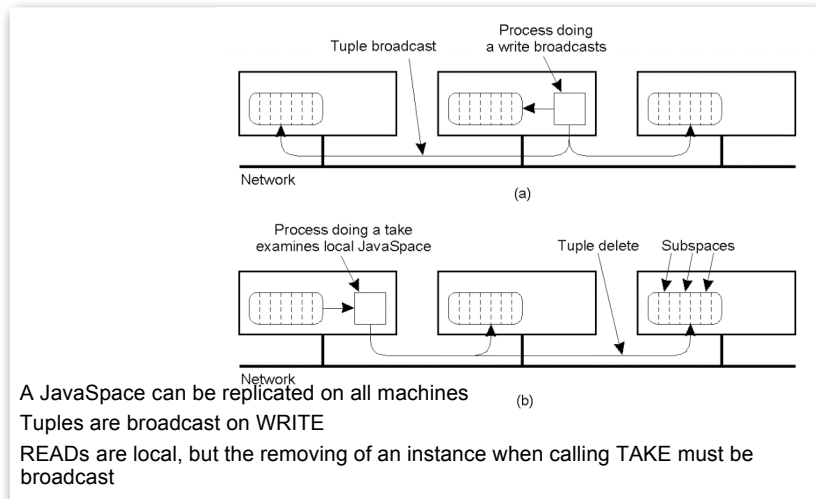
Architecture



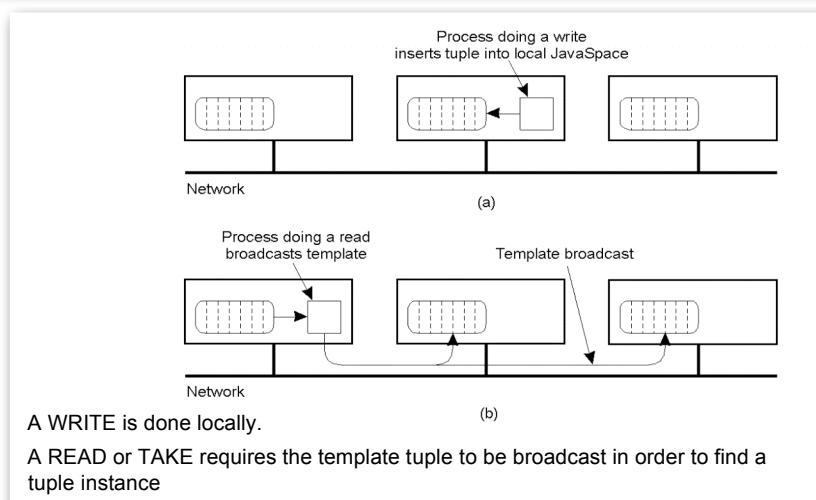
Communication events



Processes – replicated JavaSpace



Processes – unreplicated JavaSpace



The Jini lookup service

Field	Description
ServiceID	The identifier of the service associated with this item.
Service	A (possibly remote) reference to the object implementing the service.
AttributeSets	A set of tuples describing the service.

The Jini lookup service

Tuple Type	Attributes
ServiceInfo	Name, manufacturer, vendor, version, model, serial number
Location	Floor, room, building
Address	Street, organization, organizational unit, locality, state or province, postal code, country

Jini/CORBA don't hack it...

For adaptive and/or pervasive computing CORBA/Jini make bad assumptions

- Largely static and pre-configures services (naming, trading, etc..)
- A well-behaved computing environment
- Transparent and synchronous invocations
- No isolation between objects
- No independence between devices
- Distributed garbage collection

However, they provide a decent programming infrastructure in systems with limited dynamism, where services can be used to manage adaptation

Adaptive middleware 51

- Largely static and pre-configures services (naming, trading, etc..)
- A well-behaved computing environment
- Transparent and synchronous invocations
- No isolation between objects
- No independence between devices
- Distributed garbage collection

Adaptive middleware

iQueue

More of a framework than middleware

Aimed at providing data composition

input of each data composer is defined by an abstract data specification

The data resolver will receive periodic advertisement from available data sources.

Tries to bind source to data through a matching process.

Manages arrival and disappearance of data sources dynamically


```
graph TD
    CS[composer specifications] --> CM[Composer manager]
    CM -- credentials --> SM[Security manager]
    CM -- data specification --> BM[Binding manager]
    BM -- data-source handles --> CM
    BM -- port --> PM[Port manager]
    PM -- data-source descriptor --> BM
    PM -- credentials --> SM
    PM -- advertisements --> DR[Data resolver]
    DR -- data-source descriptors --> BM
    DR --> PPM[peer port managers in a distributed iQueue network]
    PPM --> PM
    PPM --> PTPM[peer topology managers in a distributed iQueue network]
    PTPM --> TM[Topology manager]
    PTPM --> CM
```

Figure 2. Internal components and data flows of iQueue. Broken lines represent components of a distributed iQueue system.

Adaptive middleware

52

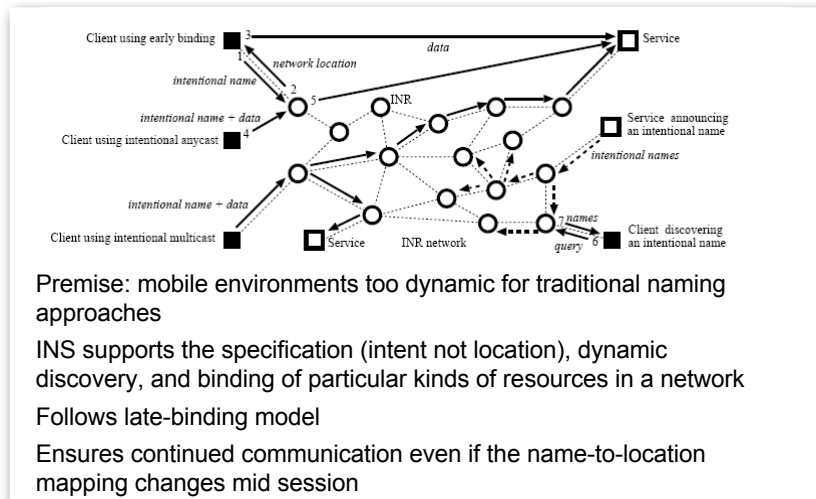
Manages arrival and disappearance of data sources dynamically



Adaptive middleware

52


Intentional Naming System



iROS

iROS is a middleware platform for a specific class of ubiquitous computing environment: interactive workspaces

- Tuple space co-ordination
- DataHeap provides type- and location-independent storage
- iCrafter service framework for user control of resources




EQUIP

Another tuple based system

Unique features is its integration of general event systems and tuple (shared data service)

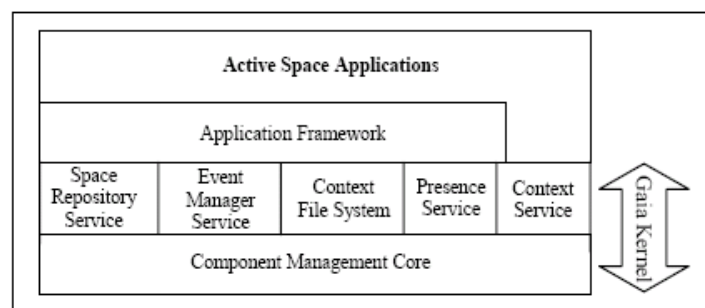
- Support replication of tuple spaces
- Another unique element is support patterns



GAIA – 1

GAIA – A middleware for active spaces

This system deserves a closer look as it represents the closest thing to an exemplar middleware infrastructure for adaptive and pervasive computing



GAIA – 2

Space repository

- An evolved Trader mechanism.
- Allows applications to query space for resources by function/attribute

Event Manager

- Adopts principle that ubiquitous systems are loosely coupled and events are the appropriate model for communication.
- Events are managed via channels which are many-to-many mappings of sinks to sources.
- Channels are generated by factories based on templates/properties.

Context Service

- Provide a way for applications to query for or register interest in certain contextual elements (sensors for instance).
- Context modelled as 4-tuple and mapped to event channel.
- In some sense a distributed and evolved version of Context Toolkit.

GAIA – 3

Presence Service


- Maintains information about digital and physical entities currently available in a given space.
- Uses heart beat beaconing to maintain current view
- After each heartbeat it informs appropriate services/applications of changes

Context File Service

- Stores files tagged with context as detected by space
- For example, to determine which files have the context of **location == RM2401 && situation == meeting**
- associated with them, one may enter the directory.
/location:/RM2401/situation:/meeting

Application Framework

- Active spaces entail a user-centric, resource-aware, multi-device, context-sensitive, mobile application model.
- Extends Model–View–Controller [25] and introduces new functionality to export and manipulate the bindings of the application components;
- Policies customize several aspects of applications including instantiation, mobility, reliability, and number and composition of components and their bindings)



Peer-to-peer

In a *really* dynamic environment you may not have *any* infrastructure

- Must provide all services using the nodes themselves
- Potentially extremely adaptive, but fewer guarantees on service

Peer-to-peer (P2P) systems

- All nodes are equal, providing (at least) routing and (possibly) other services

Two approaches

- Provide P2P “native” in the MAC layer
- Provide a P2P overlay on top of a standard transport

Allows more scope for optimisation and features specific to P2P

The more popular approach at the moment, piggy-backing onto TCP/IP over WiFi or wired networks

General requirements

Node discovery and management

- Locate a new node and integrate it into the network
- Manage a node leaving the network, possibly without notice
- Discover the services available within the network

Security and trust

- How can you trust a node you only just met?
- Need trust along the entire path, not just the end-points
- Only limited scope for encryption

Routing

- Need to compute routes between nodes that we may not know about, in the presence of frequent failure as the topology changes


Pastry – 1

Developed by Microsoft, Pastry provides a routing and service discovery model overlayed onto TCP/IP

- Each node has a 128-bit id
- Nodes form a ring
- Each node retains a leaf set consisting of n nodes along the ring from itself

Node ids are assigned by a trusted source, with certificates

- Node needs a node id, certificate and the address of another node, and can then join the ring



Pastry – 2

To route a message a node does one of the following

- Forward message to a node whose id shares a longer prefix than the current node
- If no such node known, forward to a node with a prefix at least as long as the current node but numerically closer

Messages store the node ids and IP addresses of the nodes they have visited

- Intermediate nodes can update their routing tables, discovering any nodes that are “nearby” in node id space

Sun's JXTA product takes a similar approach

Content in P2P networks

P2P networks are often used for content storage and retrieval

- Each node stores some fraction of the content files but can access the entire network's content

...and there are plenty of legal uses for this technology, so P2P content networks are not *all* used for copyright theft...

For example, Pastry provides a file store based on the hashes of files

- A distributed hash table – convert a file's hash to its node id, and then route a request to it
- Inserting content routes to the node id closest to the “ideal”, and then replicates the file over nearby nodes
- Can be more or less aggressive about caching and re-distribution

If the closest node disappears, requests will route to a nearby node and so still find the content unless *all* the nearby nodes have left

Construct

Construct is UCD's contextual systems infrastructure

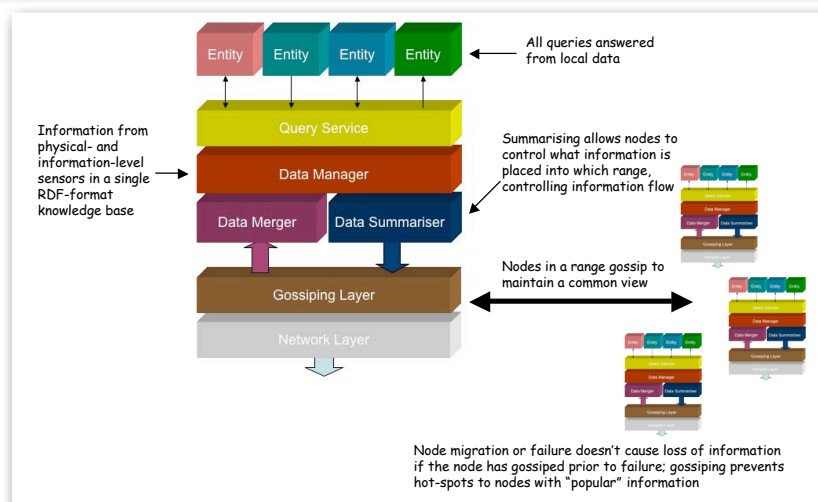
- Provide simple, open, scalable, robust, knowledge-based management of information in a dynamic network
- Being used for pervasive systems, smart spaces, autonomic management of network traffic, ...

Leverage semantic web technology

- All information represented using RDF, easily queried and shared
- Ontologies provide structure
- Component-based, standard underlying protocols

A peer-to-peer knowledge-sharing system

Construct architecture



Gossiping – 1

Most systems store information either according to its production or according to its consumption

- Production: sensor/server stores information, anyone who wants it asks for it (typical in object systems)
- Consumption: store information where it's likely to be used or will be easy to locate

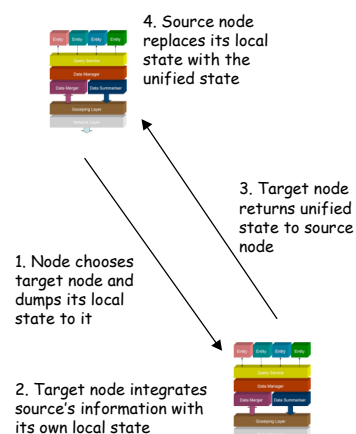
A third approach is to gossip

- Each node periodically chooses another node in the network at random and synchronises with it
- Information propagates through the network as if over a random graph – an overlay on an overlay...
- Effective way of trading space for speed and robustness

Gossiping – 2

Characteristics

- No communication or storage hot-spots – overlay behaves like a random network
- Nodes need lots of storage, as they're caching a complete system summary (although perhaps not all the details)
- Increased robustness – a node that leaves doesn't take its information with it
- Local queries only – any information is available locally
- Information aging is a critical issue, may need to use older information than is strictly available in the network



Web services

No tutorial would be complete without them...

Depending on your point of view, they either

- Re-package existing technologies in a sexy new XML wrapping
- Provide a completely different way of dealing with information

The reality is probably half-way between

- Typically *used* as an object system, so have the usual disadvantages in terms of location of services
- Typically *architected* using messages, allowing requests to be re-directed and processed on the way from client to server, as well as increasing asynchrony (at the expense of programmer headaches)
- Change the viewpoint from which we develop services: focus on exposing business tasks

Do web services affect adaptation?

At present web services are typically almost indistinguishable (from a systems perspective) to CORBA systems

No real excuse for not exposing *any* distributed system as a web service

In the future we may see more innovative uses

- Message re-direction, on-the-fly processing, ...
- Collection of meta-data on requests to inform changes

These things are possible in CORBA if you're a serious wizard, but don't appear very often in mainstream applications

Main impact seems to be in ease of integration

- Still need to decide *how* to adapt, but the actual process may well be simpler, especially across enterprise boundaries

Issues for middleware systems

Description

- It should be possible to describe components and their interactions in a way that explicitly prescribes their abstract roles in a system

A core challenge relates to how we build, share and use understandable descriptions

We need to describe not only the information but also the system and its configuration.

We need to be able to reason about the semantics of these descriptions (cf. semantic web technologies)

Issues for middleware systems

Composition

- It should be possible to describe a system as a composition of independent components and connections

Current work on composition talks of composition rules, policies, aspects, etc. With a focus on composition from known sets.

- Adaptive systems will additionally have a variety of composition rules (for the user, applications developer, system, hardware)

There should be some structure for relating between the levels.

What are the semantics of these rules – how do we describe a closure so that we can reuse/decompose the composition?

Issues for middleware systems

Dynamic composition

- It should be possible to reuse components, connectors, and architectural patterns even if they've been developed for another purpose!

We need to be able to describe families of systems, their semantics and constraints from open sets

- Open-adaptive *versus* closed-adaptive

Typically, existing composition approaches use closed or parameterised sets.

How do we support dynamic composition and still maintain a robust, predictable system?

Issues for middleware systems

Trade-offs

- We must be willing to sacrifice the notion of optimising *everything*, making *everything* efficient
- Know what we need to optimise against, and trade-off everything else

Dynamic node populations imply either accepting information loss or taking the cost of replication

Using overlays can gain robustness from random distribution but lose performance with longer routes

At the current state of the art we can probably only handle adapting against one or two criteria

Issues for middleware systems

Interference

- We should be able to deal with changing and conflicting resource requirements in the environment

JINI, and others, adopt a notion of *leasing resources*.

However, they give no solutions to the *free market economy* of the adaptive systems world

- Some resources must be held for a complete transaction
- May not be replacements

Even the simple case pretty much is intractable

Issues for middleware systems

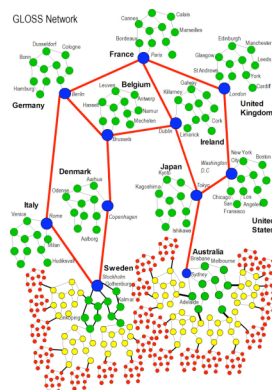
Global *versus* local

- We should be able to utilise the appropriate set of local and global resources to achieve the task

Adaptive systems are not just local interactions.

Ability to move between local environments and retain context

So, where is the information and computation placed?



Issues for middleware systems

Control

- We should be able to describe a how the system is controlled with respect to a variety of changing parameters and from a variety of view-points

Locating services and resources

Manage resources from the perspective of the group/region/domain/...

Coordinate the progress from a certain viewpoint

Be able to express partial requirements

Detect and recover from failure

Issues for middleware systems

Viewpoint

- We should be able characterise the usual interactions styles

Optimise for the usual interactions

However, this typically inhibits the ability of the system to adapt to the unusual.

Equally, the vocabulary for describing adaptive systems varies between domain

- Is it possible to have a common framework within which the semantics can be debated?

Issues for middleware systems

Context

- We should be able to contextualise interactions in order to adapt the infrastructure, information, or its delivery, to the semantics of use

Relates very much to viewpoint. How do we codify the behavioural characteristics of the user

What is the peripheral variable set for this user, doing this task, in this situation?

- What do we adapt *to*?

A core challenge in relating hardware sensed context with their semantics of use at that time

- Hard to adapt to something that hasn't happened – but when it *does* happen we may lose service, at least temporarily

Fundamental problems

Semantic multiplicity

- Lots of events are “the same” in the higher-level view. How do you find the exact service you want dynamically?

Security and privacy

- Adaptation must not reveal information that needs to be concealed

Sensorisation and sensor fidelity

- There's a temptation to collect too little information about on-going behaviour
- Anything dealing with physical phenomena is noisy, so (for example) people are identified incorrectly, or not observed

Latency

- By the time you work something out it may be too late to do anything about it

Placement of information and computation

- Where does the system do the computation, how does it get the correct information, and how does it achieve this?

What are the trends? – 1

Increasing use of peer-to-peer

- Infrastructural solutions are often too inflexible, non-scalable and expensive to deploy – how can you dimension them accurately?
- Shift the costs onto the users

Increasing need for end-to-end properties – especially confidentiality and information flow restriction

- Adaptation is information, like anything else...

Events scale well for systems, but not for programmers

- Hard to build well-integrated systems from raw events
- Middleware may use them internally, but needs another abstraction above

What are the trends? – 2

Focus on system properties *not* technology

- Take ideas from wherever they appear – you can do peer-to-peer in CORBA, with sufficient creativity
- Don't expect a packaged "adaptivity solution" that's suitable for your needs any time soon

Open-adaptivity is the only game in town

- Except for *really* closed systems – and there are increasingly few...

Conflicts are inevitable

- Adaptations **will** disagree
- How do we do decision-making in open environments?

What are the trends? – 3

Keep all expressions at the highest possible level

- Code should be the last level of expression
- Logic and rules can be analysed and evolved in a well-founded way

Robustness is not just a non-functional requirement

- Services need to remain available
- Peer-to-peer and gossiping build robustness into the core of the system

Collect context

- Almost *any* piece of (meta-)data can be used to inform a decision on adaptivity
- Instrument to make this possible going forward – the costs will massively outweigh the disadvantages

Concluding remarks

For us, adaptive systems represent the most significant challenge for middleware and systems research of recent years

- Unquestionably it requires a diversity of fundamental science and engineering research to realise.

Our hot topics:

- Light weight, real-time and adaptive (reflective) middleware
- Semantics and design
- New models of failure and recovery (Recovery Oriented Computing?)
- Update and versioning
- Quality – rephrased maybe to be quality of context and service
- Explicit management of uncertainty
- Trust – monitoring, security, privacy ...
- Sensor middleware

Thank you!

More information on this tutorial can be found
on the web by navigating from this course's
home page:

<http://www.simondobson.org/teaching/content/courses/adaptive-middleware.shtml>

More on our research activities
can be found on the SRG home
page:

<http://srg.cs.ucd.ie/>