Design Issues

- Name spaces
- Name Resolution
- The domain name system

Name Spaces

- A name space is a collection of all valid names recognized by a particular service
- Allow simple but meaningful names to be used
- Potentially infinite number of names
- Structured
  - to allow similar subnames without clashes
  - to group related names
- Allow re-structuring of name trees
  - for some types of change, old programs should continue to work
- Management of trust

Name Space

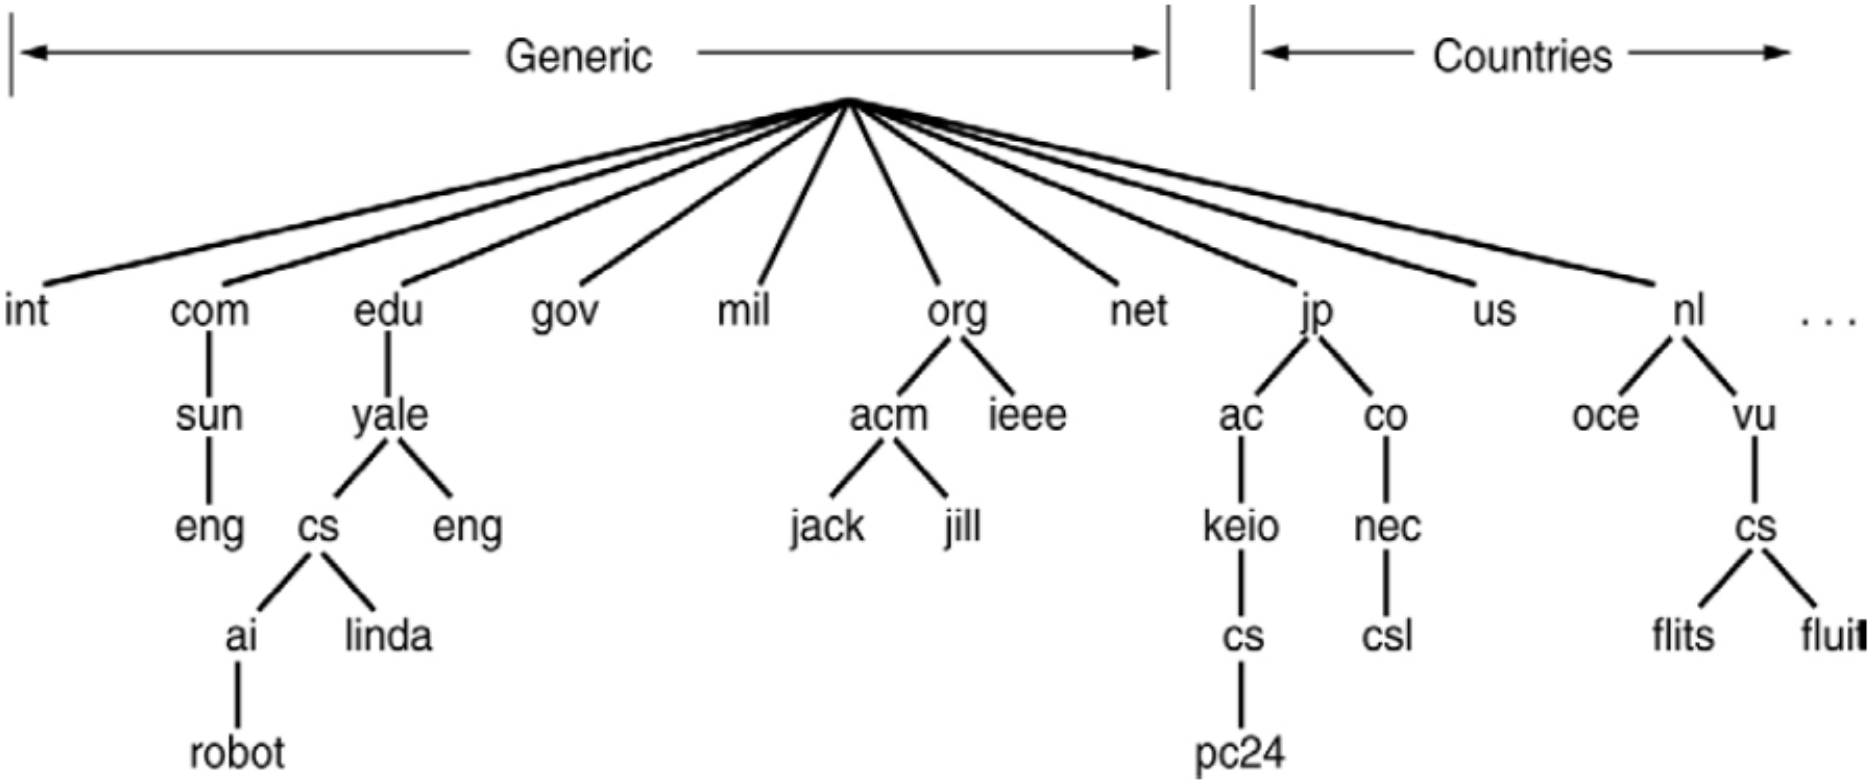
Generic
Countries
int
com
edu
gov
mil
org
net
jp
us
nl
. . .
sun
yale
acm
ieee
ac
co
oce
vu
eng
cs
eng
jack
jill
keio
nec
cs
ai
linda
cs
csl
flits
robot
pc24

Name Resolution

- Resolution is an iterative process whereby a name is repeatedly presented to the naming contexts.
- The name is first presented to some initial naming context; resolution iterates as long as further context and derived names are output.
- Example1: /etc/passwd in which ‘etc’ is presented to context / and ‘passwd’ is presented to context /etc.
- Example 2: www.dcs.qmw.ac.uk in which the alias is resolved to another domain name such as copper.dcs.qmw.ac.uk which is further resolved to produce IP address.

Name Servers and Navigation

- Any name service stores a very large database.
- Data is partitioned into servers according to its domain.
- Partitioning of the data implies that the local name server cannot answer all the enquiries without the help of other name servers.
- Process of locating naming data from among more than one name server in order to resolve a name is called navigation. Ex: Iterative Navigation model(DNS)

Iterative navigation

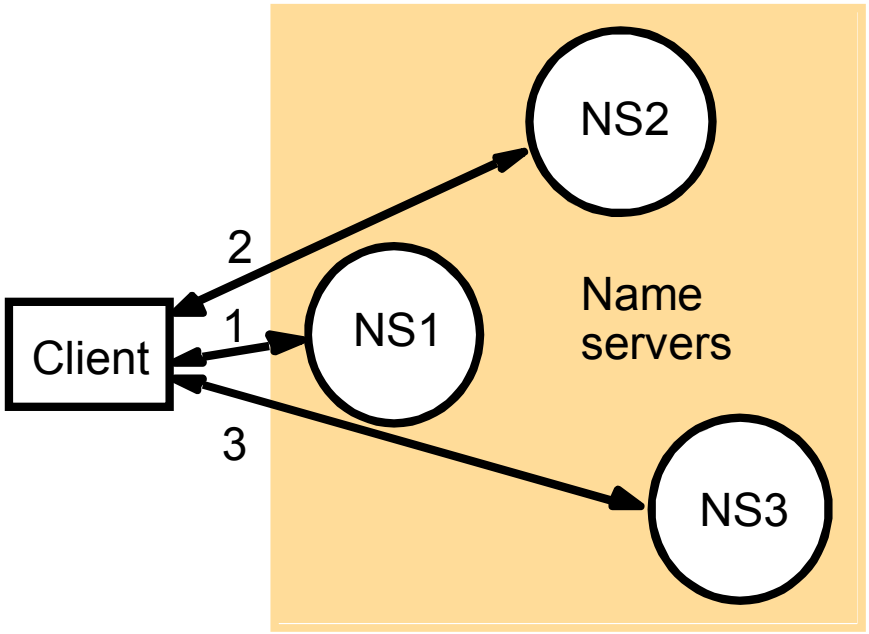

A client iteratively contacts name servers NS1–NS3 in order to resolve a name

Non-recursive and recursive server-controlled navigation

Figure 9.3

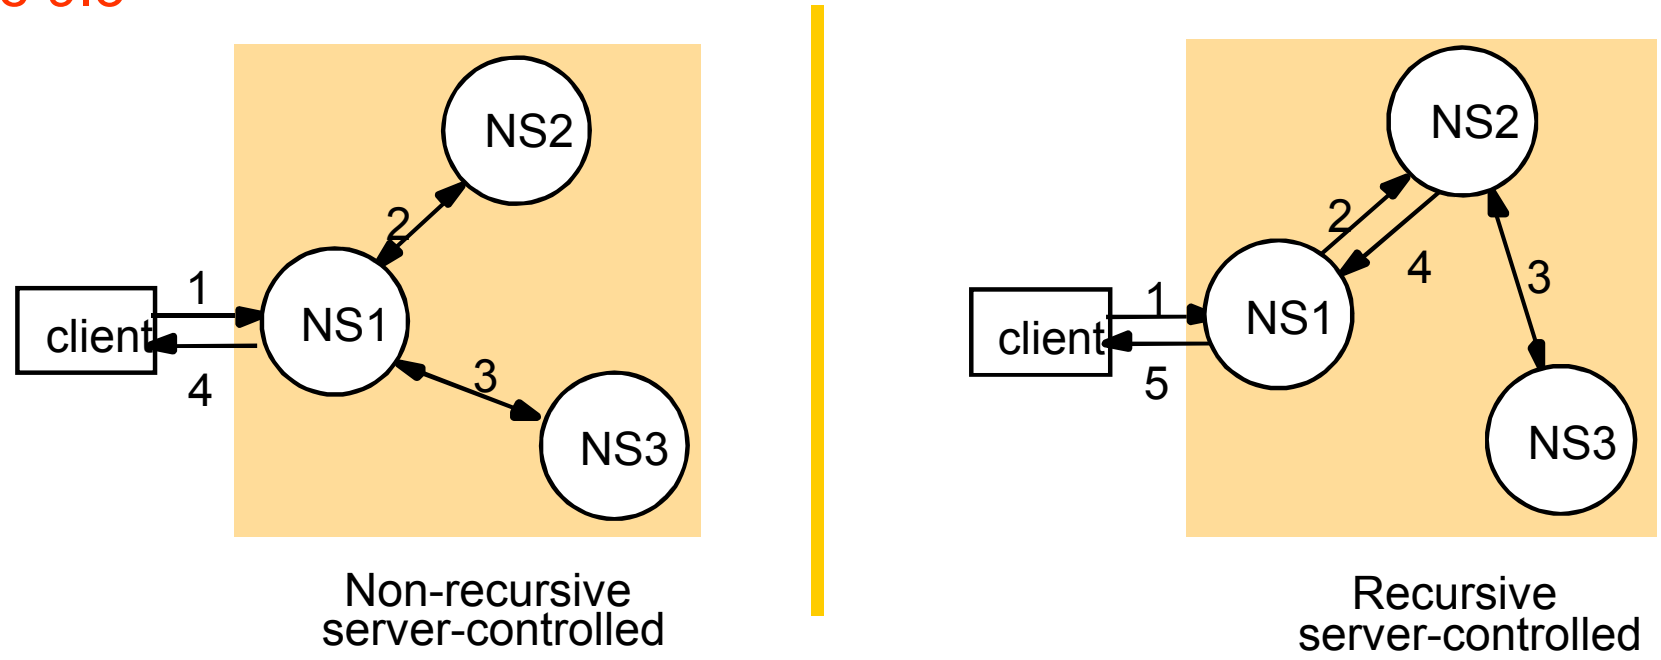

A name server NS1 communicates with other name servers on behalf of a client

DNS offers recursive navigation as an option, but iterative is the standard technique. Recursive navigation must be used in domains that limit client access to their DNS information for security reasons.

*

Caching

- Client name resolution software and servers maintain a cache of previous name resolutions.
- How long a resolver caches a DNS response (i.e. how long a DNS response remains *valid*) is determined by a value called the time to live.
- Server may use data from its own cache or other server cache it is authorized to access.
- Caching is key to performance and fault tolerance.

The Domain Name System

- A distributed naming database
- Name structure reflects administrative structure of the Internet
- Rapidly resolves domain names to IP addresses
  - exploits caching heavily
  - typical query time ~100 milliseconds
- Scales to millions of computers
  - partitioned database
  - caching
- Resilient to failure of a server
  - replication

Parts of a domain name

- Usually consists of two or more parts (technically *labels*), separated by dots. .
- The rightmost label conveys the top-level domain.
- Each label to the left specifies a subdivision, or subdomain of the domain above it.

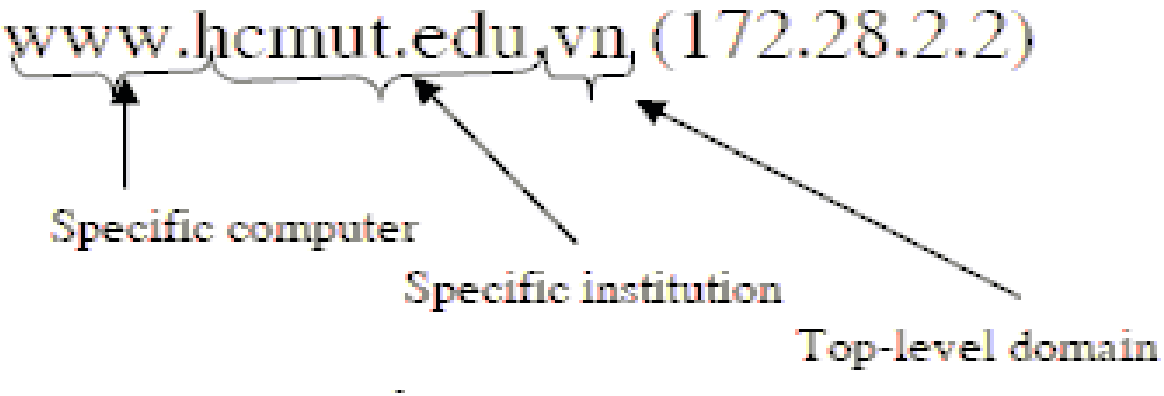

DNS name servers

- The Domain Name System consists of a hierarchical set of DNS servers
- Each domain or subdomain has one or more authoritative DNS servers that publish information about that domain and the name servers of any domains "beneath" it
- The hierarchy of authoritative DNS servers matches the hierarchy of domains.
- At the top of the hierarchy stand the root nameservers: the servers to query when looking up (*resolving*) a top-level domain name

Basic DNS algorithm for name resolution

- Look for the name in the local cache
- Try a superior DNS server, which responds with:
  - another recommended DNS server
  - the IP address (which may not be entirely up to date)

DNS name servers

Figure 9.4

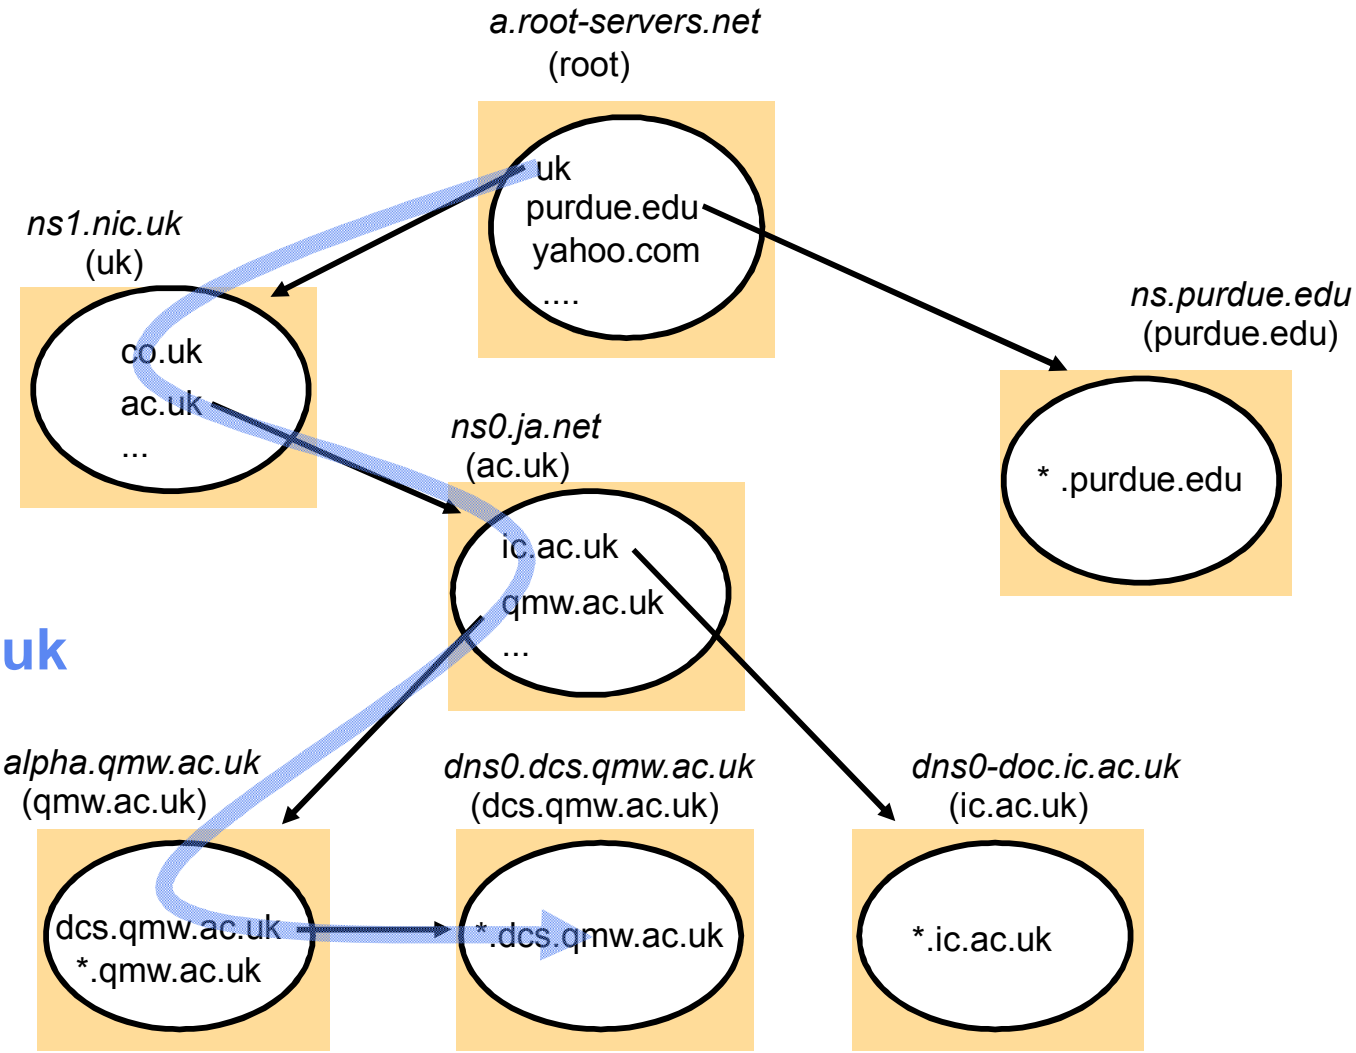

Note: Name server names are in italics, and the corresponding domains are in parentheses. Arrows denote name server entries

authoritative path to lookup:

jeans-pc.dcs.qmw.ac.uk

*

DNS in typical operation

Without caching

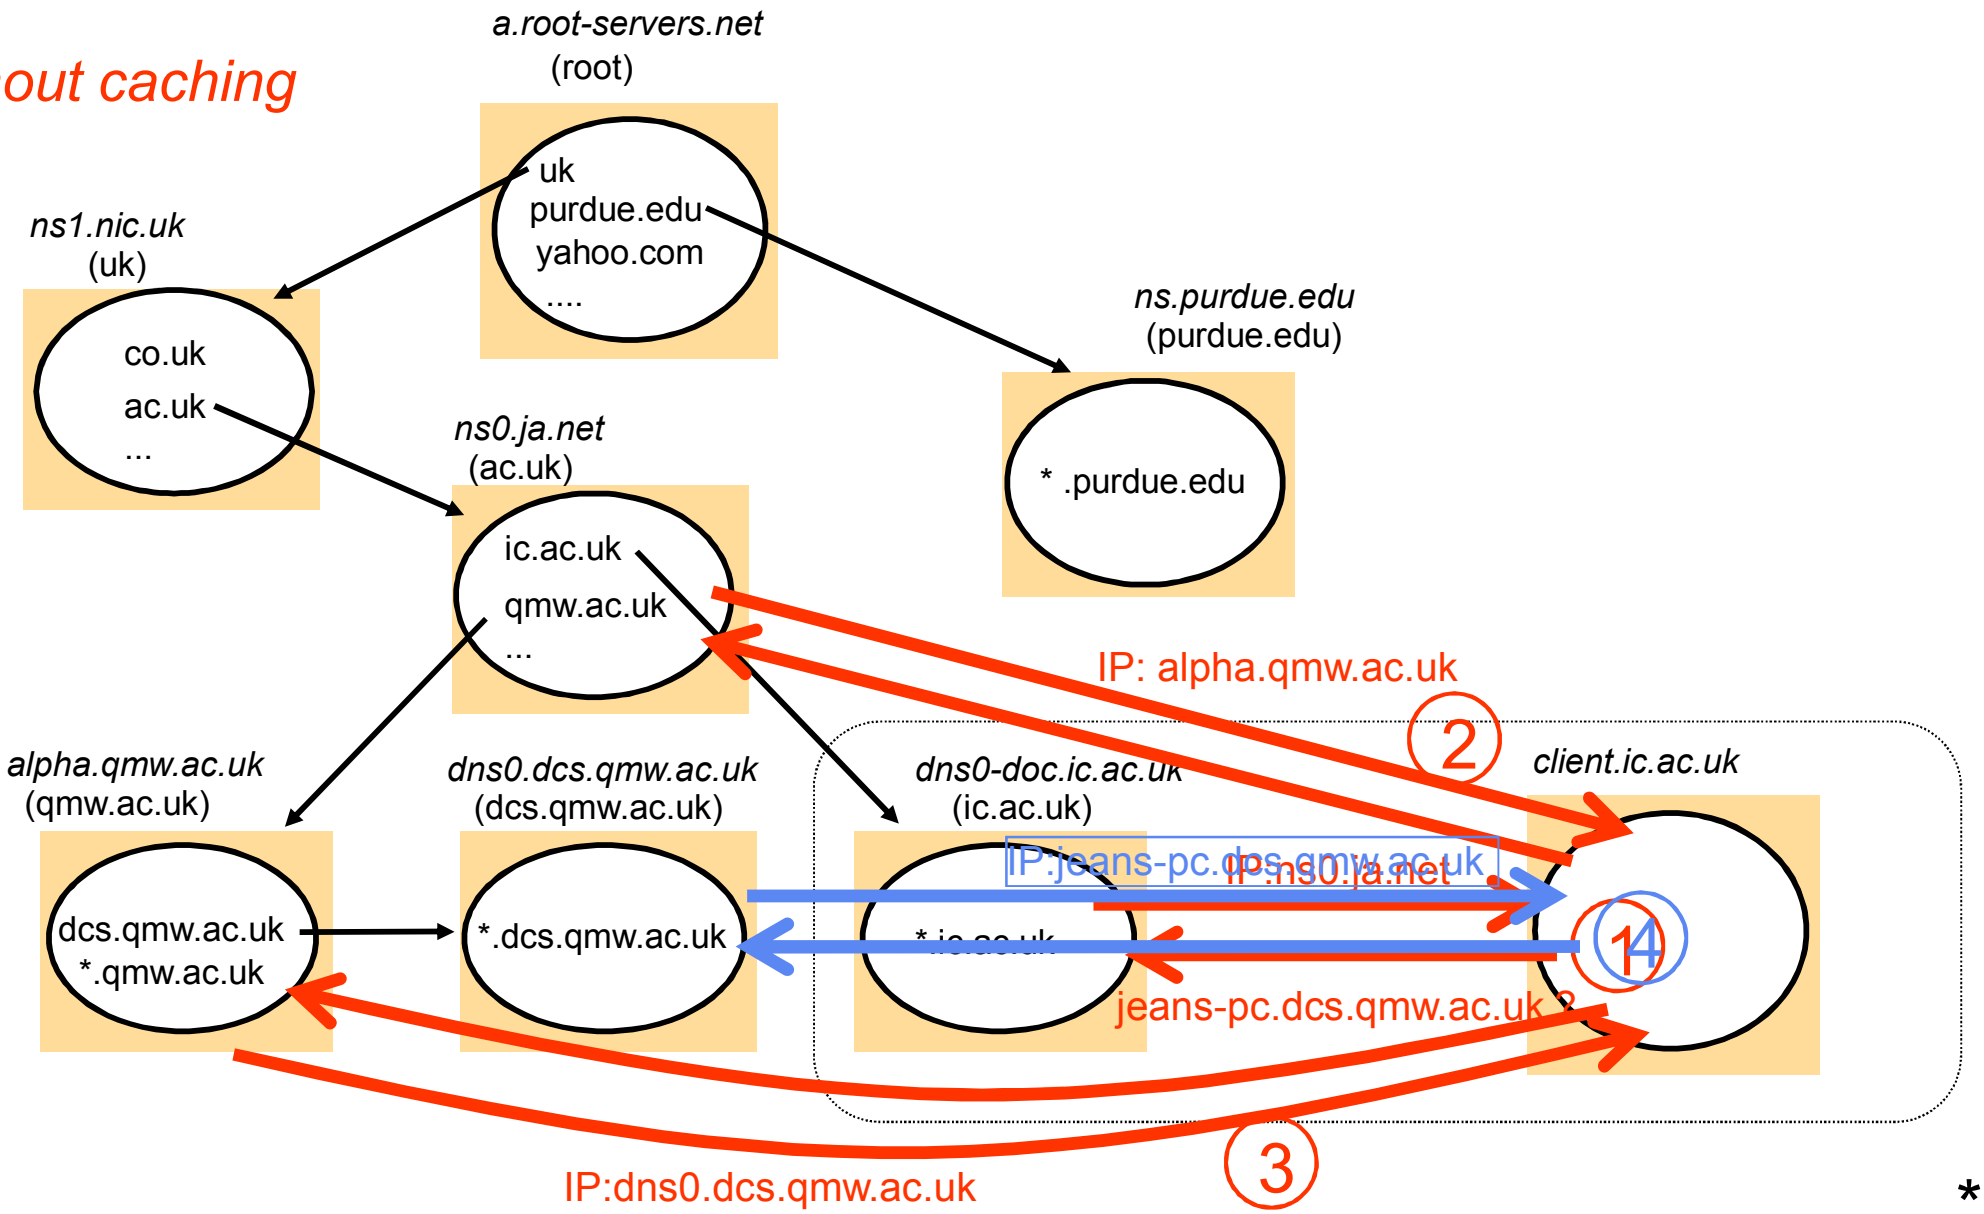

*

Resource Records

Type	Meaning	Value
SOA	Start of Authority	Parameters for this zone
A	IP address of a host	32-Bit integer
MX	Mail exchange	Priority, domain willing to accept e-mail
NS	Name Server	Name of a server for this domain
CNAME	Canonical name	Domain name
PTR	Pointer	Alias for an IP address
HINFO	Host description	CPU and OS in ASCII
TXT	Text	Uninterpreted ASCII text

Example of resource record

```
hcmut.edu.vn.   IN SOA  hcmut-server.hcmut.edu.vn. webmaster.hcmut.edu.vn. (
                2004110800      ; serial
                7200            ; refresh
                3600            ; retry
                604800          ; expire
                86400 )         ; minimum
hcmut.edu.vn.                   IN      NS      vnuserv.vnuhcm.edu.vn.
hcmut.edu.vn.                   IN      NS      server.vnuhcm.edu.vn.
hcmut.edu.vn.                   IN      MX  10  webmailserv.hcmut.edu.vn.
hcmut.edu.vn.                   IN      MX  0   vnuserv.vnuhcm.edu.vn.
hcmut-server.hcmut.edu.vn.      IN      A       172.28.2.2
stu-mailserv.hcmut.edu.vn.      IN      A       172.28.2.3
webmailserv.hcmut.edu.vn.       IN      A       172.28.2.4
pop3.student.hcmut.edu.vn.      IN      CNAME   stu-mailserv.hcmut.edu.vn.
mailhost.student.hcmut.edu.vn.  IN      CNAME   stu-mailserv.hcmut.edu.vn.
www.student.hcmut.edu.vn.       IN      CNAME   stu-mailserv.hcmut.edu.vn.
```

DNS issues

- Name tables change infrequently, but when they do, caching can result in the delivery of stale data.
  - Clients are responsible for detecting this and recovering
- Its design makes changes to the structure of the name space difficult. For example:
  - merging previously separate domain trees under a new root
  - moving subtrees to a different part of the structure (e.g. if Scotland became a separate country, its domains should all be moved to a new country-level domain.

Global Name Service (GNS)

- Designed and implemented by Lampson and colleagues at the DEC Systems Research Center (1986)
- Provide facilities for resource location, email addressing and authentication
- When the naming database grows from small to large scale, the structure of name space may change

  → the service should accommodate it
- Cache consistency ?

GNS Structure

- Tree of directories holding names and values
- Muti-part pathnames refer to the root or relative working directory (like Unix file system)
- Unique Directory Identifier (DI)
- A directory contains list of names and references
- Leaves of tree contain value trees (structured values)

GNS directory tree and value tree