Domain Name System
&
Peer-to-Peer Networking

H.-Arno Jacobsen

University of Toronto
jacobsen@eecg.toronto.edu

www.eecg.toronto.edu/~jacobsen

Overview

- reflection on the Internet and peer-to-peer applications
- the domain name system
- architecture and design of DNS

The "Evolution" of the Internet

- from a peer-to-peer network (ARPANET) to a client/server network (many clients !!)
- introduction of firewalls + NAT, to hide netw. behind
  - given NAT, a machine does not have a valid address at all
- millions of home users connected via slow modems via SLIP / PPP
- web browser boom (e.g., a Web-TV-box)
  - no need for **continuous** connection
  - no need for **permanent** addressno support for **multiple users**
- dynamic IP address assignment

What does this mean for peer-to-peer applications ?

- scalable, secure Internet required due to commercialization
- brought **millions of clients** on the Internet

However, peer-to-peer applications, require

- participants **serve** resources as well as **use** them
- i.e., "clients" to become **identifiable peers**
- must be **reachable** consistently

Names rather than numbers

- names for us; numbers for machines
- I.e., ease to read vs. ease to store, process, manipulate …
- where do names intervene (essentially use of DNS)
  - `ftp.eecg.toronto.edu`
  - `http://www.eecg.toronto.edu/~jacobsen/bio.html`
  - `jacobsen@eecg.toronto.edu`
  - `telnet chocolate.eecg.utoronto.ca`

HOSTS.txt - *the early days*

- name-to-address mapping table was one file
- managed by Stanford research Institute (SRI)
- updates submitted by e-mail
- hosts.txt distributed by e-mail / ftp / UUCP
- sometimes a week to reach the "end" of the Net
- the Internet grew, grew, grew, and GROWs …
- became too "expensive" to manage
  - HOSTS.txt became too large
  - centralized file vs. distributed nature of Net
  - was not scalable

DNS: The Domain Name System

- domain-name-to-addr-mapping
  - used to resolve names for (e-mail, ftp, http, telnet, …)
- addr-to-domain-name-mapping
  - 123.232.23.2 is (??)
  - a.k.a. in-addr.arpa domain

DNS: Design Goals

- provide same information as HOSTS.txt
- manage database in distributed manner
- no obvious size limits for names, …
- provide acceptable performance
- extensible (wrt. information stored)
- encapsulate other name spaces
- be independent of network topology and operating system

The Name Space

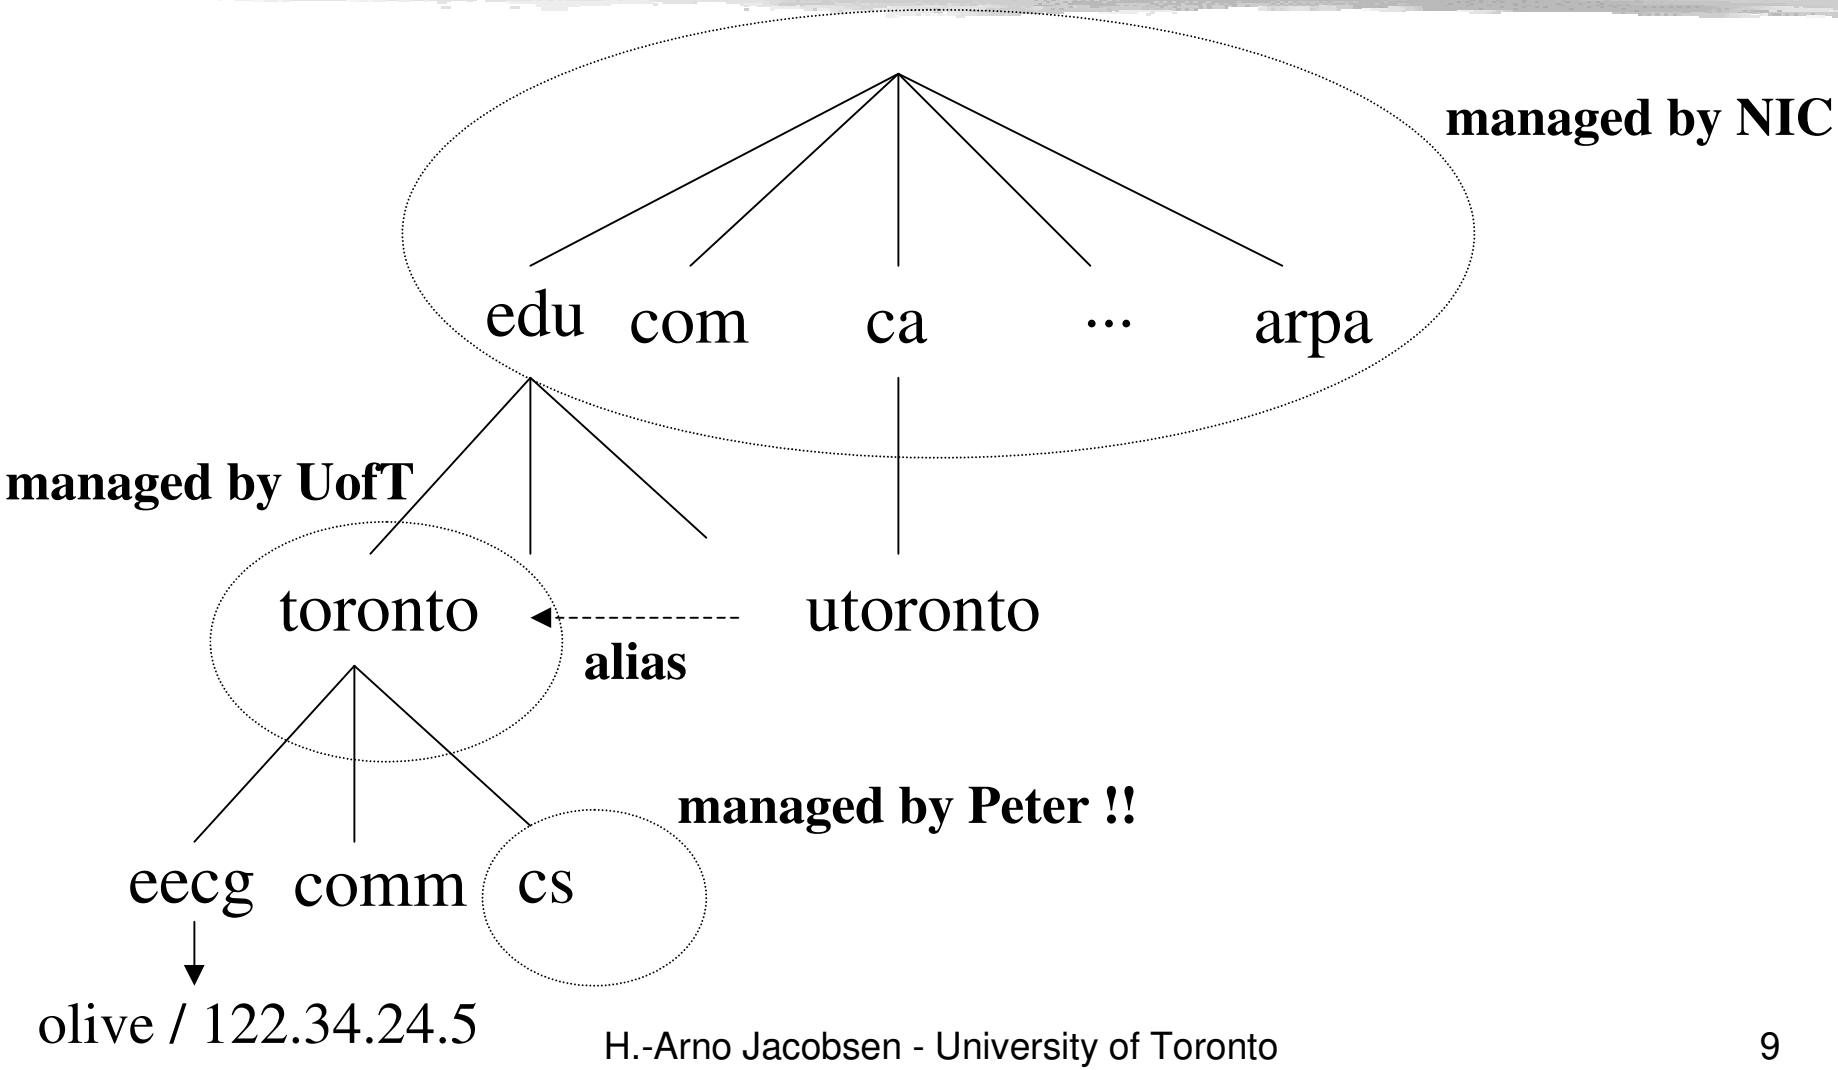

Architecture and Design

- DNS consists of
  - name servers
    - manage zones of name space
    - constitute databases with information
    - replicated for availability reasons
    - primary master server
    - secondary master server
    - used to be 7 root name servers
  - resolvers used by applications
    - clients that pass queries to name server(s)

Architecture continued

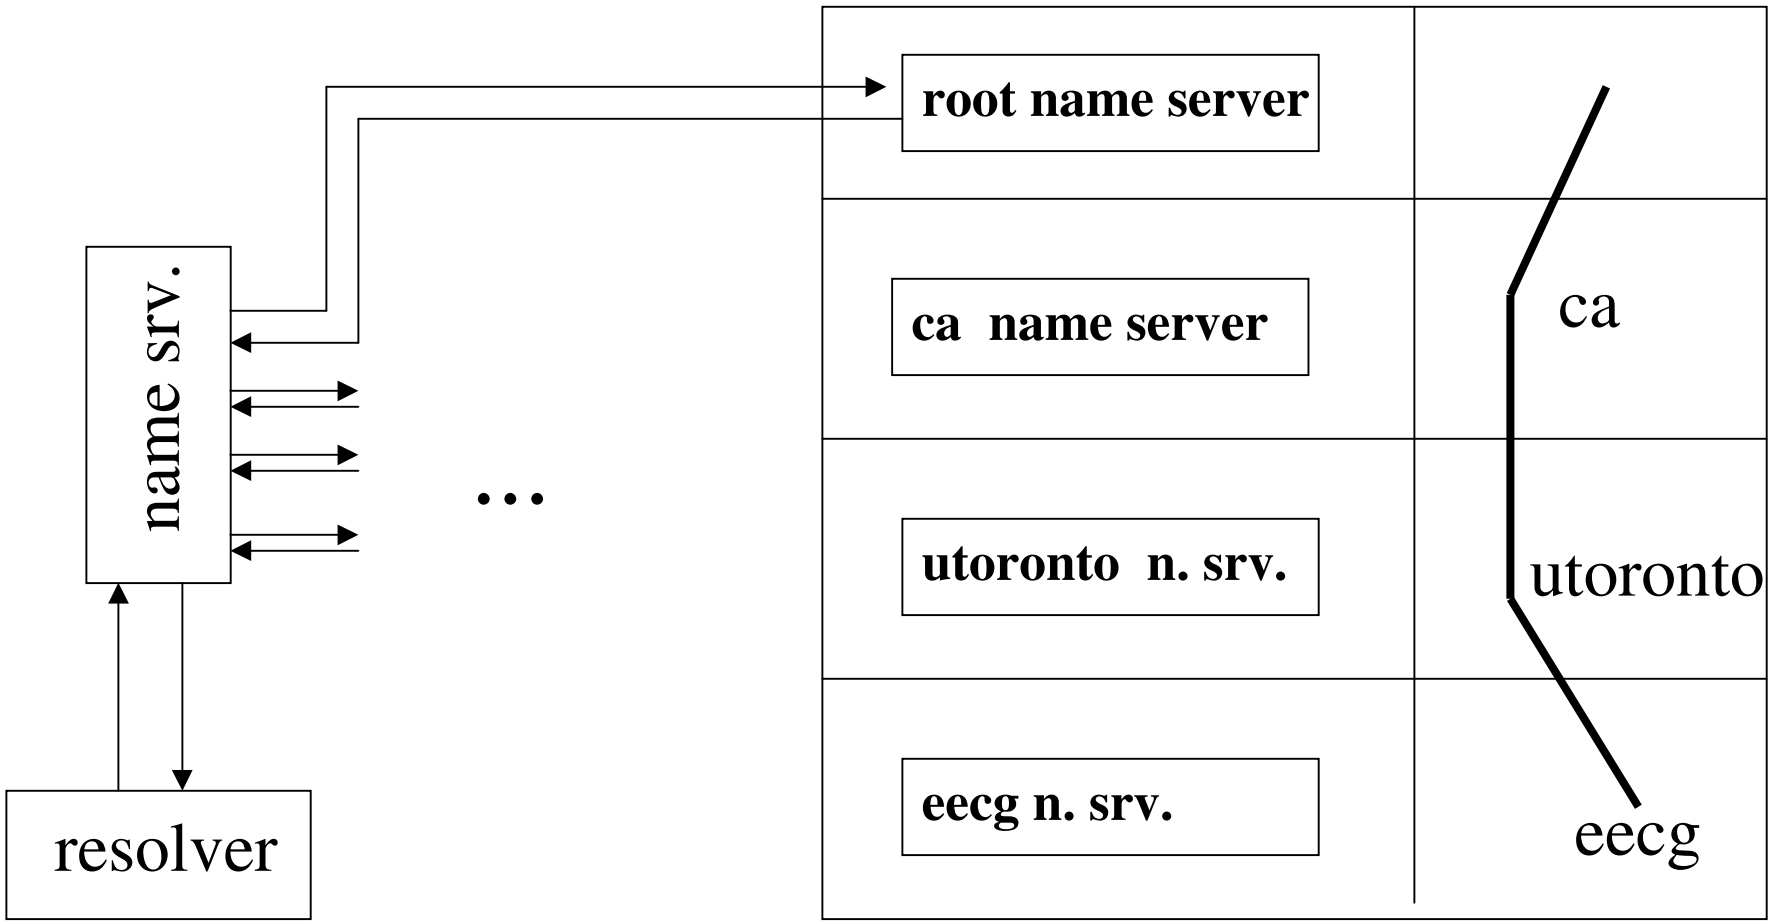

Architecture continued

- name servers deploy a cache
  - items in cache are flaged with a time to live filed (TTL)
  - "negative caches", for erroneous queries

Summary

- DNS: the Domain Name System replaced `HOSTS.txt`
- maps names to addresses, addresses to names, some other services
- uses replication to guarantee availability
- a corner stone of Internet, *cannot do without ?*
- *or can ?*
  - see emerging peer-to-peer applications addressed in subsequent seminars

References

- Computer Networks. L. Peterson, B. S. Davie. Morgan kaufman, 1996.
- DNS and BIND P. Albitz and C. Liu. O'Reilly & Associates, 1997.
- Development of the Domain name system. P. Mockapetris, K. Dunlap. ACM, 1988.