ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Overlay Networks

Enhancing the Internet ...

Seminar in Distributed Computing 2007

Ramon Schwammberger
ramons@ethz.ch

Contents

Part One

- overlay networks
- INS - intentional naming system
- i3 - Internet Indirection Infrastructure
- Active Names

Part Two

- delivery modes:
  - anycast
  - multicast
- late binding
- inter node routing
- caching
- soft state
- scalability
- security

Overlay Network

- build a new layer which forms an overlay network on top of the “IP network”

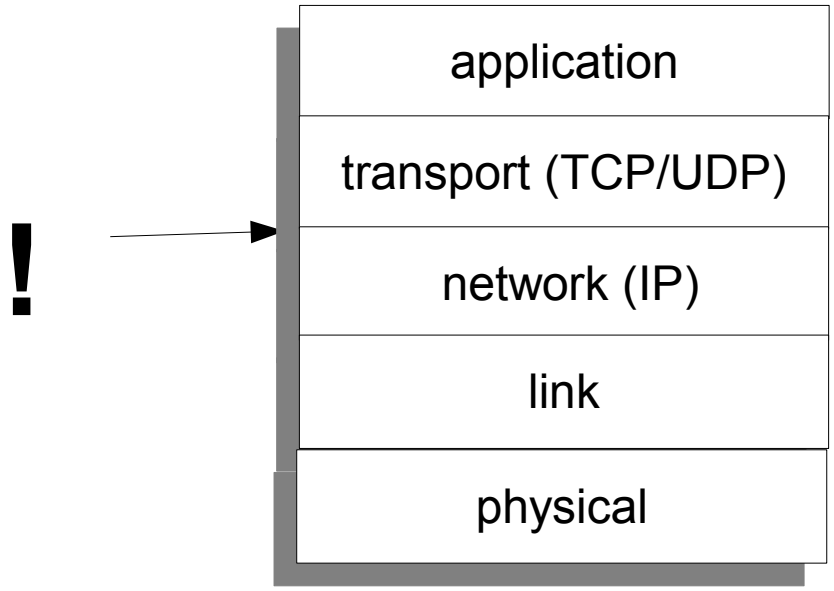

goals - overview

INS

- expressiveness
- responsiveness
- robustness
- easy configuration

i3

- multicast
- anycast
- mobility
- end-to-end principle

Active Names

- customization / extensibility
- composability
- efficient resource management
- location independent execution

From Metanet to INS

- use a central network
- have a name and an address for everything "model the world"

"Metanet View"

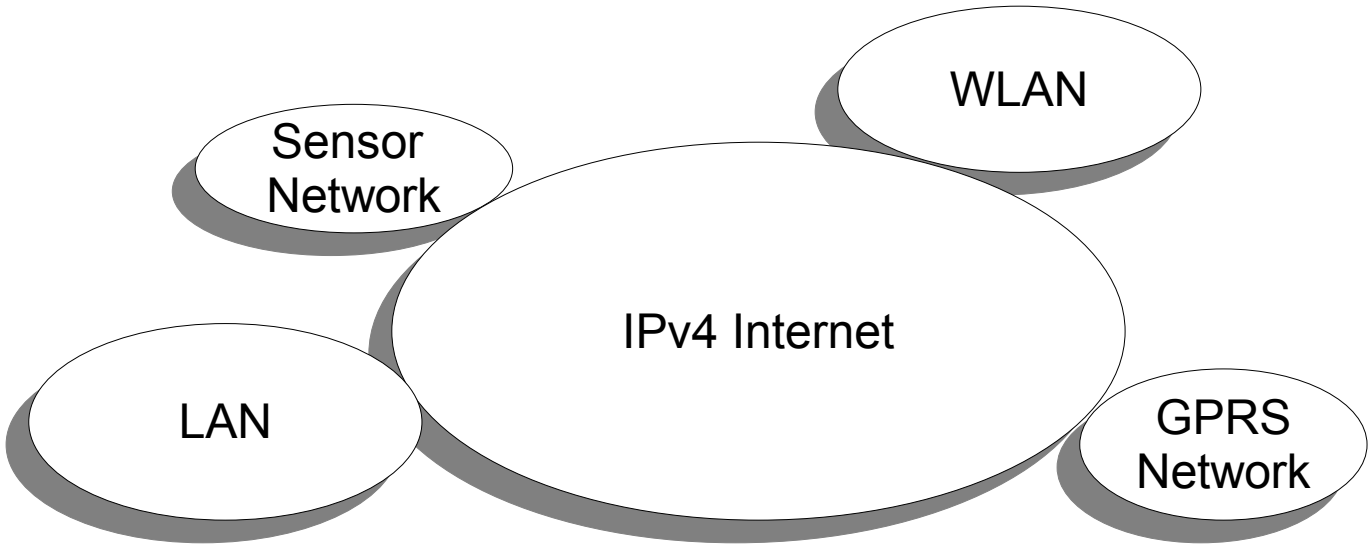

From Plutarch to INS

"INS View"

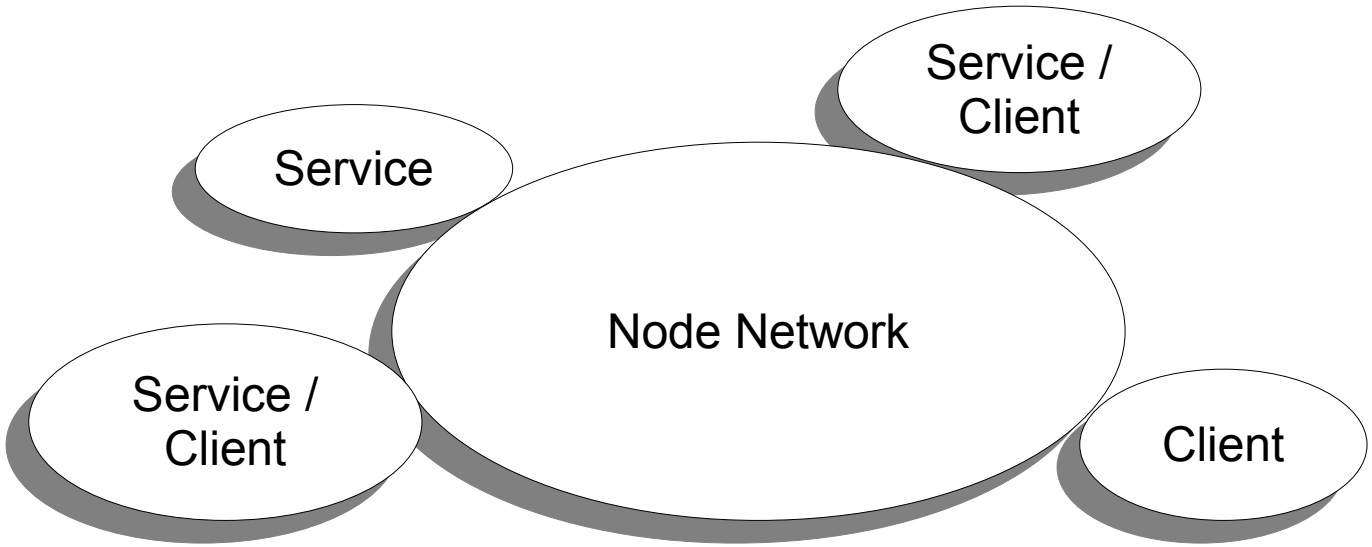

INS – resolver network

- INR ●
- DSR ○

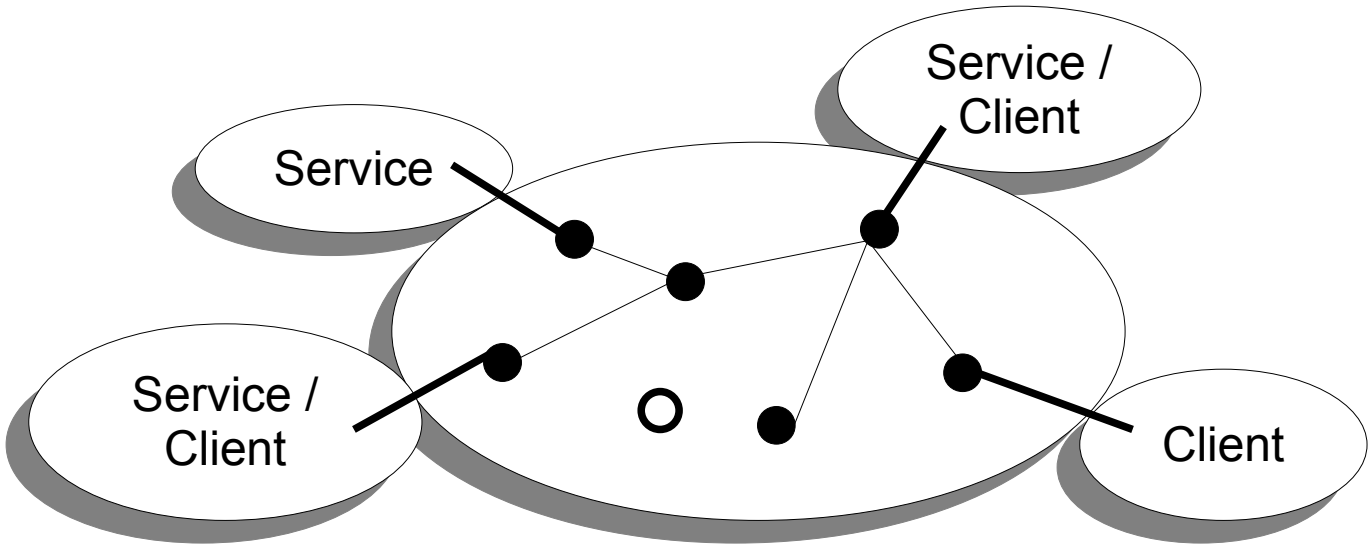

INS – name-specifier

- building blocks: attribute=value pairs

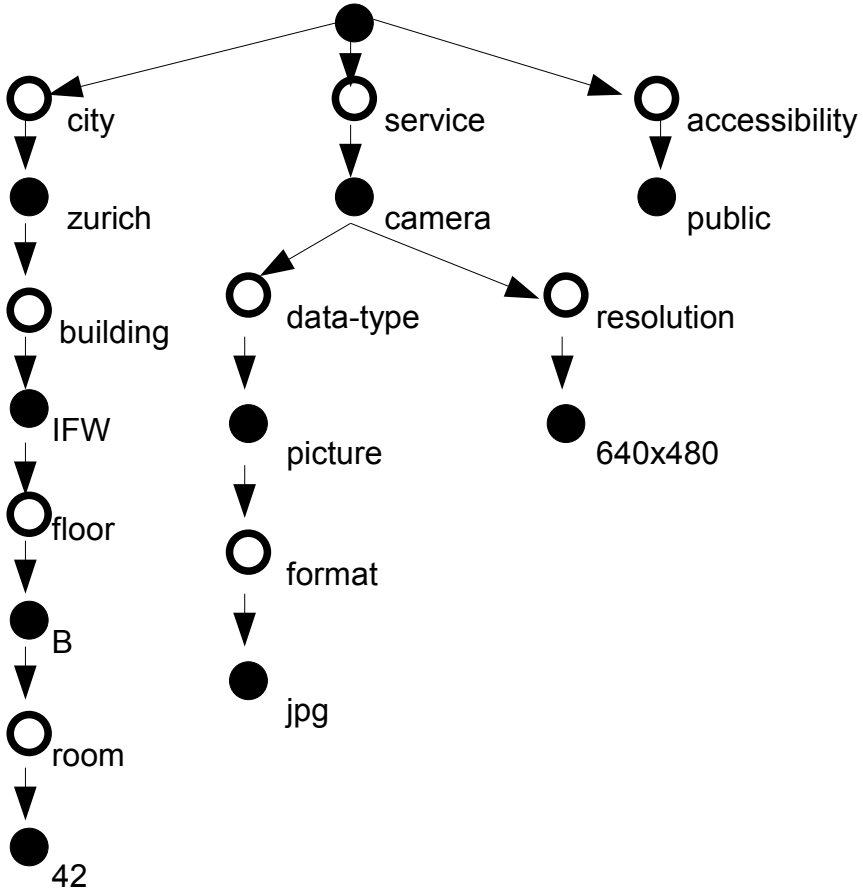


```
[city=zurich
      [building=IFW
            [floor=B
                  [room=42]]]

[service=camera
      [data-type=picture
            [format=jpg]]
      [resolution=640x480]]

[accessibility=public]
```

INS - example

INS message:

{ B | D | destination pointer | source pointer | body}

B – binding flag (early or late)

D – delivery flag (true -> multicast)

destination

[service=camera [entity=transmitter]] [room=510]

source

[service=camera [entity=receiver]] [id=r] [room=500]

From INS to i3

- introduce rendez-vous based sending / receiving
- remove “model the world”
  - use a simple naming scheme
  - name rendez-vous points not services
- change the architecture of the resolver network

goals - overview

INS

- expressiveness
- responsiveness
- robustness
- easy configuration

i3

- multicast
- anycast
- mobility
- end-to-end principle

Active Names

- customization / extensibility
- composability
- efficient resource management
- location independent execution

i3 – rendez-vous point

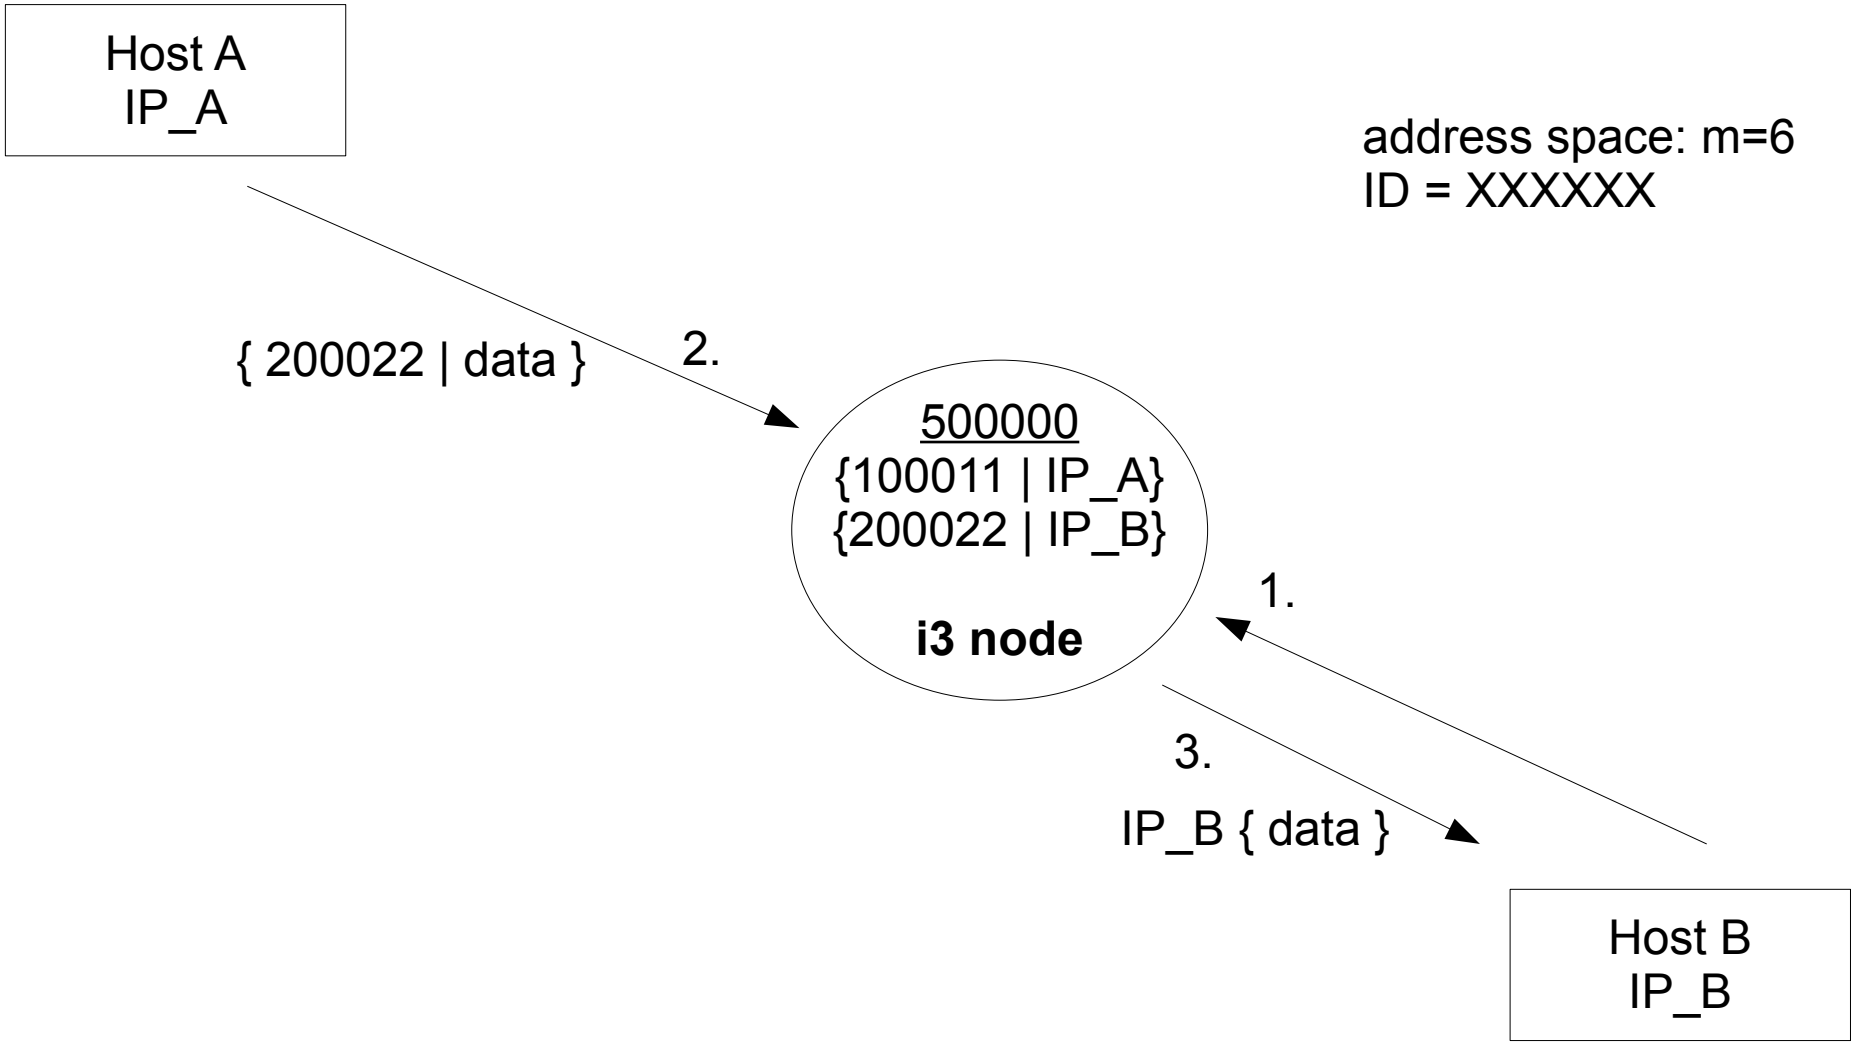

i3 - details

- packet header

  { ID | source | data }

  { ID_1, ID_2, ... , ID_N | source | data }

- trigger

  { ID | receiver }

  { ID | receiver1, receiver2, ... , receiverN }

Plutarch to Active Names

- Plutarch's interstitial functions

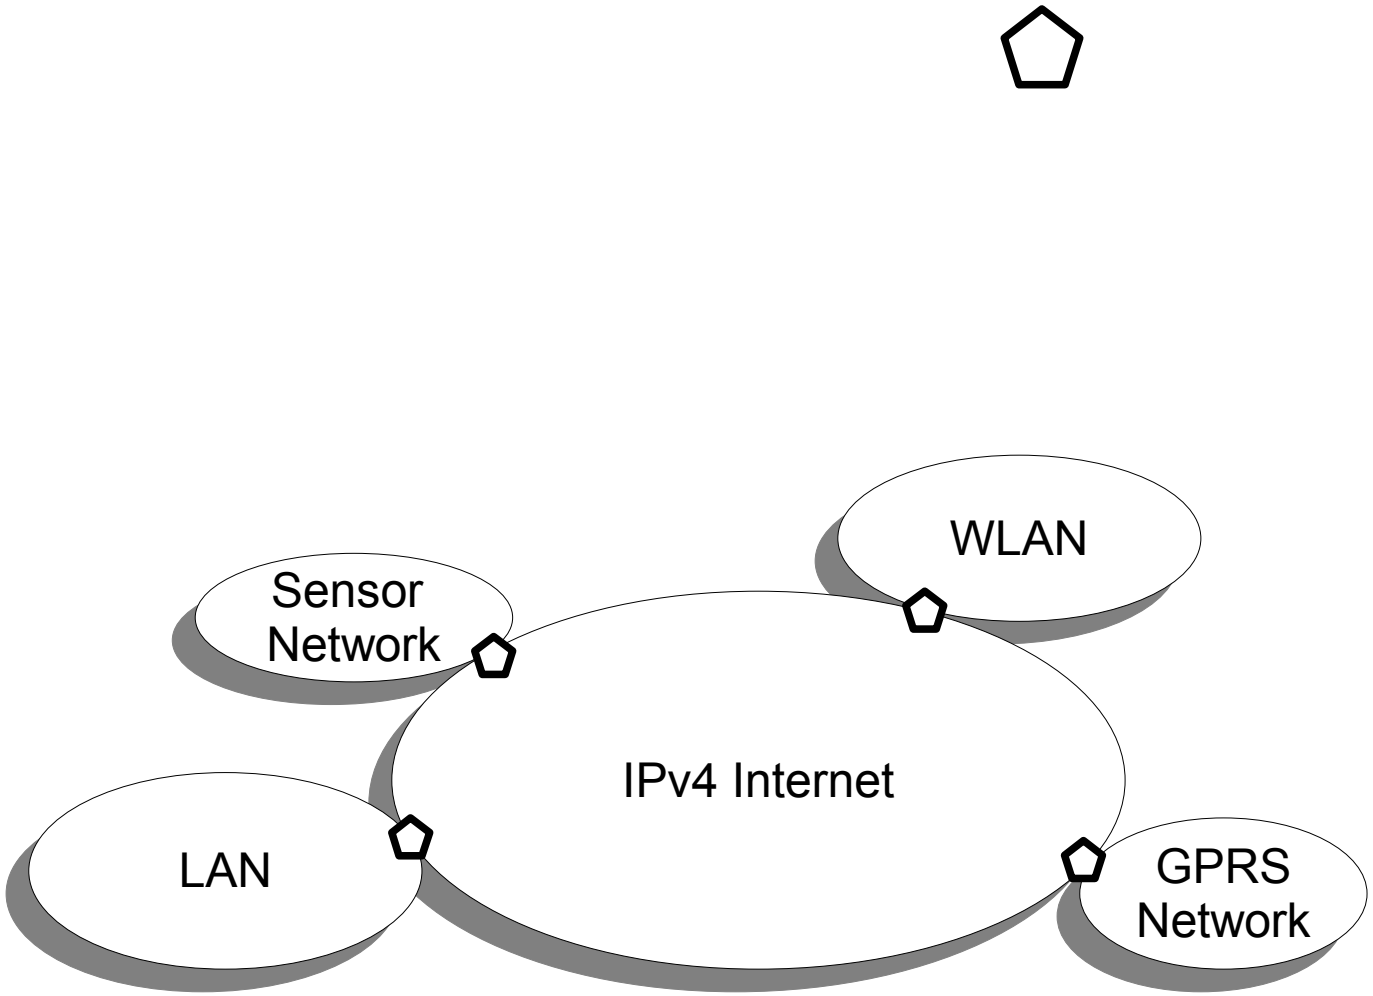

Plutarch to Active Names

- Active Names running programs

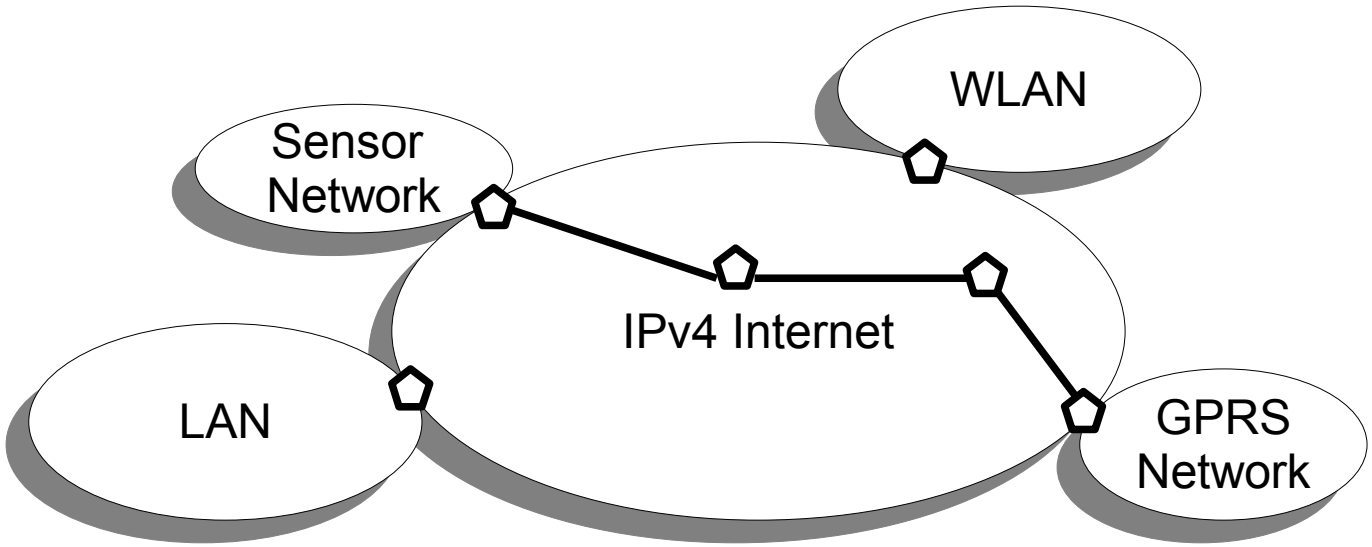

goals - overview

INS

- expressiveness
- responsiveness
- robustness
- easy configuration

i3

- multicast
- anycast
- mobility
- end-to-end principle

Active Names

- customization / extensibility
- composability
- efficient resource management
- location independent execution

Active Names - details

- hierarchical name space
  - programs define how remaining name shall be resolved
- micro kernel approach for resolver
  - loader
  - execution environment
  - interface for remote invocation of programs

Active Names - details

- after methods
  - describe the “path” from the service back to the client
  - can be reordered by any program

Active Names - example

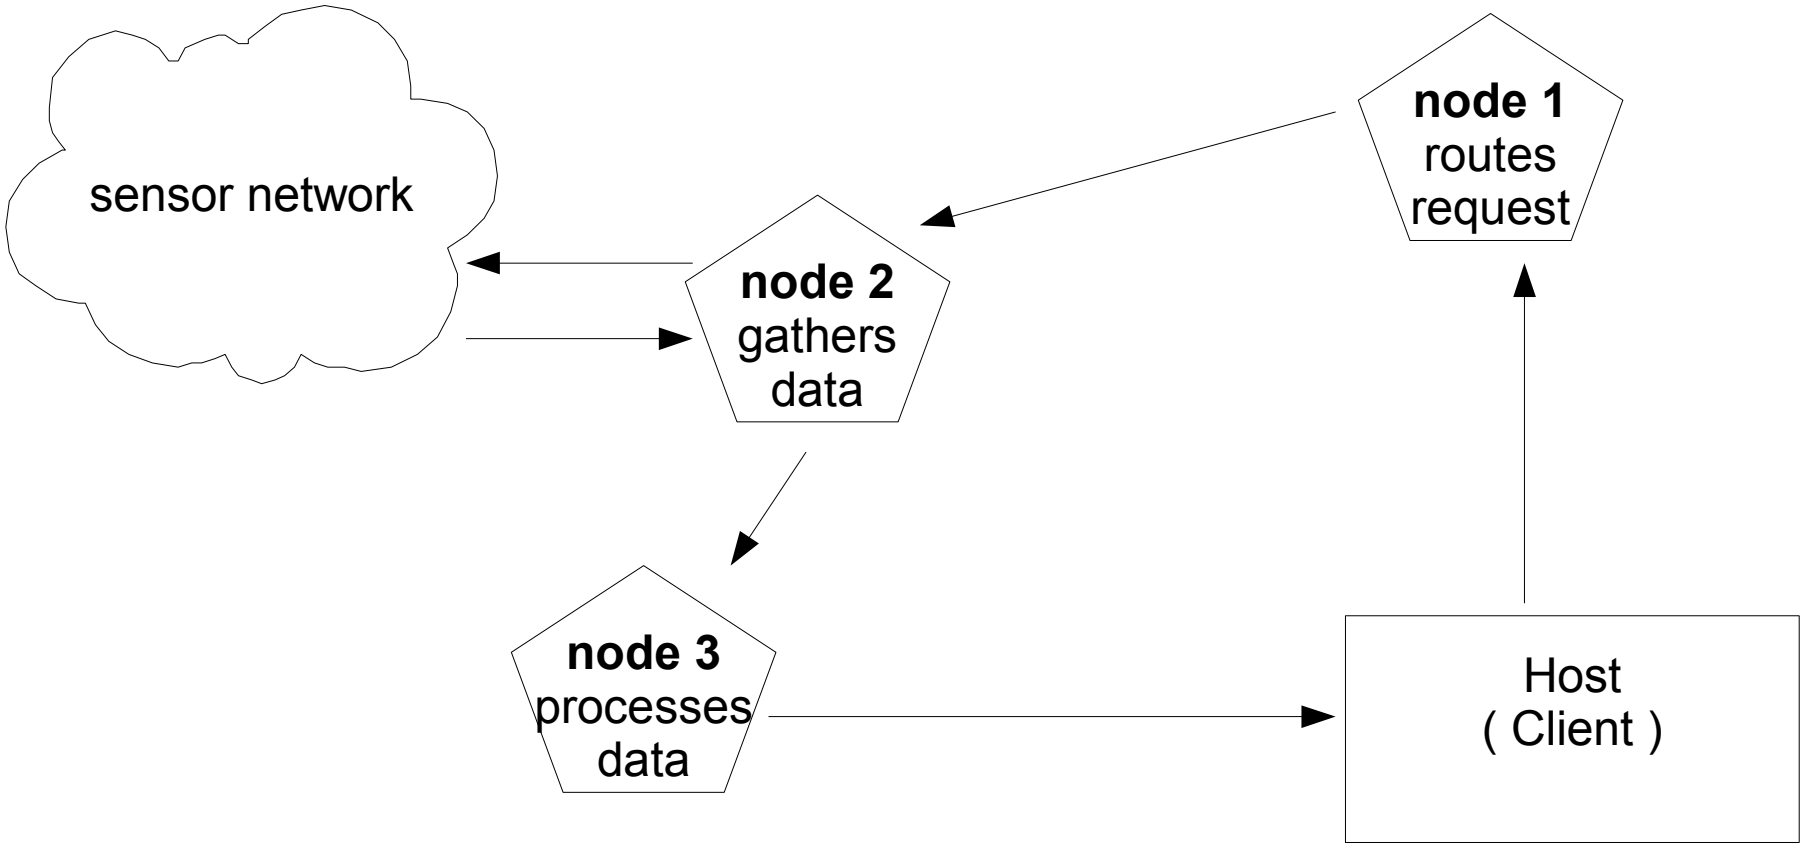

Contents

Part One

- overlay networks
- INS - intentional naming system
- i3 - Internet Indirection Infrastructure
- Active Names

Part Two

- delivery modes:
  - anycast
  - multicast
- late binding
- inter node routing
- caching
- soft state
- scalability
- security

delivery modes

INS

- switch between anycast and multicast (groupcast)

-> delivery flag

i3

- anycast
- multicast (groupcast)
- anycast between groups

-> address space xxxxyy

x-part for service

y-part for anycast

Active Names

- anything that can be programmed

late binding

INS

- switch between late and early binding

-> binding flag

i3

- always kind of late binding
- goes always through rendez-vous point

Active Names

- is programmable
- most useful with late binding (send data for nodes)

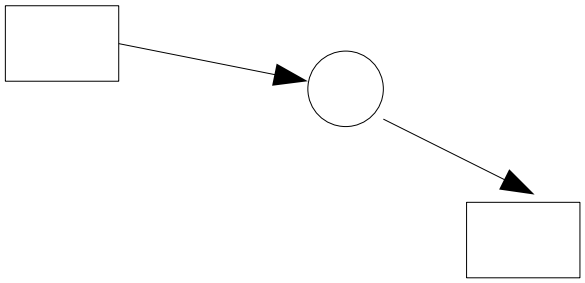

inter node routing

INS

- always to next hop in INR spanning tree
- flooding to announce services

i3

- distributed hash table (Chord)

Active Names

- name tells the program where to send the packet to

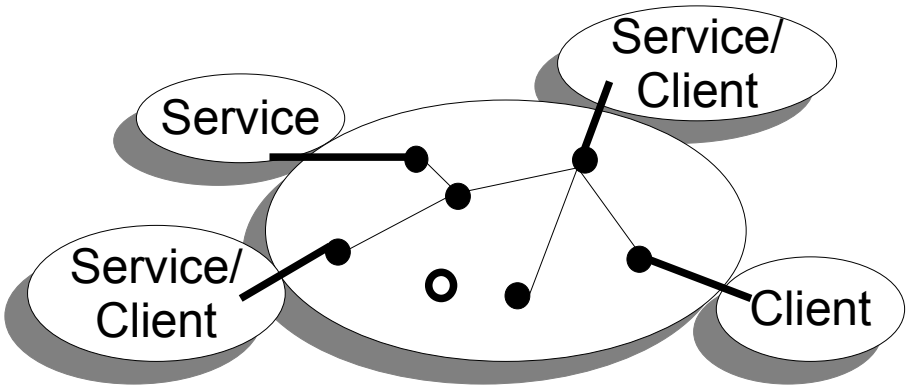

caching

INS

- cache content

{ B | D | dst ptr | src ptr |
cache lifetime | body }

i3

- cache IP of node which is responsible for ID

Active Names

- is programmable
- (content caching)

soft state

INS

- service advertisements are soft state
- easy recovery
- easy logout
- up to date

i3

- triggers are soft state
- up to date
- easy recovery

Active Names

- state is handled by programs

scalability

INS

- not designed for wide area application
- DSR not hierarchical
- scales to several thousand services

i3

- designed for wide area

Active Names

- hierarchical name spaces scale

INS

- not addressed in the paper
- needs to be explored

i3

- target: at least as secure as IP layer
- suggestions for
  - eavesdropping
  - trigger hijacking
  - DoS attacks

Active Names

- depends on “sandbox”
- depends on resolver programs

conclusion

INS

- good for services
- small scale
- good prototype

i3

- useful enhancement of the internet
- wide area
- good prototype

Active Names

- most adaptive solution
- complexity?
- security?

Thank you
for
your attention