

CS 268: Project Suggestions

Ion Stoica
February 6, 2003

Project Related with Internet Indirection Infrastructure (*i3*)

- Goal: provide an uniform abstraction for basic communication primitives:
 - Anycast
- Next: overview of *i3*

istoica@cs.berkeley.edu

2

Motivations

- Today's Internet is built around a point-to-point communication abstraction:
 - Send packet "p" from host "A" to host "B"
- This abstraction allows Internet to be highly scalable and efficient, but...
- ... not appropriate for applications that require other communication abstractions:
 - Multicast
 - Anycast
 - Mobility
 - ...

istoica@cs.berkeley.edu

3

Why?

- Point-to-point communication abstraction implicitly assumes that there is one sender and one receiver, and that they are placed at fixed and well-known locations
 - E.g., a host identified by the IP address 128.32.xxx.xxx is most likely located in the Berkeley area

istoica@cs.berkeley.edu

4

Key Observation



- All previous solutions use a simple but powerful technique: **indirection**
 - Assume a logical or physical indirection point interposed between sender(s) and receiver(s)
- **Examples:**
 - IP multicast assumes a logical indirection point: the IP multicast address
 - Mobile IP assumes a physical indirection point: the home agent

istoica@cs.berkeley.edu

5

Our Solution

- Add an efficient indirection layer (IL) on top of IP
 - Transparent for legacy applications
- Use an **overlay network** to implement IL
 - Incrementally deployable; don't need to change IP




istoica@cs.berkeley.edu

6

Internet Indirection Infrastructure

- Change communication abstraction: instead of point-to-point, exchange data by name
 - Each packet is associated an identifier ID
 - To receive a packet with identifier ID, receiver R maintains a trigger (ID, R) into the overlay network



istoica@cs.berkeley.edu

7

Service Model


- Best-effort service model (like IP)
- Triggers are periodically refreshed by end-hosts
- Reliability, congestion control, and flow-control implemented at end-hosts

istoica@cs.berkeley.edu

8

Mobility

- Host just needs to update its trigger as moves from one subnet to another
- Both sender and receiver can be mobile
- Can eliminate the "triangle routing problem"




istoica@cs.berkeley.edu

9

Mobility

- Host just needs to update its trigger as moves from one subnet to another
- Both sender and receiver can be mobile
- Can eliminate the "triangle routing problem"




istoica@cs.berkeley.edu

10

Multicast

- Unifies multicast and unicast abstraction
 - Multicast: receivers insert triggers with the same identifier
- An application can dynamically switch between multicast and unicast




istoica@cs.berkeley.edu

11

Composable Services

- Use a stack of IDs to encode the successions of operations to be performed on data (e.g., transcoding)
- Advantages
 - Don't need to configure path
 - Load balancing and robustness easy to achieve




istoica@cs.berkeley.edu

12

Composable Services (cont'd)

- Both receivers and senders can specify the operations to be performed on data




istoica@cs.berkeley.edu

13

Anycast

- Generalize the matching scheme used to forward a packet
 - Until now we assumed exact matching
- Next, we assume:
 - Exact matching on the most significant l bits of ID
 - Longest prefix matching on the remaining bits (ID size = m)




istoica@cs.berkeley.edu

14

Anycast (cont'd)

- Anycast is simply a byproduct of the new matching scheme
 - Each receiver in the anycast group inserts IDs that differ only in the last $l-m$ bits



istoica@cs.berkeley.edu

15

Anycast (cont'd)

- Highly flexible: the least significant $l-m$ bits of ID are application specific
- Two examples:
 - Load balancing
 - Proximity

istoica@cs.berkeley.edu

16

Idea 1: Load Balancing

- Assumptions:
 - N servers of capacity C_i , $1 \leq i \leq N$
 - M clients downloading files from these servers
- Goal: come up with an algorithm to insert triggers and set up their identifiers such that to balance the load in the presence of server failures

istoica@cs.berkeley.edu

17

Idea 2: Transcoding Application

- Design a transcoding application
 - From one video format to another (e.g., MPEG \rightarrow H.263), or
 - From one data format to another (e.g., HTML \rightarrow WML)
- Note: the goal of the project is not to design the transcoder, but to demonstrate the service composition function

istoica@cs.berkeley.edu

18

Idea 3: Migrate-able End-to-End Protocols

- Design a congestion control mechanism (e.g. TCP) such that it is possible to change the receiving machine in the middle of the transfer!
- A and B open a connection (A receiver; B source)
- A changes to A'
- B continues to send data to A' without creating a new connection
- Challenge: transparently transfer the receiver state from A to A'

istoica@cs.berkeley.edu

19

Other Project Ideas

istoica@cs.berkeley.edu

20

Idea 4: Reducing (elimination) Multicast State in Routers

- Today each router maintain state for each multicast group that has traffic traversing it
- Problem: state is hard to maintain and manage → not scalable
- Extreme solution: maintain all receiver addresses in each packet
 - Routers don't need to maintain any state, but
 - Packet headers can become very large → huge overhead
- Solution: design an algorithm in between
 - Maintain some state in routers and some in packets
- Note: you can think either at the IP or application layer

istoica@cs.berkeley.edu

21

Forwarding in Low Energy Wireless Networks

- Problem: each node cannot afford to remain ON all the time
 - a node can communicate/receive data only when it is ON
- Two nodes can communicate only when both of them are simultaneously ON
- A node stores a packet in transit until it finds the next hop ON

istoica@cs.berkeley.edu

22

Ideas 5 & 6


- Assume routing tables are known
- Assume that each node is independently switching between ON and OFF states
- Idea 5:
 - Study the tradeoff between the fraction of time a node is ON and the time to deliver a message and the amount of storage required by a node
- Idea 6:
 - Design a self-synchronization algorithm and study its properties (i.e., a distributed algorithm that will result in all nodes being ON at the same time)

istoica@cs.berkeley.edu


23

Idea 7: Implement Round Robin at the Application Layer

- Problem: flow isolation (UDP can kill TCP)



- Solution outline:



istoica@cs.berkeley.edu

24

Idea 8: N-TCP

- Design a congestion control algorithm that provides a throughput equivalent to N individual TCPs between the same source and destination

istoica@cs.berkeley.edu

25

Ideas 9 & 10: Edge Control

- Consider a network domain in which you can only control all edge nodes, but not core nodes
- Idea 9: Derive an efficient measurement algorithm to infer the (approximate) topology and link capacities
- Idea 10: Assuming that you know the domain topology, what kind of services can you provide and how
 - Bandwidth and loss guarantees
 - What about delay?

istoica@cs.berkeley.edu

26

Next Step

- You can either choose one of the projects we discussed during this lecture, or come up with your own
- Pick your partner, and submit a one page proposal by February 13. The proposal needs to contain:
 - The problem you are solving
 - Your plan of attack with milestones and dates
 - Any special resources you may need

istoica@cs.berkeley.edu

27