IP MULTICAST

Adiseshu Hari, T. V. Lakshman and Gordon Wilfong
Nokia Bell Labs

DIMACS Workshop on Algorithms for Data Center Networks
Rutgers University, NJ

Why is IP Multicast not deployed in public networks?

- Denial-of-Service (DoS) attack amplification
- Complex Control Plane
- Large Forwarding state
  - Non aggregable

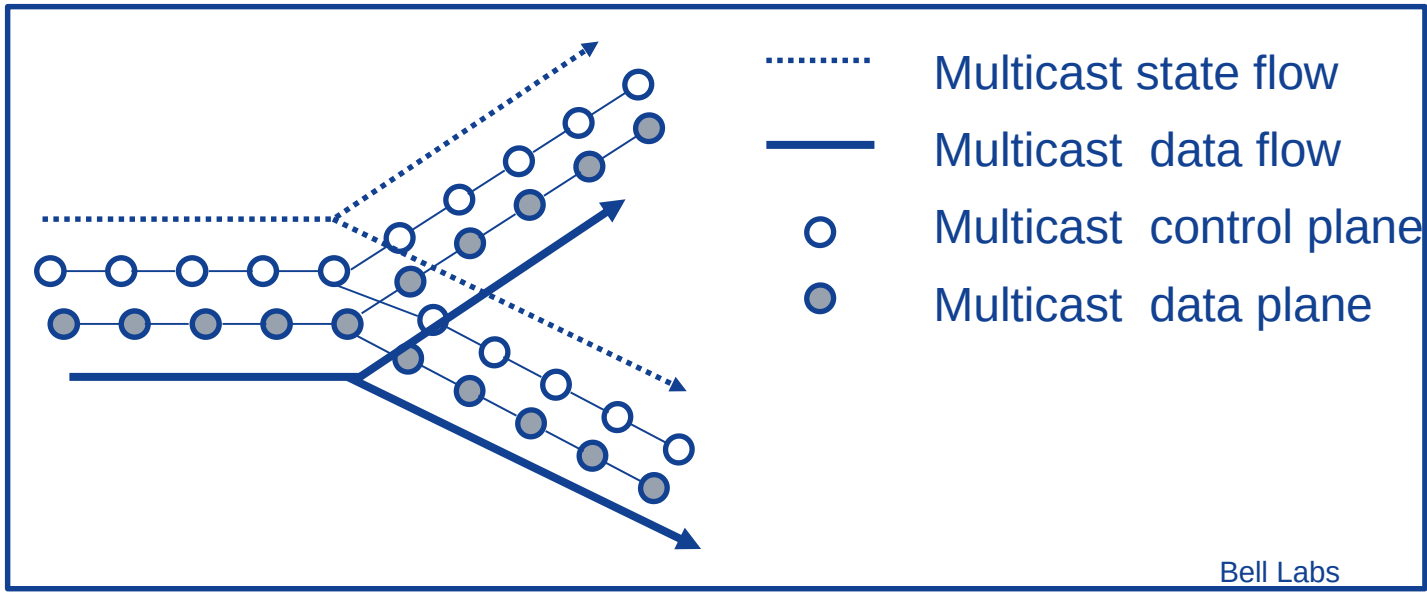

Can SDN help with Multicast?

- Denial-of-Service (DoS) attack amplification ✔
- Control state ✔
- Forwarding state ✘

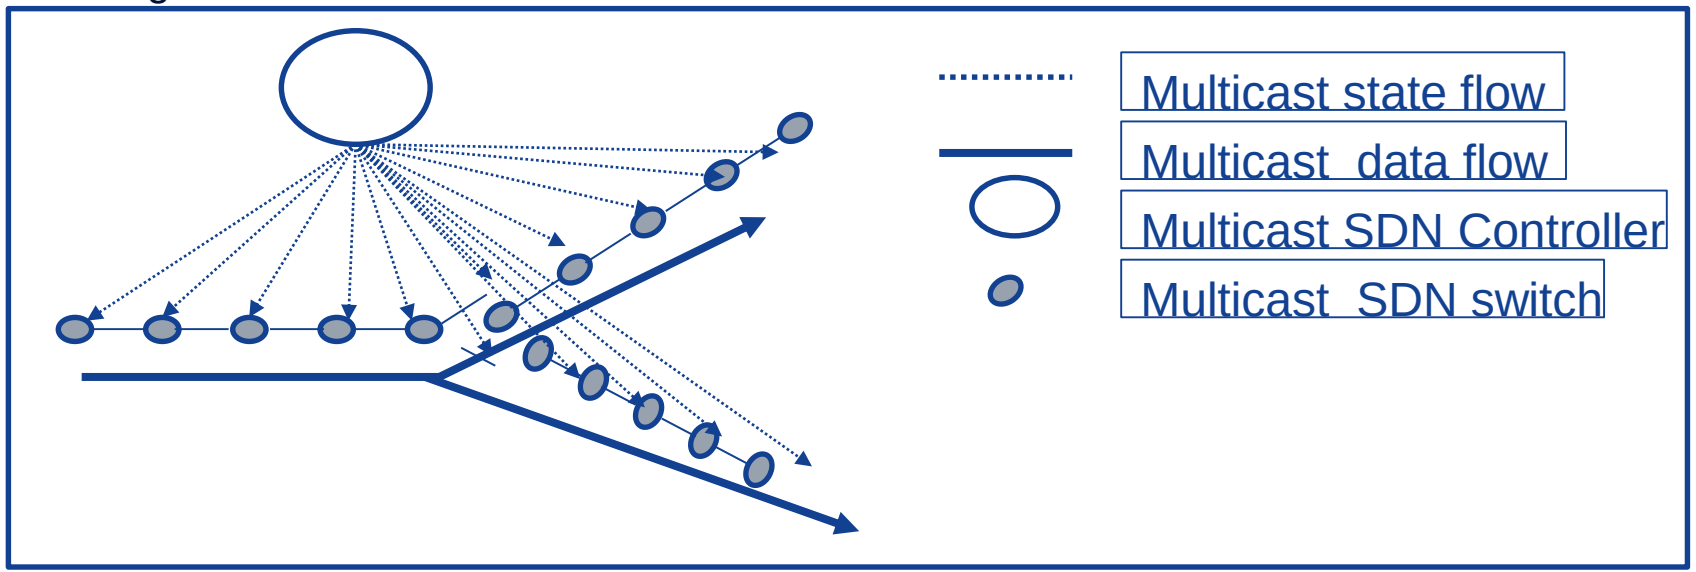

Can we eliminate multicast forwarding state in SDN?

- Eliminate unicast forwarding state in SDN:
  - Path Switching: per-flow routing without per-flow state
  - New data path suitable for SW switches and programmable packet processors
  - Encode path in the packet headers
  - DIMACS 2016

Eliminating unicast forwarding state in SDN using Path Switching

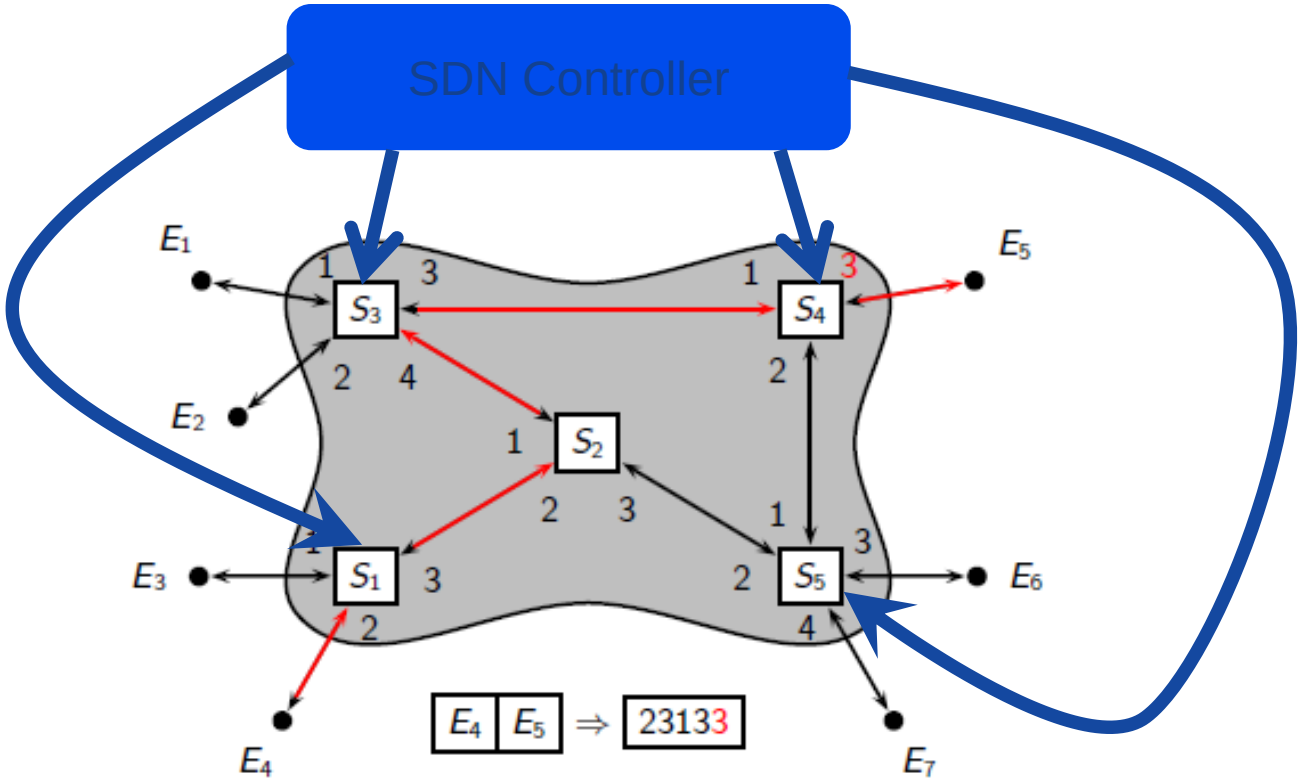

Can we eliminate multicast forwarding state in SDN?

Can we extend Path Switching to encode multicast paths?

Can we create an efficient encoding of a multicast path?

- No blowup in packet size (e.g. using bitmaps)
- No blowup in storage state (e.g,. encode each multicast tree by a unique identifier)

Can we reduce multicast forwarding state in SDN?

- Unicast Branching (UB)
  - Use branching nodes in the network to replicate unicast flows.
  - Use SDN Flow Table at ingress and egress
  - Use SDN Group Table at branching nodes
- Reduces multicast forwarding state from linear to sublinear in number of forwarding nodes

Unicast Branching (UB) Reference Diagram

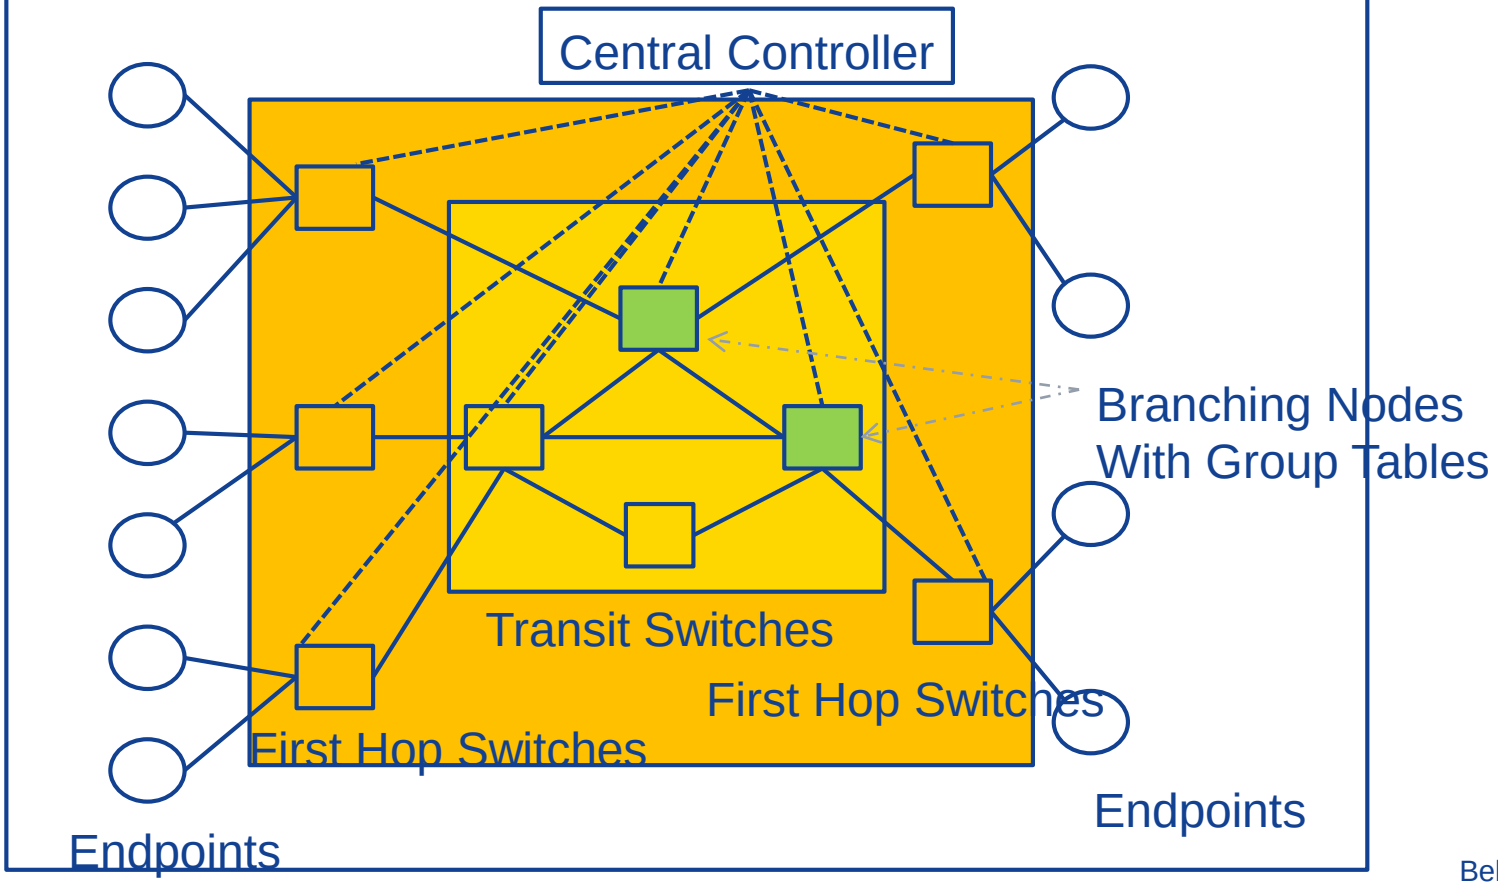

Reducing multicast forwarding state in SDN using Unicast Branching (UB)

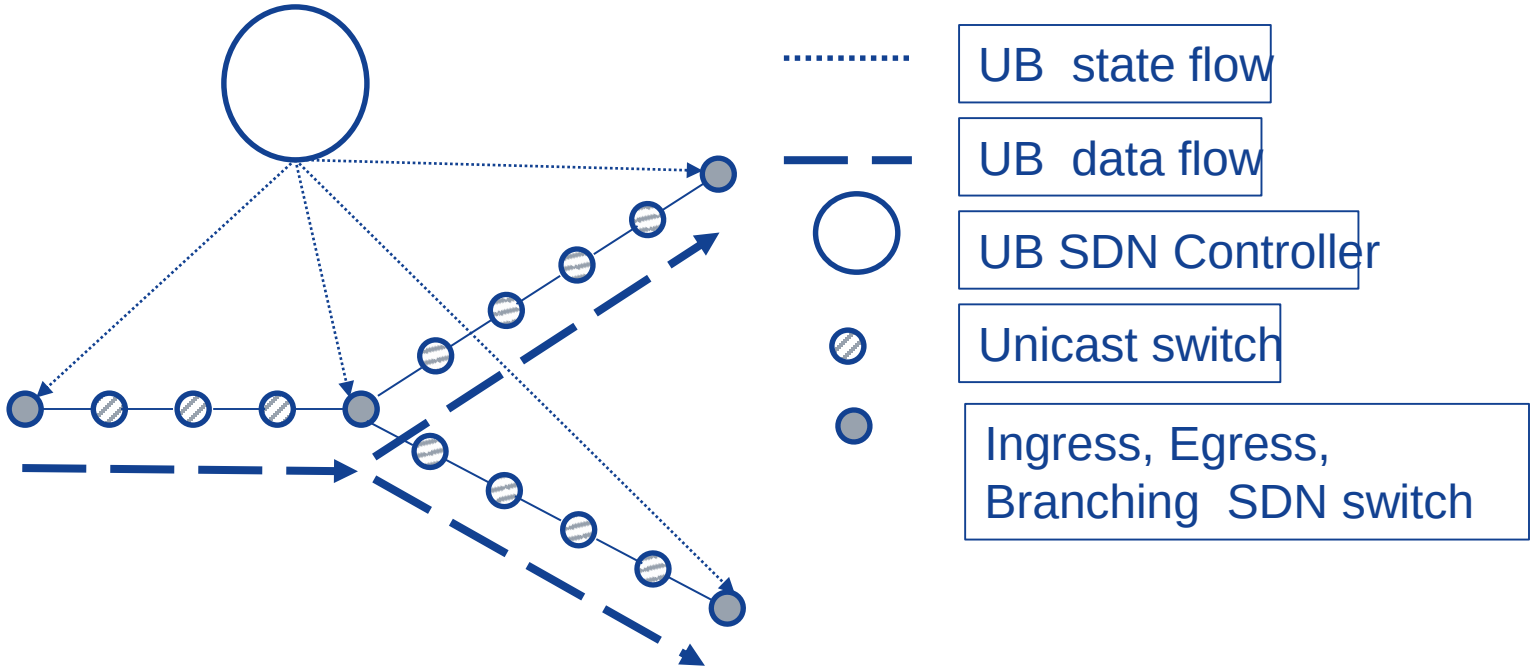

Added advantages of Unicast Branching (UB)

- Tunable knob to switch between unicast replication and full multicast
- Allows for an NFV based implementation
- Allows Traffic Engineered branches
  - Fast Reroute, Per branch QoS
- Works at all protocol layers – protocol agnostic
  - Ethernet, IP, MPLS

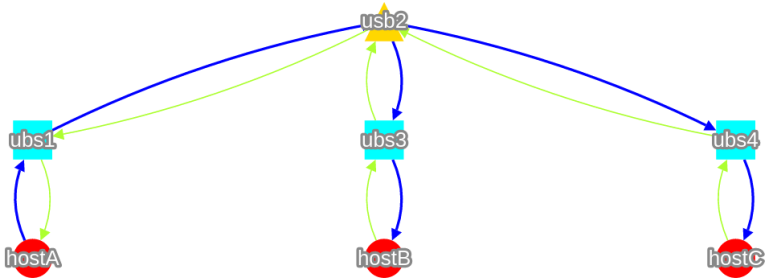

- Enables unicast only protocols like Segment Routing and TCP to be multicast capable*
  - HTTP Adaptive Streaming multicast
  - Efficient content caches
- Enables Policy Based Multicast

* Requires stateful NFV elements, not just SDN switches for branching points

Policy Based Multicast

- Policy based networking: Rules for non default routing
  - Geofencing
  - QoS
  - Membership filtering
- UB enables Policy Based Multicast
  - Number, location and type of branching nodes

Where are the Algorithms?

Building Efficient Policy Based Multicast Trees

- Problem 1 definition:
  - Given an ingress node, a set of egress nodes and a set of branching nodes, build an “optimal” multicast tree.
  - What is “optimal”
  - Usual definition is based on link cost.
  - Steiner tree problem (NP-complete)

Building Multicast Trees using UB – Major Issue

- UB based multicast tree is not a tree!!!
  - It is a “configuration”
- Cannot apply Steiner tree approximation solutions directly.
- Problem: How to create minimum cost configurations?

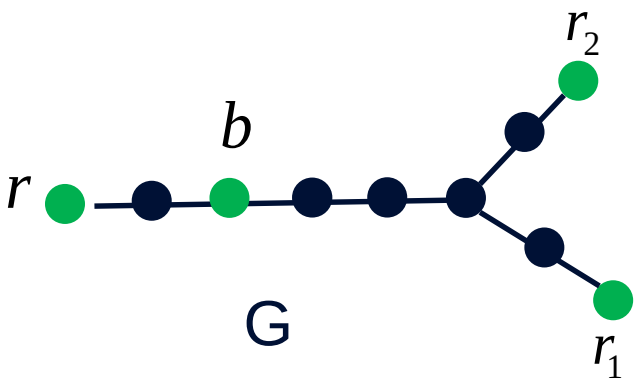

Transformation to Steiner tree problem on H

Define :

- Edge-weighted graph H = (O,E). O is set of branching nodes (including terminals)
- e=(b,b') ∈ O, w(e) = length shortest path containing no internal O nodes

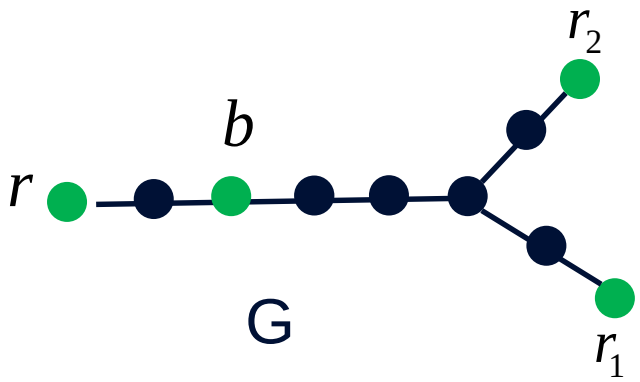

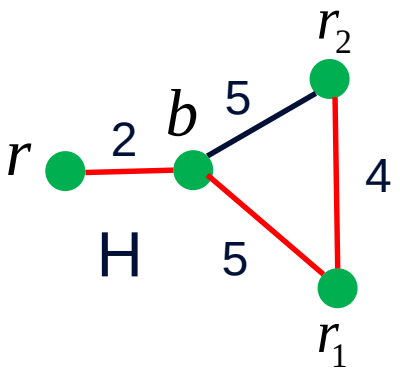

Theorem: Minimum cost configuration problem in G is equivalent to Steiner tree problem in H

Minimum cost configuration problem

Theorem: There is a polynomial-time 1.39-approximation algorithm for min cost configuration problem. [BGRS10]

Theorem: The minimum cost configuration problem is APX-hard.
Proof: Follows from APX-hardness of Steiner problem for complete graphs with weights 1 and 2. [BP89]

Problem 2: Minimize branching nodes

- Problem 1: Minimize cost given a set of branching nodes . MIN COST PROBLEM
- Problem 2: Minimize number of branching nodes given a fixed cost. MIN BRANCHING PROBLEM

Min Branching Problem

- For a subset X of the transit nodes, let C_X be the minimum cost valid configuration using X as the set of extra branching nodes.

- We are given a graph G = (V, E), a multicast demand d = (r, r1, r2,..., rt), a bound k and an attainable cost c.

- Does there exists a branching set X with least cost valid configuration C_X satisfying d where $|X| <= k$ and $cost(C_X) <= c$.

Does anybody

Min Branching Problem

Theorem: This problem is NP-complete.
Proof: Follows from a construction using Set Cover. ∎

Corollary: For this problem the best possible approximation is $\approx \ln n$.
Proof: Follows from bounds for Set Cover. ∎

Does anybody

Theorem: Min Branching is NP-complete

$demand = (r, r_1, \cdots, r_n)$

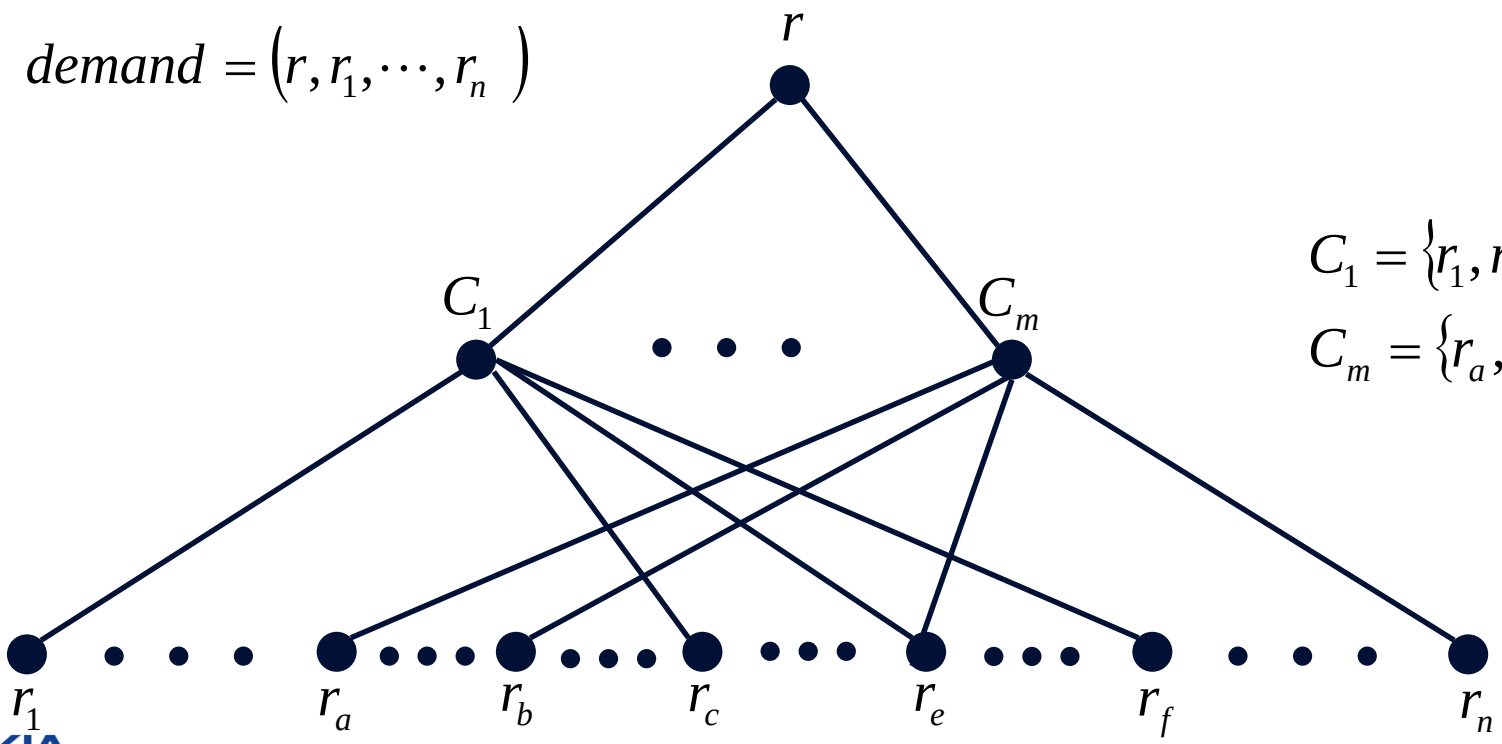

$C_1 = \{r_1, r_c, r_e, r_f\}$

$C_m = \{r_a, r_b, r_e, r_n\}$

Policy Driven Software Defined Multicast Using Efficient Selection of Unicast Branching Points

- Conclusion:
  - Unicast Branching based multicast provides for efficient, policy driven Software Defined Multicast.