

Projecting IPv6 Packet Forwarding Characteristics Under Internet-wide Deployment

Craig A. Shue and Minaxi Gupta
Computer Science Dept.
Indiana University at Bloomington
{cshue,minaxi}@cs.indiana.edu



Outline

- n Introduction
- n IP Packet Forwarding
- n IPv4 Baseline
- n IPv6 Analysis
 - Growth Projection
 - Partial Deployment
 - Impact of Deaggregation
- n Conclusion



Introduction


- n IPv6 could solve address exhaustion
- n Scalability of IPv6 warrants attention
 - Longer prefixes lead to bigger routing tables
 - Deaggregation is an increasing trend
- n Routers use a longest prefix match (LPM) to determine where to forward packets
- n We use popular LPM algorithms to determine scalability



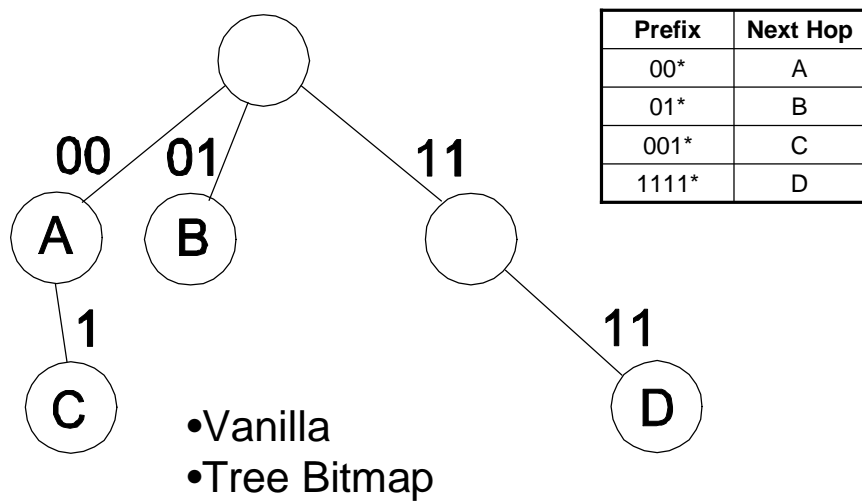
Longest Prefix Matching

- n Trie data structures frequently used for longest prefix matching
- n Optimizations performed on tries:
 - Multibit tries
 - Path compression
 - Hybrid (combines both of the above)


Traditional Trie



Multibit Trie

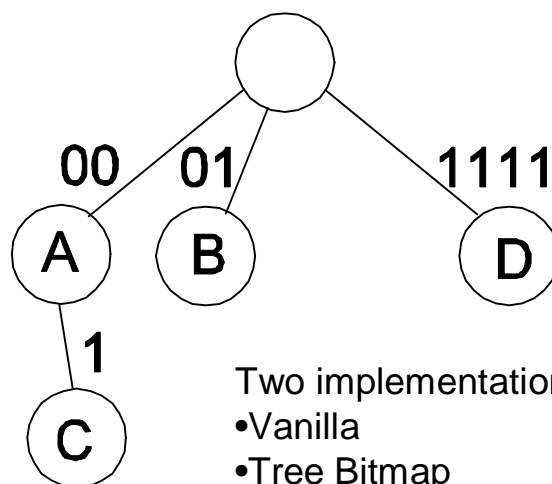


Path Compressed Trie



Prefix	Next Hop
00*	A
01*	B
001*	C
1111*	D

Hybrid Trie



Prefix	Next Hop
00*	A
01*	B
001*	C
1111*	D

Two implementations:

- Vanilla
- Tree Bitmap

IPv4 Baseline

- n We implement these LPM algorithms
- n Data from one router in Route Views Project
 - April 22, 2007
 - 233,500 prefixes
- n Evaluation Metrics
 - Routing Table Creation Time
 - Lookup Time
 - Update Time
 - Memory required to hold routing table

IPv4 Results

	Lookup Time	Memory Required
Traditional	2,643	19.364
Path Compressed	2,631	13.537
Multibit	1,779	27.826
Tree Bitmap	1,121	8.031
Hybrid	1,731	20.615
Hyrid (Tree Bitmap)	1,153	5.080

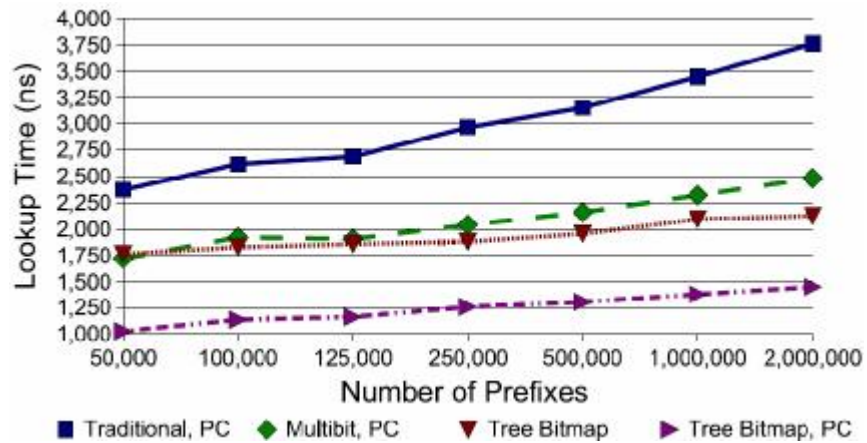
IPv6 Evaluation

- n Current deployment of IPv6 is limited
 - Insufficient data for reasonable comparison
- n Compare with current number of IPv4 prefixes
- n Project IPv6 routing entries into the future
- n Performance when IPv6 is partially deployed
- n Examine impact of deaggregation on IPv6


IPv6 Results

	Lookup Time		Memory Required	
	IPv4	IPv6	IPv4	IPv6
Traditional	2,643	5,296	19.364	88.238
Path Compressed	2,631	2,951	13.537	19.073
Multibit	1,779	3,411	27.826	109.857
Tree Bitmap	1,121	1,905	8.031	43.974
Hybrid	1,714	2,063	20.843	26.746
Hybrid (Tree Bitmap)	1,153	1,278	5.080	7.368

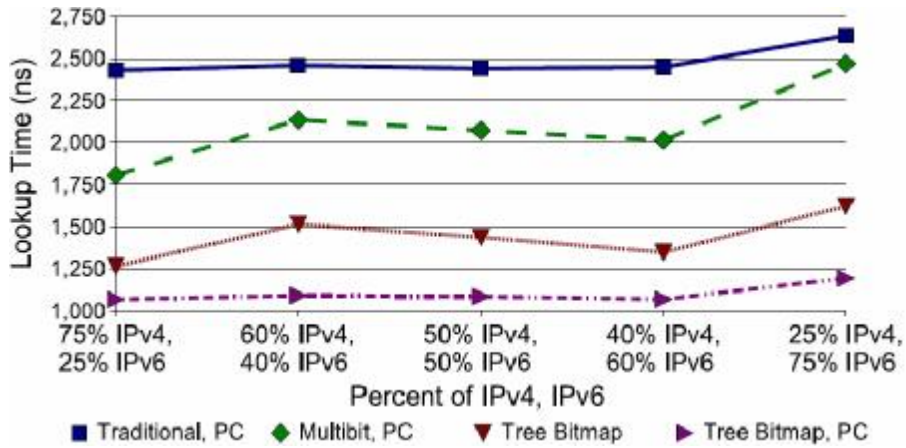
Projecting IPv6 – Lookup Times




Projecting IPv6 – Memory Usage



Partial Deployment – IPv4 and IPv6 Lookup Times



Partial Deployment – IPv4 and IPv6 Memory Requirements





Deaggregation Factors

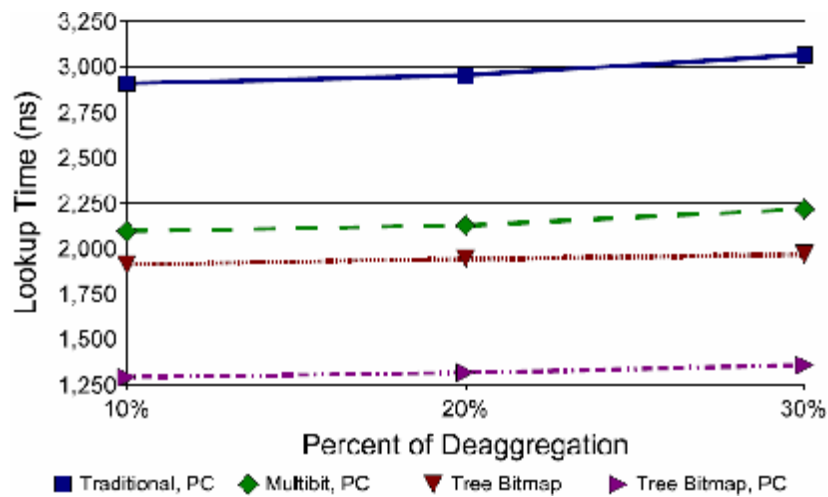
- n Load balancing
 - Prefixes split to announce routes differently
- n Multihoming
 - Subprefix created to announce via multiple routes
- n Failures to aggregate
 - Two prefixes can be aggregated, but for some reason are not



Simulating Deaggregation

- n Load balancing
 - Split an existing prefix in half
- n Multihoming
 - Randomly select subprefix of an existing prefix
- n Failure to Aggregate
 - Take existing prefix, but create new entry by toggling last bit

Impact of Deaggregation on IPv6



Conclusions

- n Using today's forwarding algorithms, IPv6 would take:
 - 67% longer to perform lookups
 - 4.5 times as much memory
- n Path compressed Tree Bitmaps are best
 - Scalable: sub-linear growth in lookup times, but linear growth in memory requirements
- n Mileage may vary: projections of future growth cannot account for unanticipated events

