

IP Reachability Differences: Myths and Realities

Dr. Dan Massey
Colorado State University

He Yan, Ben Say, Dave Oko, Brendan Sheridan
Colorado State University

Dr. Christos Papadopoulos
Colorado State University

Dr. Dan Pei
AT&T Research

Ideally, all ISPs provide you with reachability to same Internet

In practice: Different ISPs can reach different portions of the Internet. In other words, there are some prefixes that **can be reached via ISP A** but **cannot be reached by ISP B** (a vice versa).

The reasons why this occur are well known:
Bogon filtering causes reachability differences for newly allocated prefixes.

Bogon filtering matters, but the fact is that regional differences matter most and these differences may be political, economic, technical, or a combination of the above

Bogon filtering and regional issues matter, but the real difference is due to various customer only services.

Even if there are differences, there are very little (if any) “production” services on networks that are not globally reachable from most ISPs.

Our Objectives

1. Demonstrate reachability differences exist and quantify the extent of these differences.
2. Examine which conjectures (if any) provide a significant contribution to the differences.


Methodology

- Compare BGP Routing Tables from 25 ISPs
 - Data collected from Oregon RouteViews
 - All peers send full tables of over 310K prefixes
 - Studied from April 1 to April 20, 2010
- Study Uses ISPO As a Central Point of View
 - ISPO is a tier-1 ISP
 - Have additional data on ISPO operations
 - Compare each of the other 24 ISPs to ISPO

Types of Table Differences


- **Exact Matches:** ISP0 prefix appears ISPX table
 - Ex: Both ISP0 and ISP2 have a route for 129.82/16
- **Covered Prefixes:** ISP0 prefix does not appear in ISPX, but ISPX can reach all IP addresses
 - Less Specific: 131.179.0/17 does not appear in ISPX, but 131.179/16 does appear in ISPX
 - More specifics: 131.179.0/17 does not appear in ISPX, but 131.179.0/18 and 131.179.64/18 appear in ISP X
- **Remainder represents potential reachability differences.**

Exact Matches

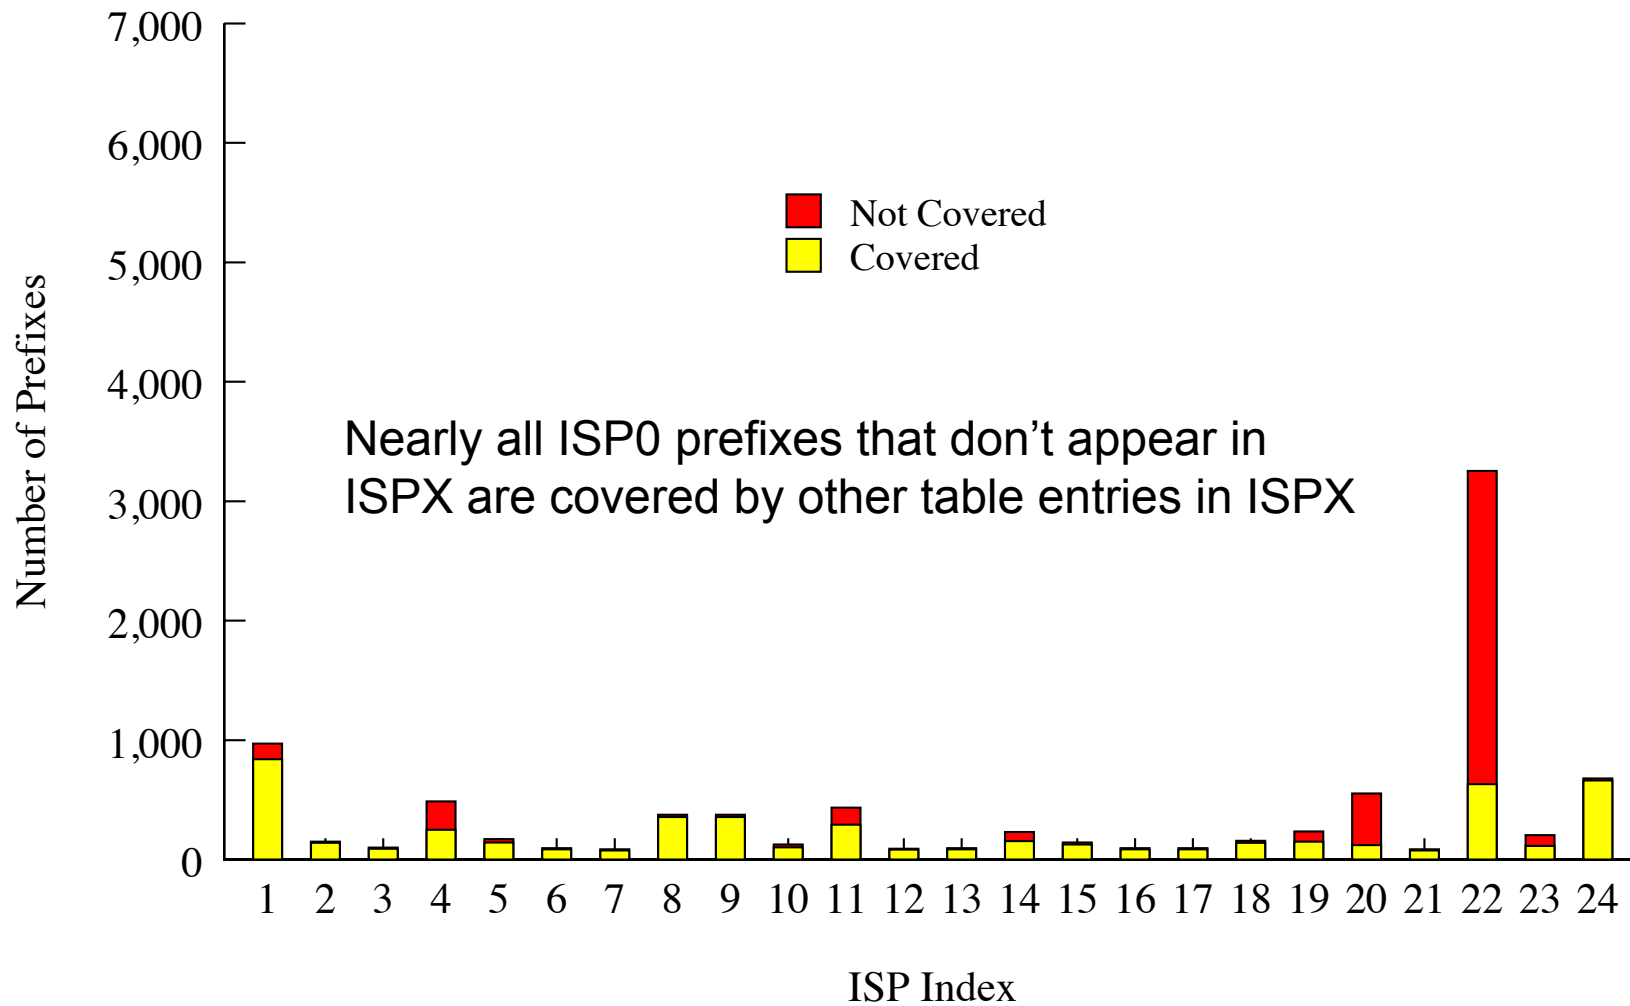


(not shown here, but 97.25% of ISPX prefixes also appear in ISP0)

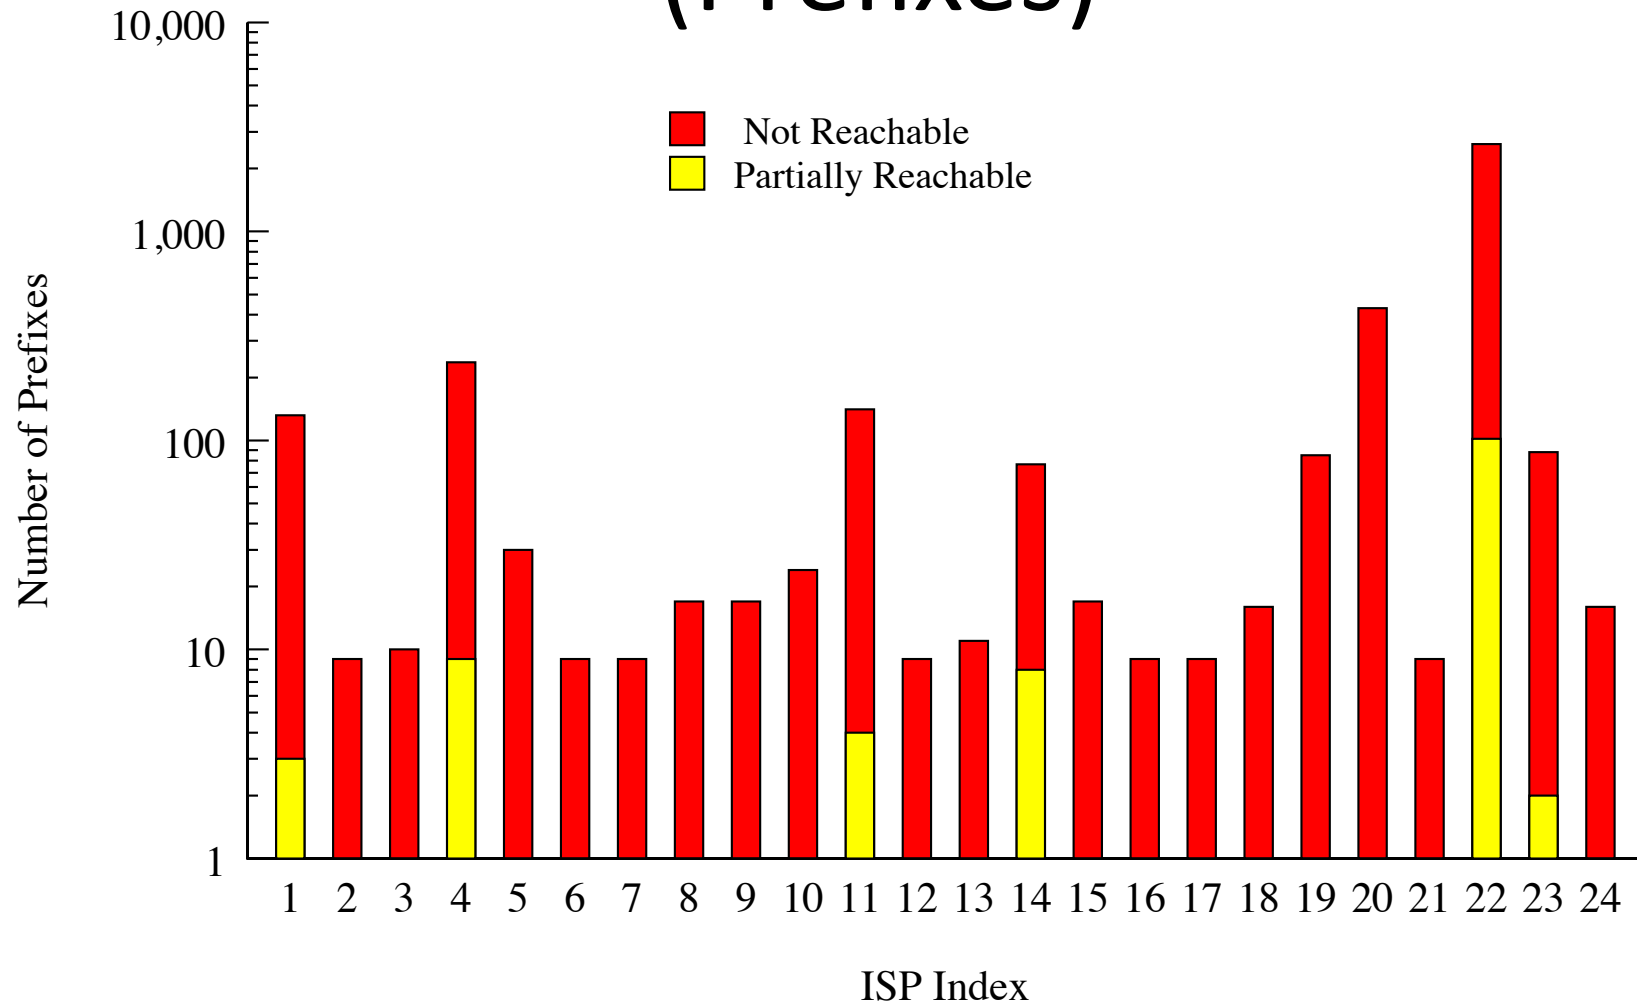
The BGP Table Differences




Covered Prefixes



Overall Reachability Differences (Prefixes)



Overall Reachability Differences (IP Addresses)



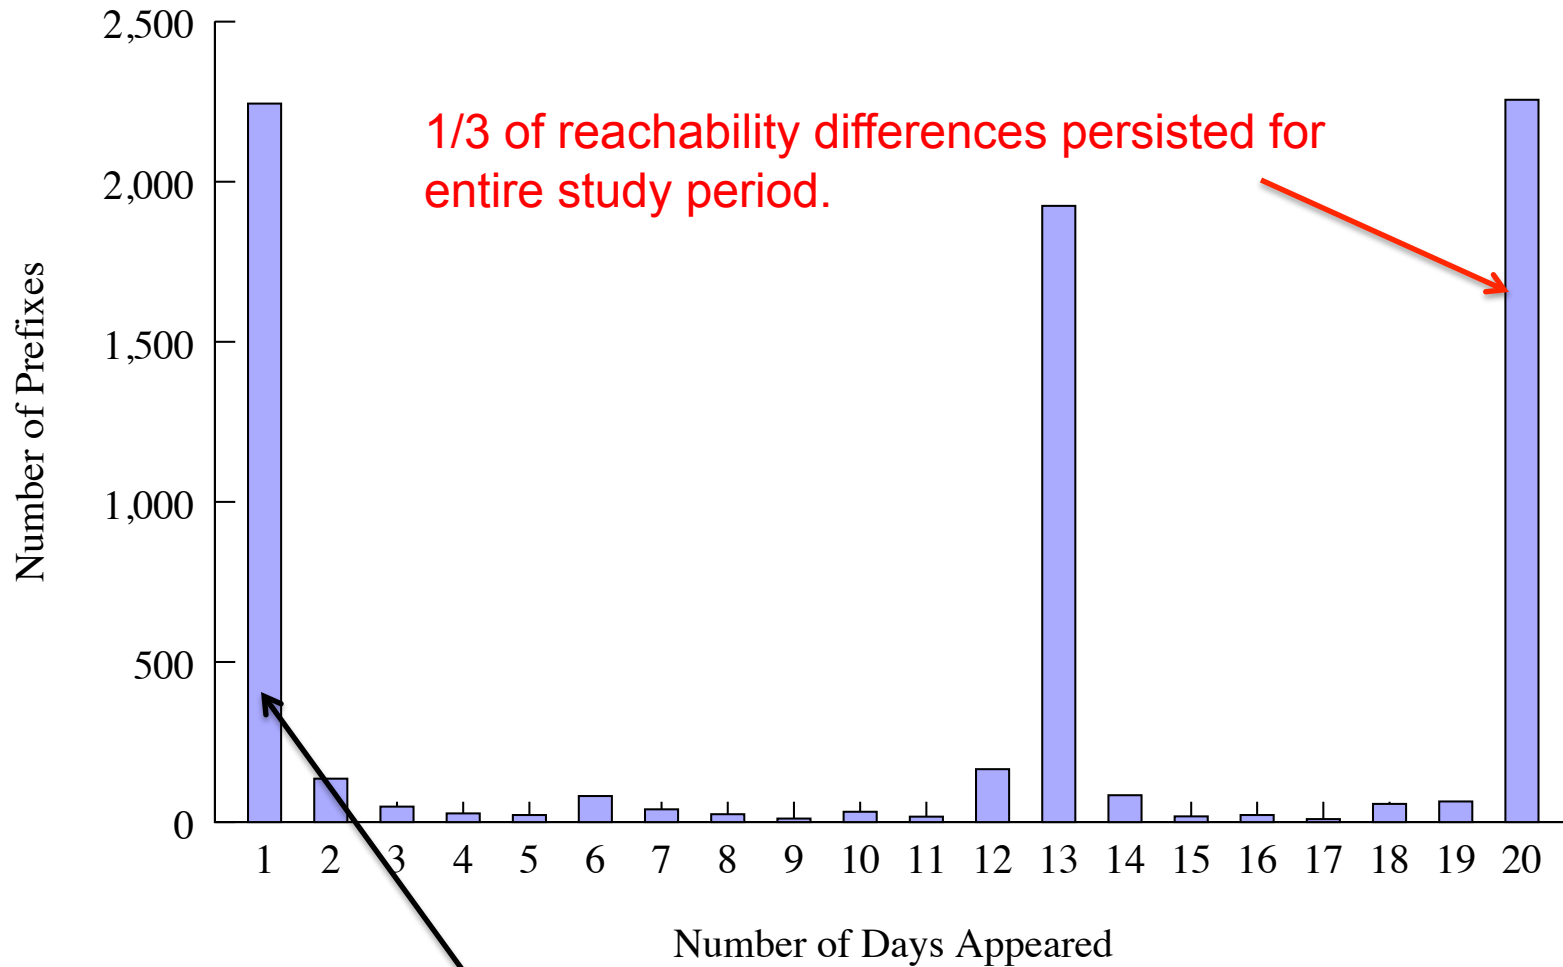
Our Objectives

1. Reachability differences do exist.
 - Small percentage (less than 3%) of total routes.
 - Can be up to 1000 prefixes and millions of IP addresses
2. Examine which conjectures (if any) provide a significant contribution to the differences.

Transient Changes

- Short Lived BGP Table Differences
 - Fast learners: IPSO learns the route first, eventually other ISPs catch up and also learn the route
 - Slow to react: ISPX learned the prefix is unreachable, but it takes longer for ISPO to withdraw the route.
- Long Term BGP Table Differences.
 - For multiple days, ISPO could reach the prefix but ISPX could not.

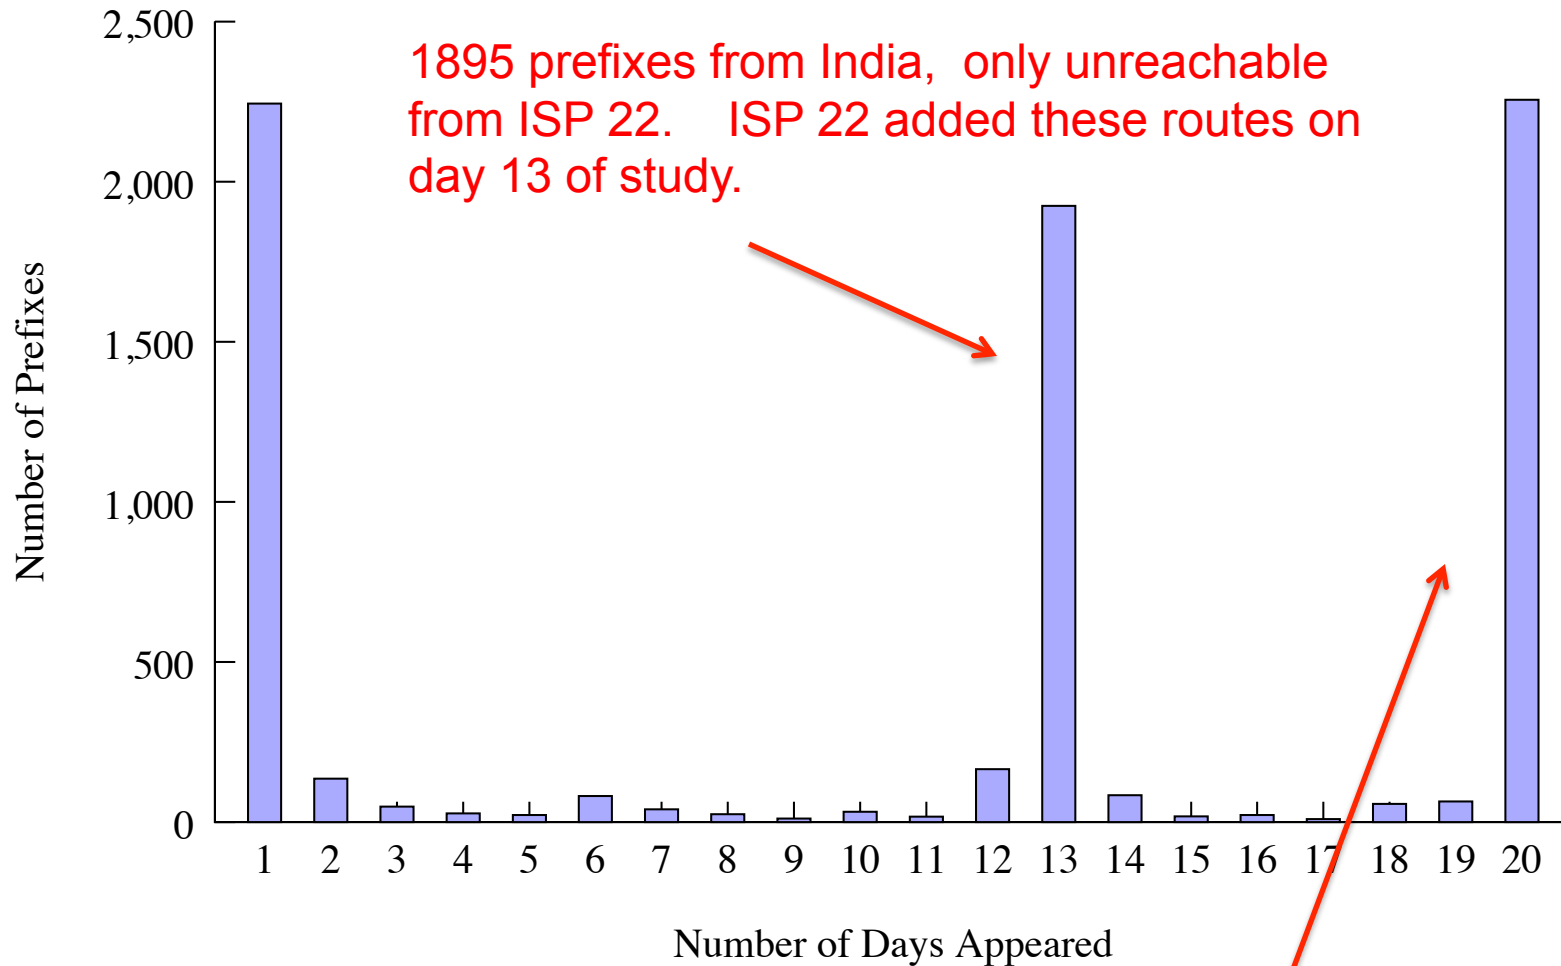
Transient Behaviors



Bogon Routes

- ISPs Apply Policies to Filter Invalid Routes
 - No universal agreement on what is “invalid”
- New Allocations and Bogons
 - Filter unallocated prefixes and update filter after allocation
 - Going back 7 months accounts for **43 prefixes**.
 - Going back years yields **295 prefixes**.
- Management Strategies and RADb
 - **92 prefixes** with no entry in RADb
 - **199 prefixes** with no origin in RADb.
- Does produce reachability differences, but only explains a small percent of differences.

Regional Differences



1895 prefixes from India, only unreachable from ISP 22. ISP 22 added these routes on day 13 of study.


Bulgaria accounted for most prefixes (509)
Great Britain was second highest (168)

Customer Services and Prefix Usage

- ISPs Provide Specialized Services to Customers
 - Confirmed ISPO has several such prefixes
 - Most are covered prefixes, but did find some uncovered prefixes
 - Accounts for **9 prefixes** in our study.
- Are Services Running on the Impacted Prefixes?
 - Scanned prefixes looking for common services
 - Limited scan to /19 or more specifics
 - Responding prefixes further checked for DNS entries
- Contacted Prefix Owners
 - Responses include reports of reachability problems, blacklist issues, network testing, and so forth
- **Valid Global Services Running on Clear Majority of Prefixes**

RouteViews Data in Real-Time

Active Prefixes



Conclusions

- Reachability Issues Do Exist
 - Over 97% of prefixes appeared as exact matches
 - But differences do occur and impact both research and operations
- Common Conjectures:
 - Transient issues, Bogons, Management, Regional Policies, Configuration Errors, Special Services all proposed as explanations
 - All are valid, but none are dominant.

Questions

massey@cs.colostate.edu