BGP update profiles and the implications for secure BGP update validation processing

Geoff Huston
APNIC

7th caida/wide measurement workshop
Nov 3-4 2006

Why?

- Secure BGP proposals all rely on some form of validation of BGP update messages
- Validation typically involves cryptographic validation, and may refer to further validation via a number resource PKI
- This validation may take considerable resources to complete.
- This implies that the overheads securing BGP updates in terms of validity of payload may contribute to:
  - Slower BGP processing
  - Slower propagation of BGP updates
  - Slower BGP convergence following withdrawal
  - Greater route instability
  - Potential implications in the stability of the forwarding plane

What is the question here?

- Validation information has some time span
  - Validation outcomes can be assumed to be valid for a period of hours
- Should BGP-related validation outcomes be locally cached?
- What size and cache lifetime would yield high hit rates for BGP update validation processing?

Method

- Use a BGP update log from a single *e*BGP peering session with AS 4637 over a 14 day period
  - 10 September 2006 – 23 September 2006
- Examine time and space distributions of BGP Updates that have similar properties in terms of validation tasks

Update Statistics for the session

Day	Prefix Updates	Duplicates: Prefix	Duplicates: Prefix + Origin AS	Duplicates Prefix + AS Path	Duplicates Prefix + Comp-Path
1	72,934	60,105 (82%)	54,924 (75%)	34,822 (48%)	35,312 (48%)
2	79,361	71,714 (90%)	67,942 (86%)	49,290 (62%)	50,974 (64%)
3	104,764	93,708 (89%)	87,835 (84%)	65,510 (63%)	66,789 (64%)
4	107,576	94,127 (87%)	87,275 (81%)	64,335 (60%)	66,487 (62%)
5	139,483	110,994 (80%)	99,171 (71%)	68,096 (49%)	69,886 (50%)
6	100,444	92,944 (92%)	88,765 (88%)	70,759 (70%)	72,108 (72%)
7	75,519	71,935 (95%)	69,383 (92%)	56,743 (75%)	58,212 (77%)
8	64,010	60,642 (95%)	57,767 (90%)	49,151 (77%)	49,807 (78%)
9	94,944	89,777 (95%)	86,517 (91%)	71,118 (75%)	72,087 (76%)
10	81,576	78,245 (96%)	75,529 (93%)	63,607 (78%)	64,696 (79%)
11	95,062	91,144 (96%)	87,486 (92%)	72,678 (76%)	74,226 (78%)
12	108,987	103,463 (95%)	99,662 (91%)	80,720 (74%)	82,290 (76%)
13	91,732	87,998 (96%)	85,030 (93%)	72,660 (79%)	74,116 (81%)
14	78,407	76,174 (97%)	74,035 (94%)	64,994 (83%)	65,509 (84%)

CDF by Prefix and Originating AS

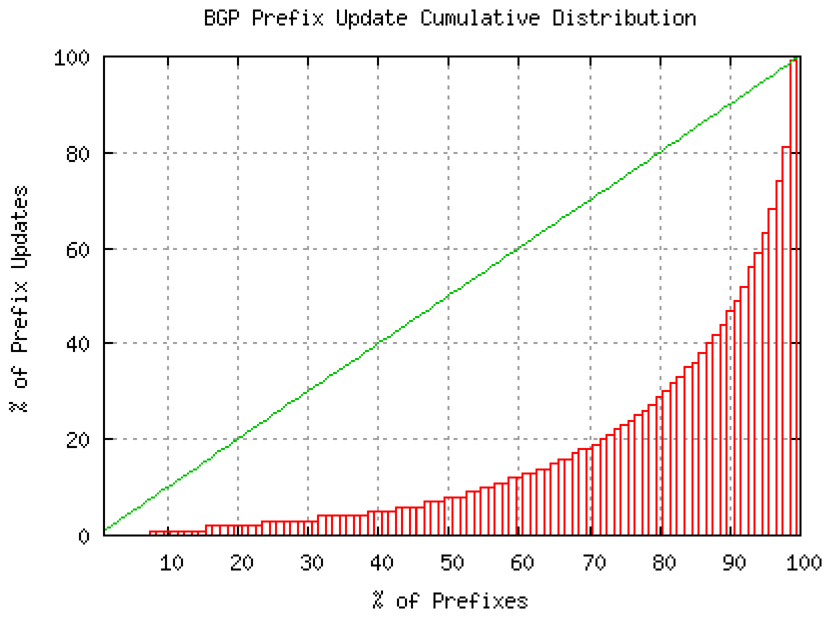

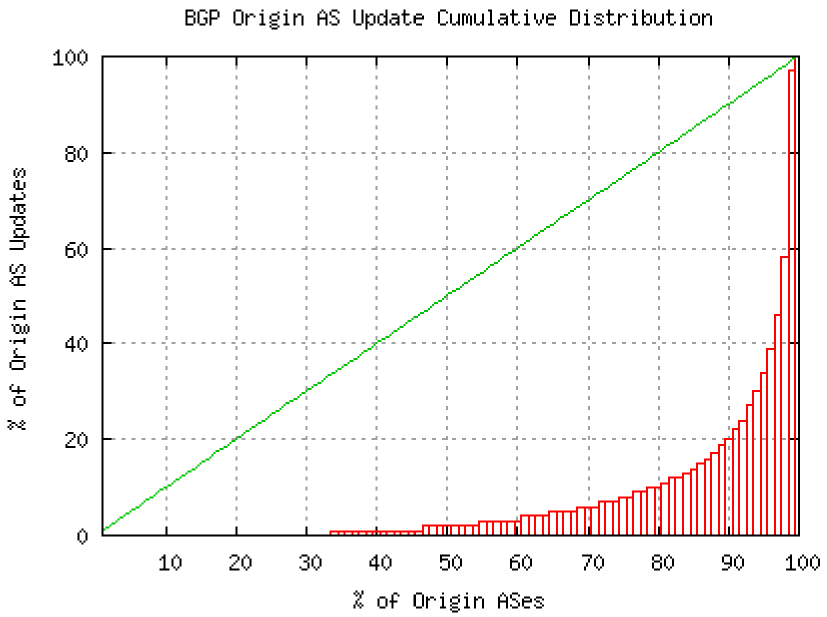

Time Distribution

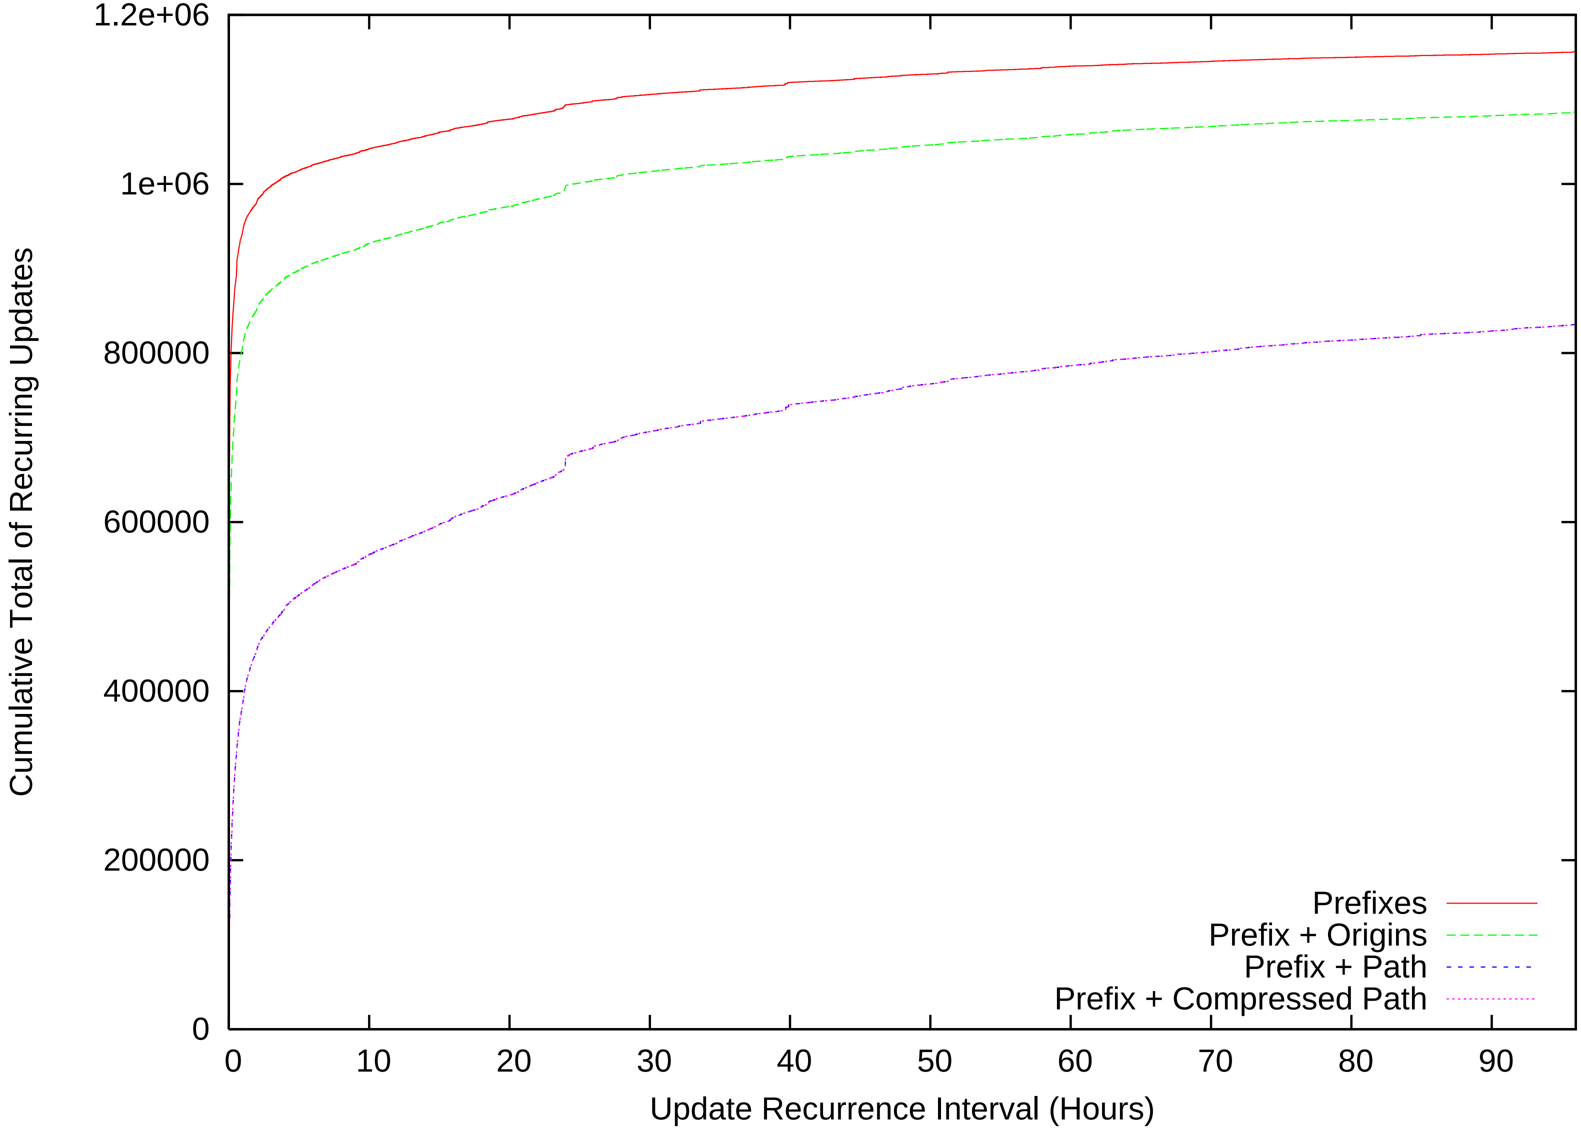

Time Spread

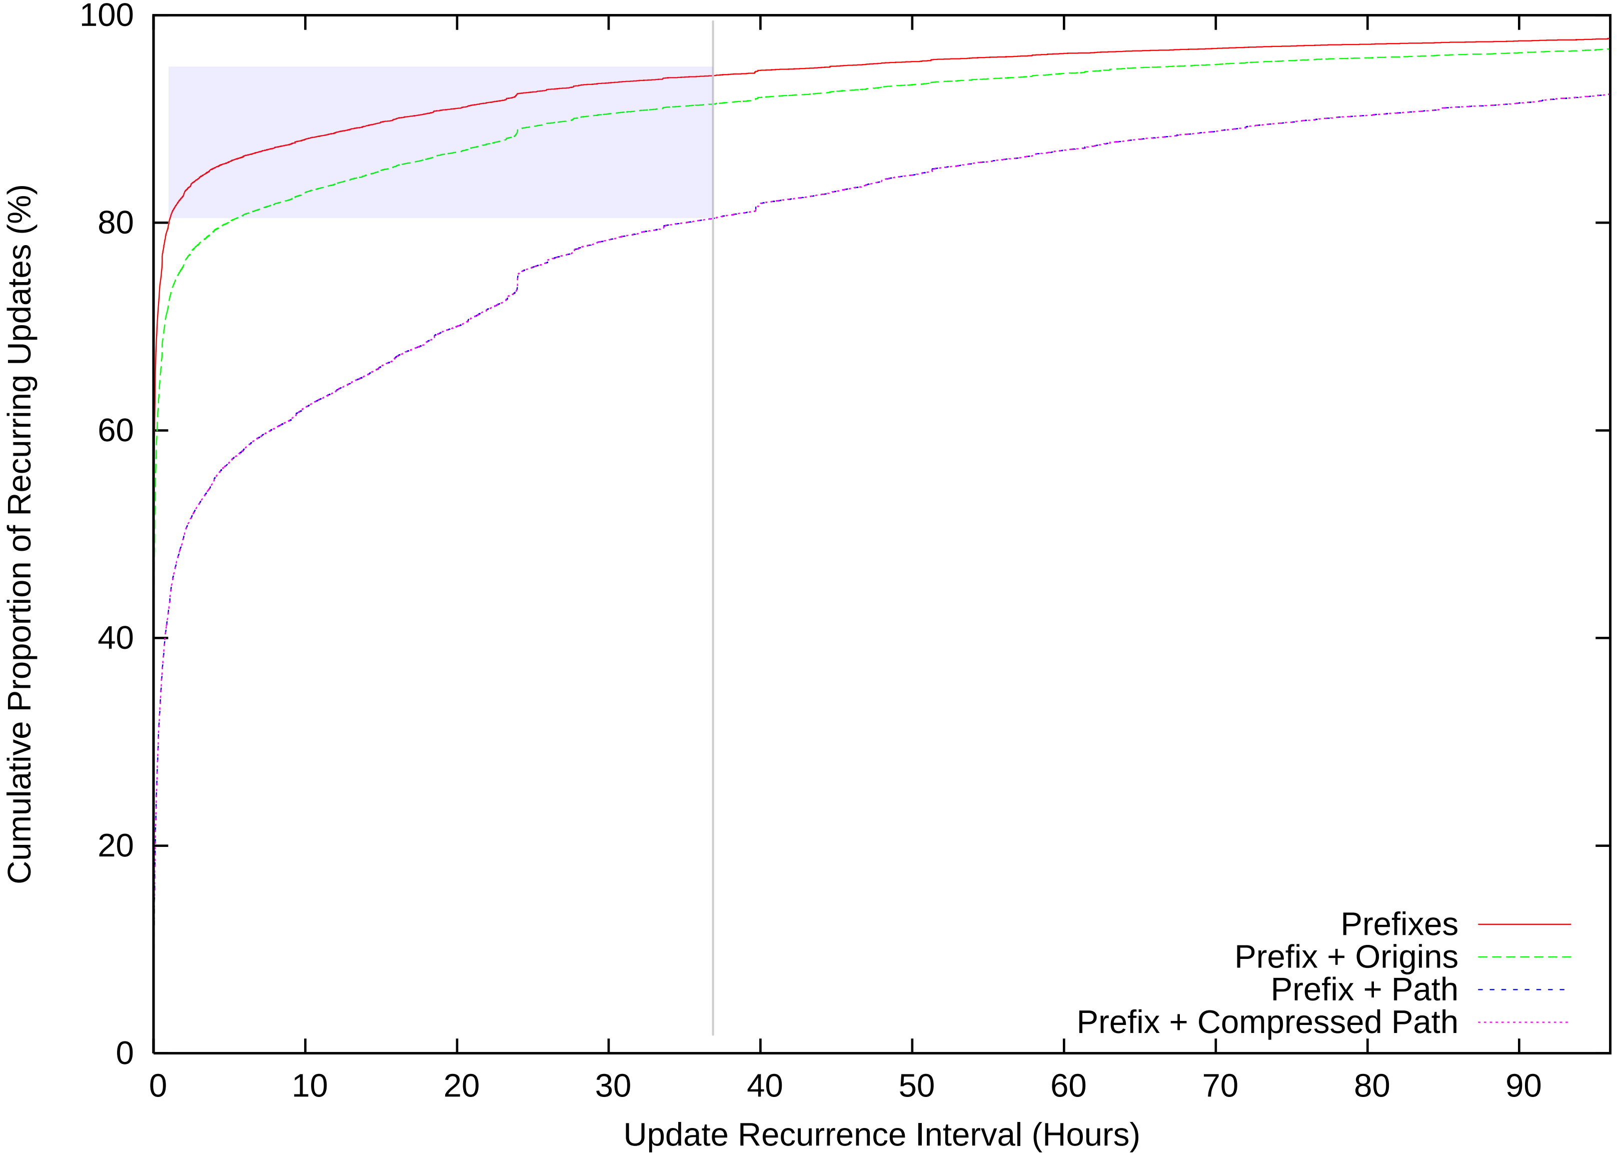
100
80
60
40
20
0
Cumulative Proportion of Recurring Updates (%)
0
10
20
30
40
50
60
70
80
90
Update Recurrence Interval (Hours)
Prefixes
Prefix + Origins
Prefix + Path
Prefix + Compressed Path

Space Distribution

- Use a variable size cache simulator
- Assume 36 hour cache lifetime
- Want to know the hit rate of validation queries against cache size

Prefix Similarity

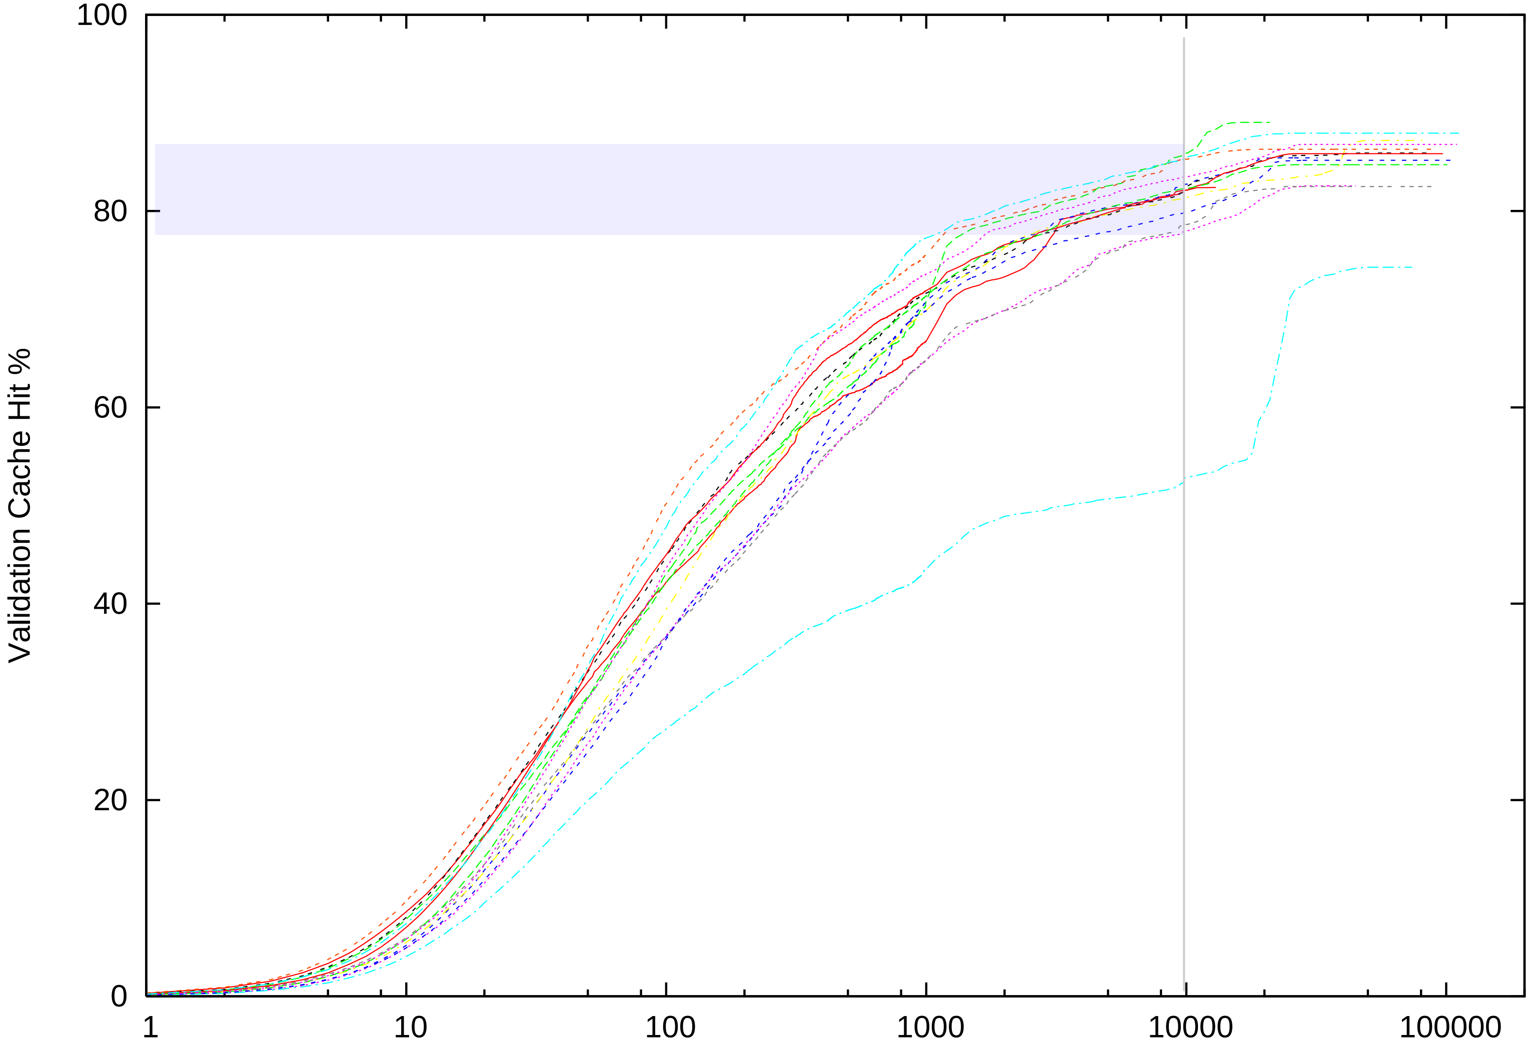

Prefix + Origin Similarity

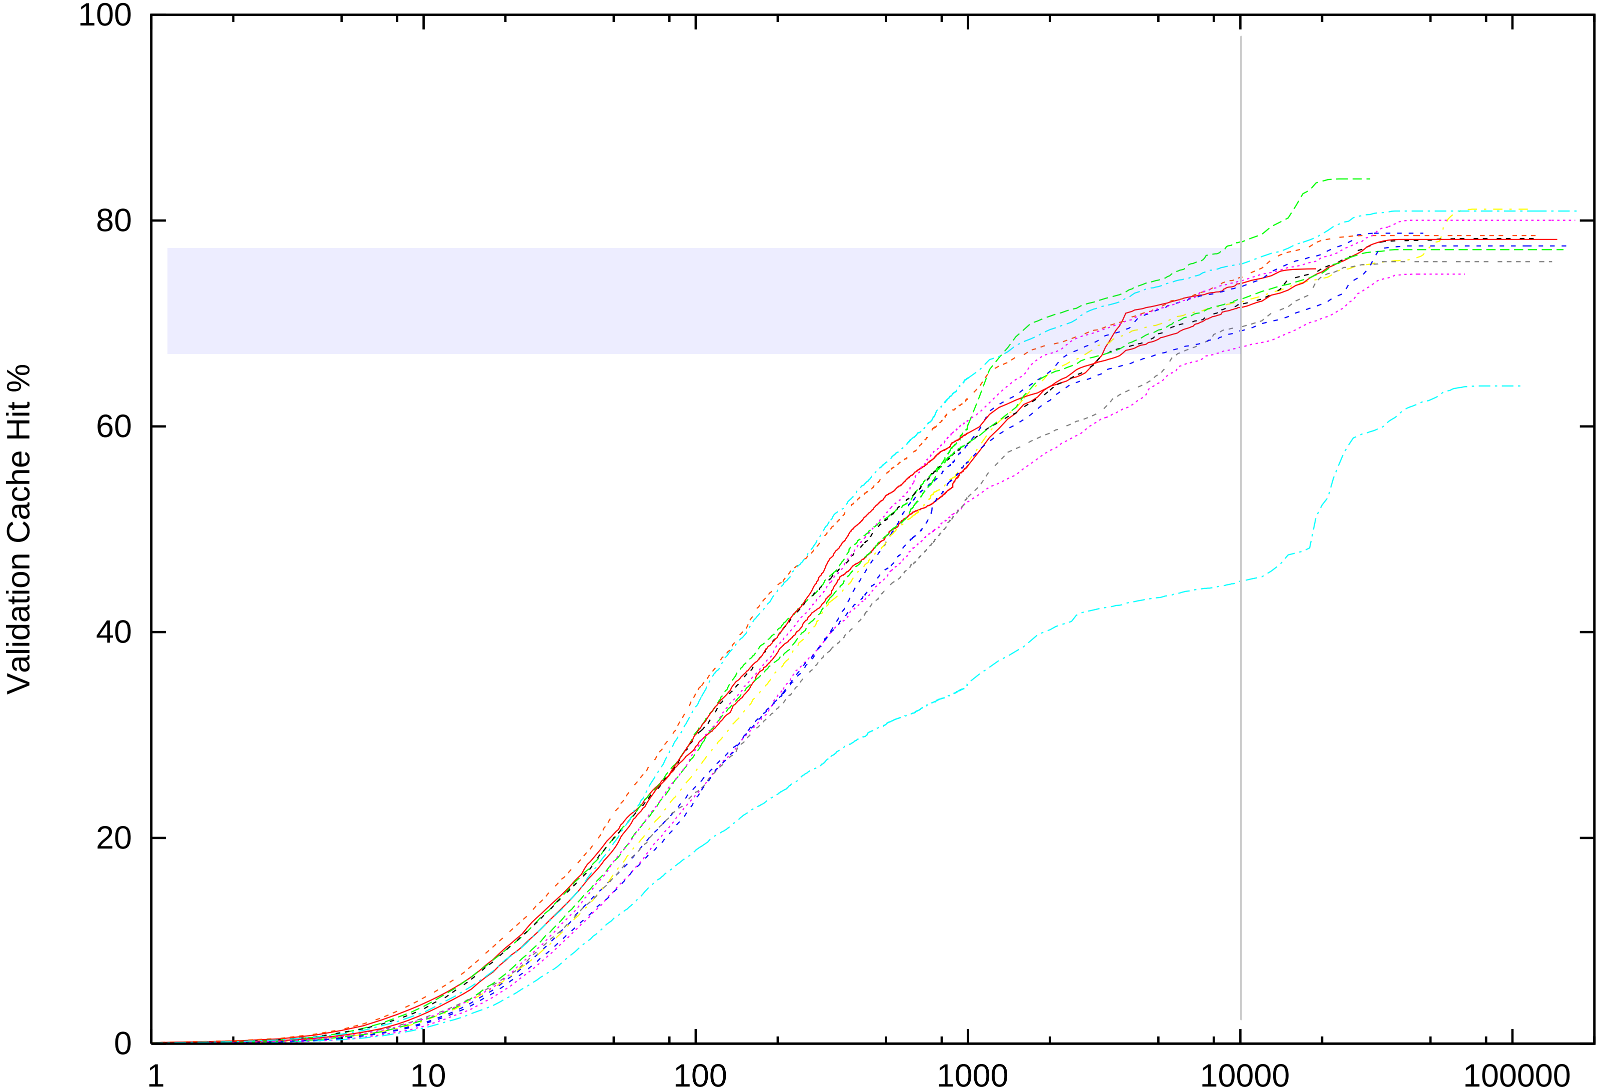

Prefix + Path Similarity

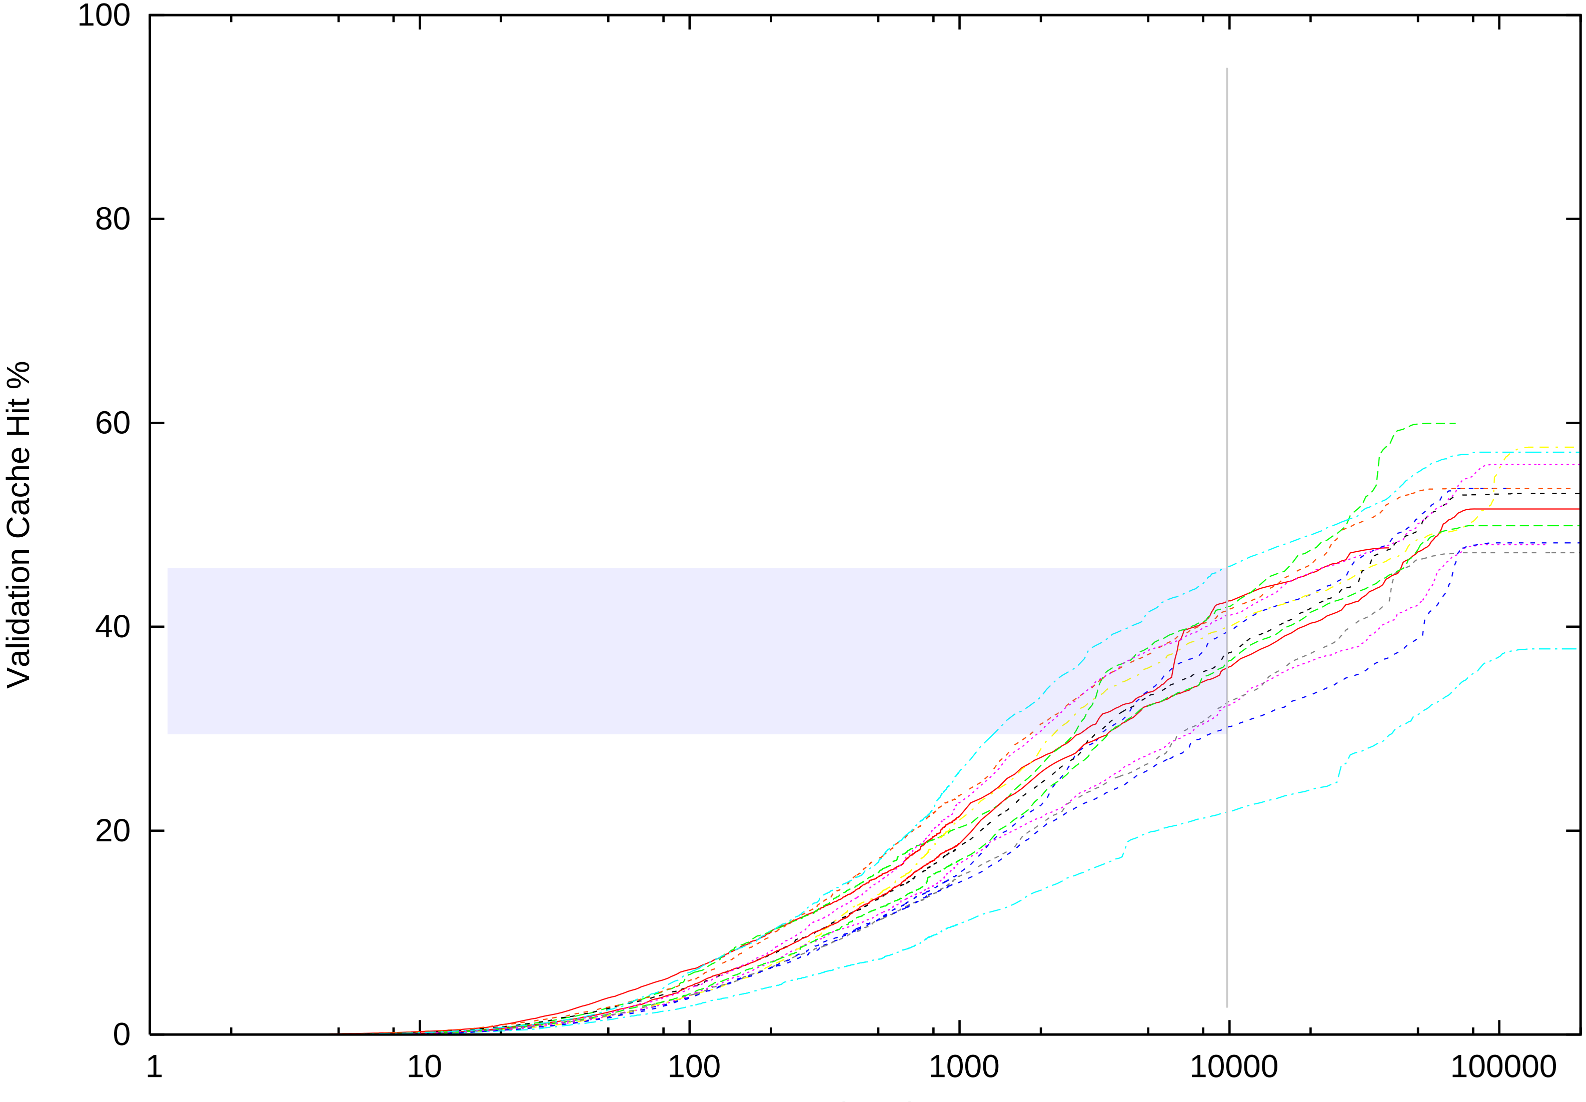

Observations

- A large majority of BGP updates explore diverse paths for the same origination
- True origination instability occurs relatively infrequently (1:4) ?
- Validation workloads can be reduced by considering origination (prefix plus origin) and the path vector as separable validation tasks
- Further processing reduction can be achieved by treating a AS path vector as a sequence of AS paired adjacencies

AS Path Similarity

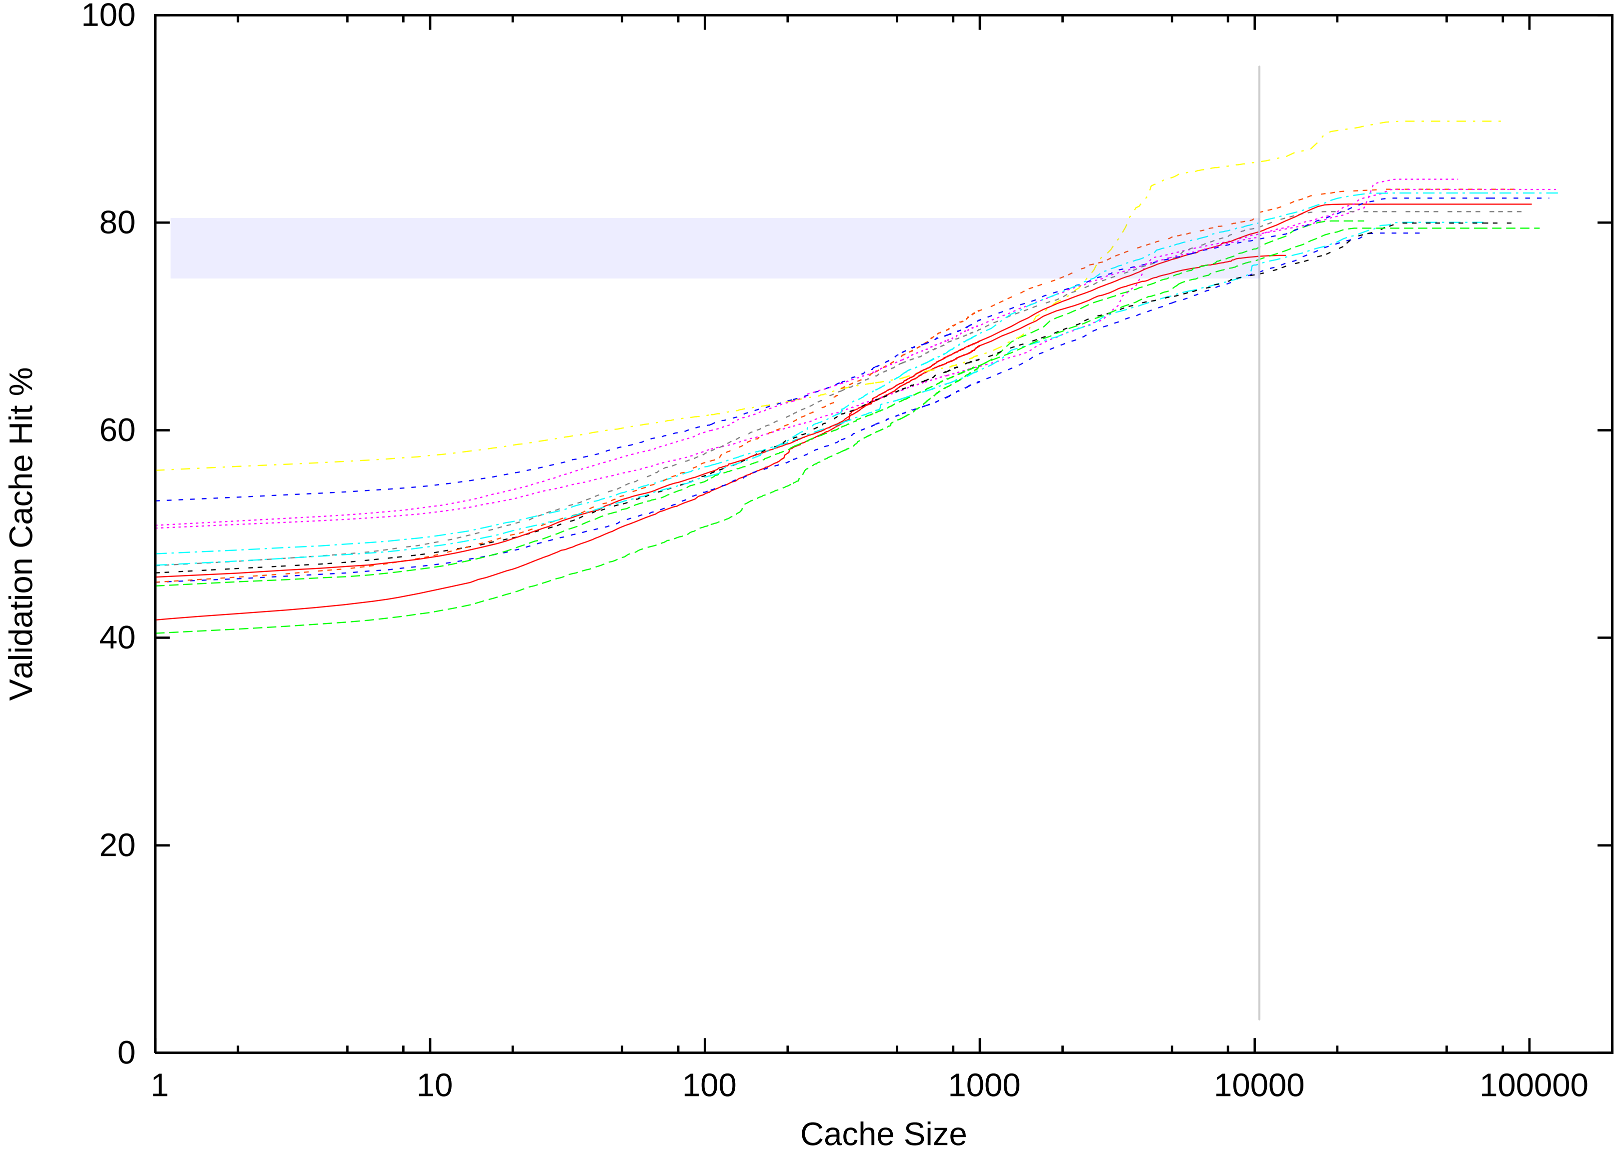

AS Pair Similarity

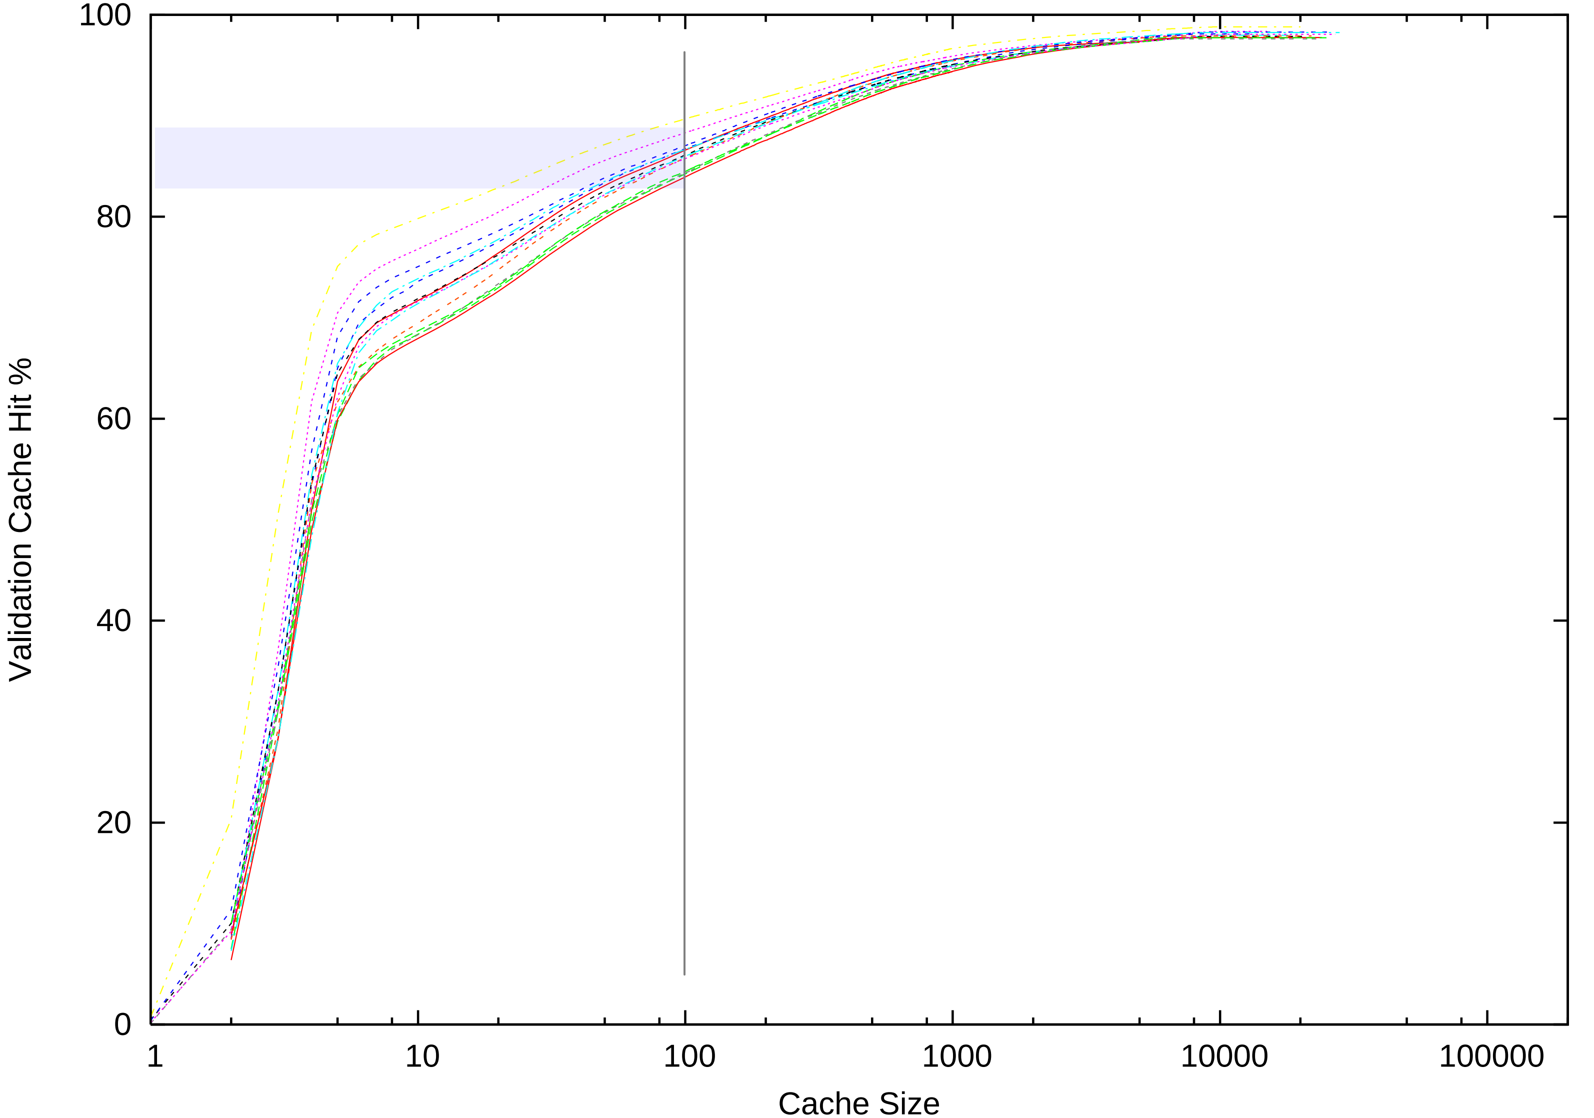

Observations

- Validation caching appears to be a useful approach to addressing some of the potential overheads of validation of BGP updates
- Separating origination from path processing, using a 36 hour validation cache can achieve 80% validation hit rate using a cache of 10,000 Prefix + AS originations and a cache of 1,000 AS pairs

What do we want from secure BGP?

- Validation that the received BGP Update has been processed by the ASs in the AS Path, in the same order as the AS Path, and reflects a valid prefix, valid origination and valid propagation along the AS Path?

or

- Validation that the received Update reflects a valid prefix and valid origination, and that the AS Path represents a plausible sequence of validated AS peerings?

Further work?

- Heaps!

Thanks

Questions?