NIST National Institute of Standards and Technology

INFORMATION TECHNOLOGY LABORATORY

Trustworthy Networking Program

A Comparative Analysis of BGP Anomaly Detection and Robustness Algorithms

Kotikapaludi Sriram, Oliver Borchert, Okhee Kim,

Patrick Gleichmann, and Doug Montgomery

National Institute of Standards and Technology

(Contact: ksriram@nist.gov; dougm@nist.gov)

NANOG-45, January 2009

This research was supported by the Department of Homeland Security under the Secure Protocols for the Routing Infrastructure (SPRI) program and the NIST Information Technology Laboratory Cyber and Network Security Program.

How this paper differs from the what we presented at NANOG-43

- The analysis presented at NANOG-43 was preliminary and was limited to RIPE IRR/RIR only
- This presentation will cover more detailed analysis with extensive data from Global IRRs/RIRs/RADB and RIPE-RIS
- Several significant additional insights
- Full paper available including detailed results and discussion

http://www.antd.nist.gov/~ksriram/NIST_BGP_Robustness.pdf

Outline of the Talk

- **Known / New BGP robustness schemes**
- **Evaluation of BGP robustness algorithms**
  - **Qualitative / comparative analysis of utility**
  - **Quantitative results**
- **Conclusions / Future Work**

"Blueprint" / Nemecis: Registry Based Algorithm

- For (p, Origin AS) pair from an update:
  - Check for existence of prefix, autnum, and route objects in RIR/IRR
  - Check for consistency between these declared objects by matching Organization, maintainer, email, etc.
  - Generate alerts if these checks fail -- full / partial consistency checks

G. Siganos and M. Faloutsos, "A Blueprint for Improving the Robustness of Internet Routing," 2005. http://www.cs.ucr.edu/%7Esiganos/papers/security06.pdf

G. Siganos and M. Faloutsos, "Analyzing BGP policies: methodology and tool," IEEE Infocom, 2004. http://www.cs.ucr.edu/~siganos/papers/Nemecis.pdf

PHAS: Prefix Hijack Alert System

- Provide alert messages if:
  - Origin AS set changes
  - New subprefix is added to observed set of subprefixes
  - Last-hop AS set changes

Mohit Lad, Dan Massey, Yiguo Wu, Beichuan Zhang and Lixia Zhang, *PHAS: A prefix hijack alert system*, North American Network Operators Group Meeting (NANOG-38), October, 2006. http://www.nanog.org/mtg-0610/presenter-pdfs/massey.pdf

Mohit Lad, Dan Massey, Dan Pei, Yiguo Wu, Beichuan Zhang and Lixia Zhang, *PHAS: A prefix hijack alert system*, in Proceedings of 15th USENIX Security Symposium (USENIX Security 2006). http://www.cs.ucla.edu/~mohit/cameraReady/ladSecurity06.pdf

PGBGP: Pretty Good BGP

Old Version of the Algorithm

- Observed {prefix, Origin AS} pairs based on update history and RIB entries over the last *h* days (*h* = 10 days) are recorded
- An update for a prefix is considered suspicious if the origin AS is new relative to the history record; the update is propagated with lower local pref
- A new subprefix (of a prefix in history record) is always considered suspicious and quarantined
- The quarantine lasts for suspicious period of *s* hours (*s* = 24 hours); if the subprefix is not withdrawn during that time, then the update is propagated

One Weakness of Old PGBGP

From NANOG discussions back in 2006

Q: Panix's first, obvious countermeasure aimed at restoring their connectivity – announcing subprefixes of their own address space – would also have been considered suspicious, since it gave two "sub-prefixes" of what ConEd was hijacking?

A: [Here] things get a little more subtle. We have considered allowing the trusted originator of a prefix to split the space among itself and those downstream of it without considering that suspicious behavior.

Note: This was part of the Q&A after the paper on PGBGP was presented by J. Karlin at NANOG-37. http://www.nanog.org/mtg-0606/pdf/josh-karlin.pdf

New Version of PGBGP

- From an updated new version of PGBGP paper:
  - "PGBGP would not interfere if an AS announces sub-prefixes of its own prefixes in order to gain traffic back during a prefix hijack."

Josh Karlin, Stephanie Forrest, and Jennifer Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," The 14th IEEE International Conference on Network Protocols, November 2006. http://www.cs.unm.edu/~treport/tr/06-06/pgbgp3.pdf

Potential Weaknesses of (New) PGBGP

- The short-span historical view (last ten days) has the following negative implications:
  - PGBGP will typically unnecessarily lower local-pref on path announcements due to multi-homing related AS origin change.
  - If a malicious user observes a prefix withdrawal by genuine origin AS and announces the prefix at that time, the malicious path propagates with a lower local-pref value and will be used *(*Effectively - *False Negative).*
  - If the prefix owner sometimes announces sub-prefixes in conjunction with multi-homing related AS origin change, PGBGP will quarantine the announcements.

New Integrated Approach

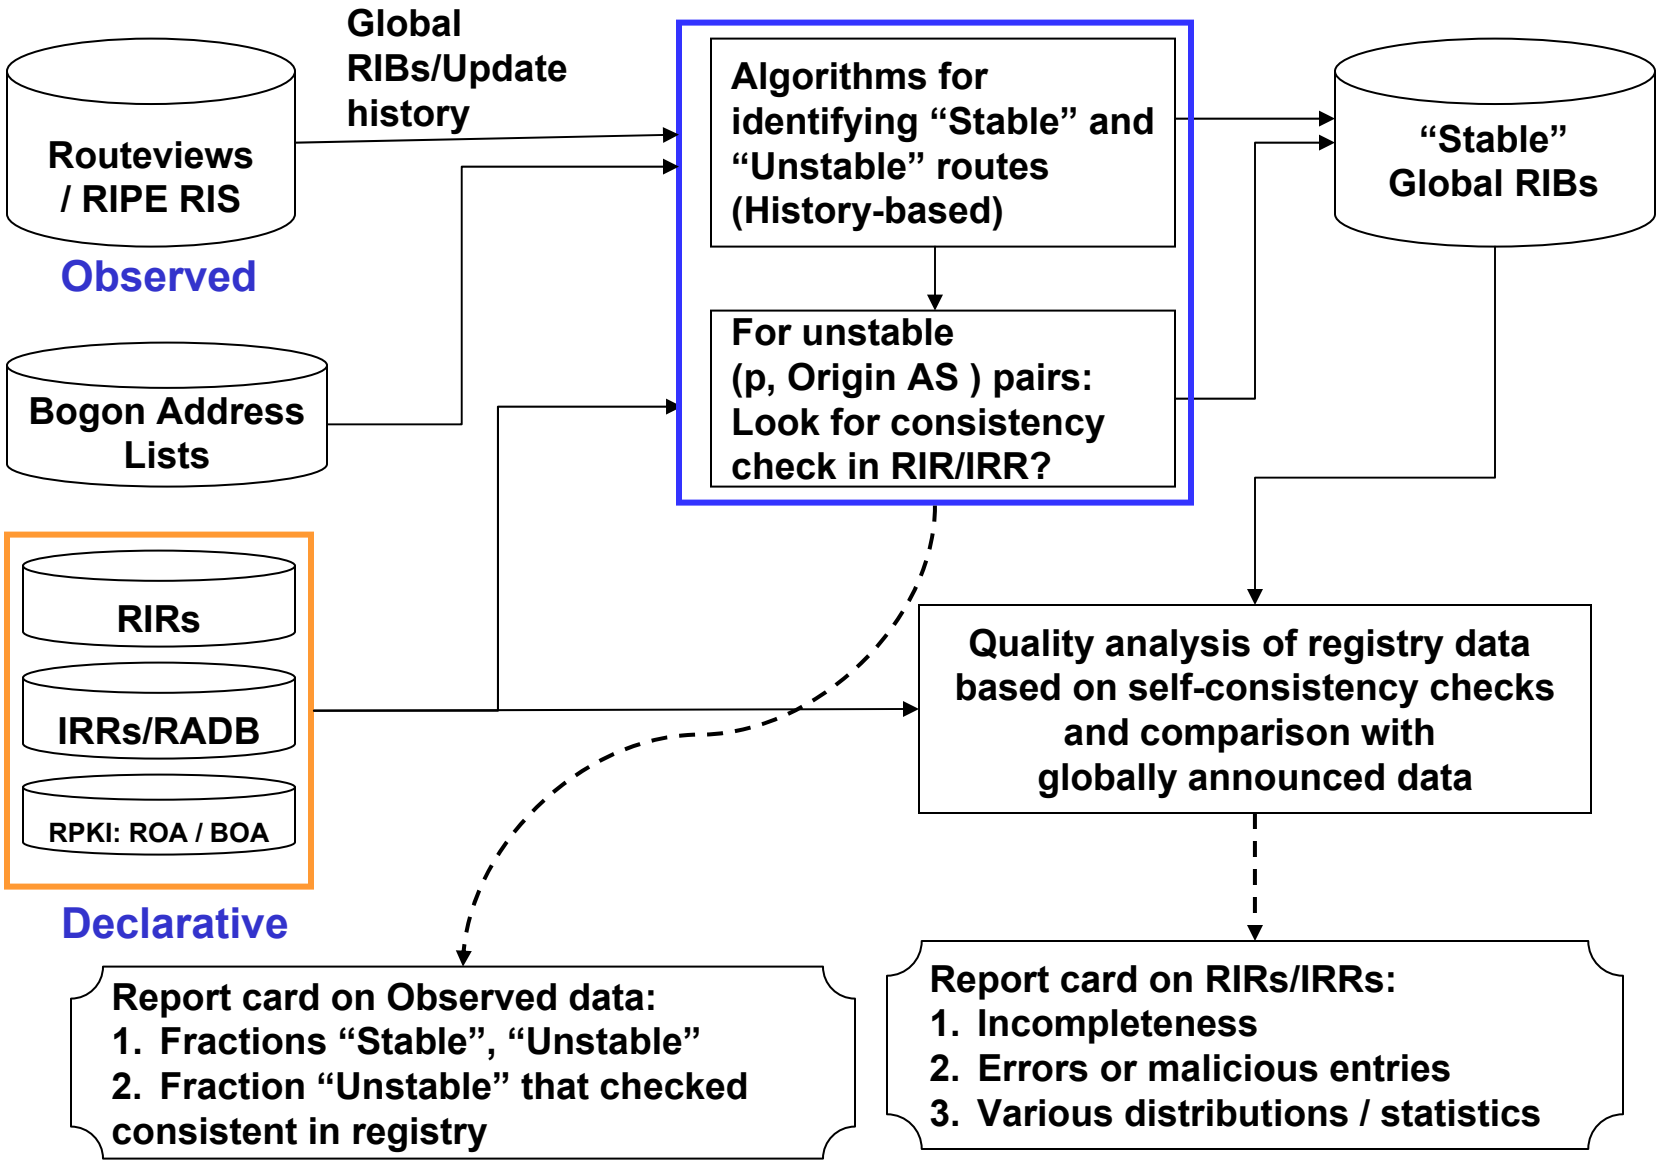

ROA: Route Origin Attestation

BOA: Bogon Origin Attestation

Checking Consistency of a Registered Route with Corresponding Inetnum and Aut-Num

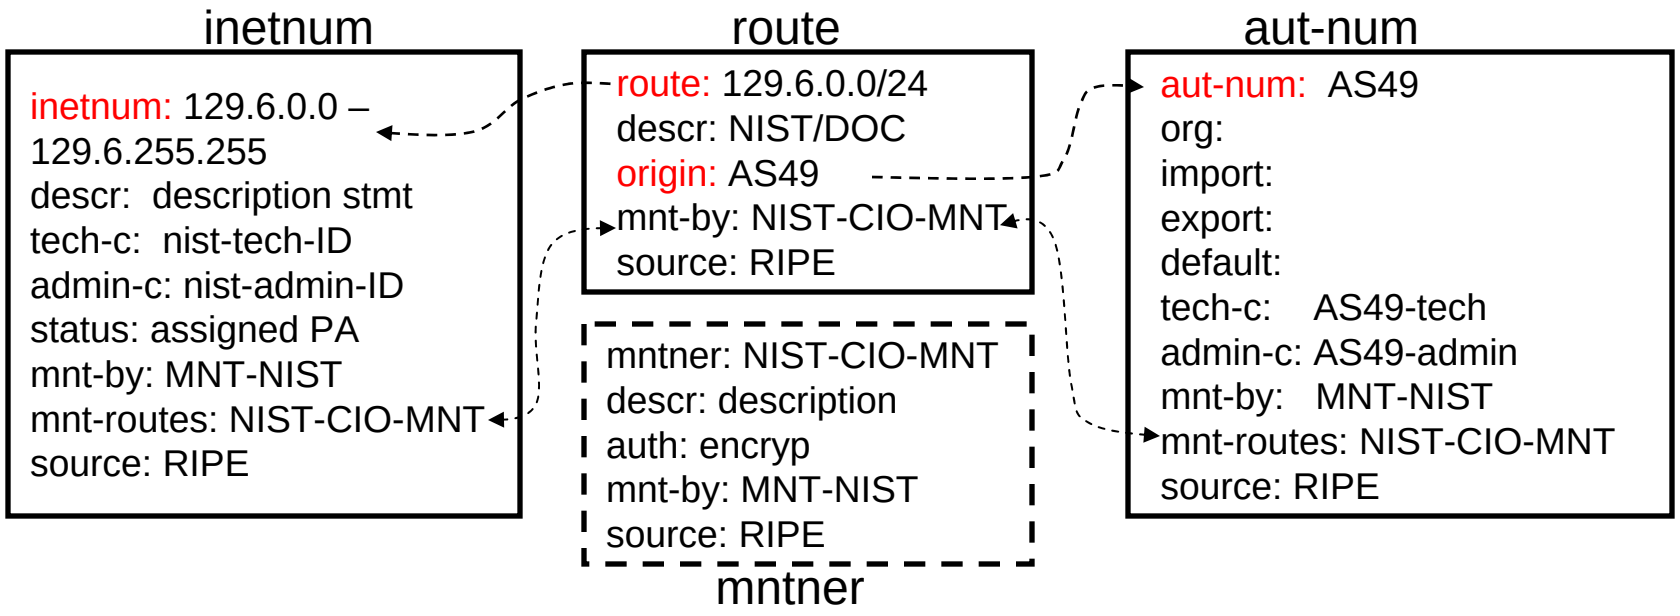

Checking Registry Consistency of Registered Routes (Algorithm)

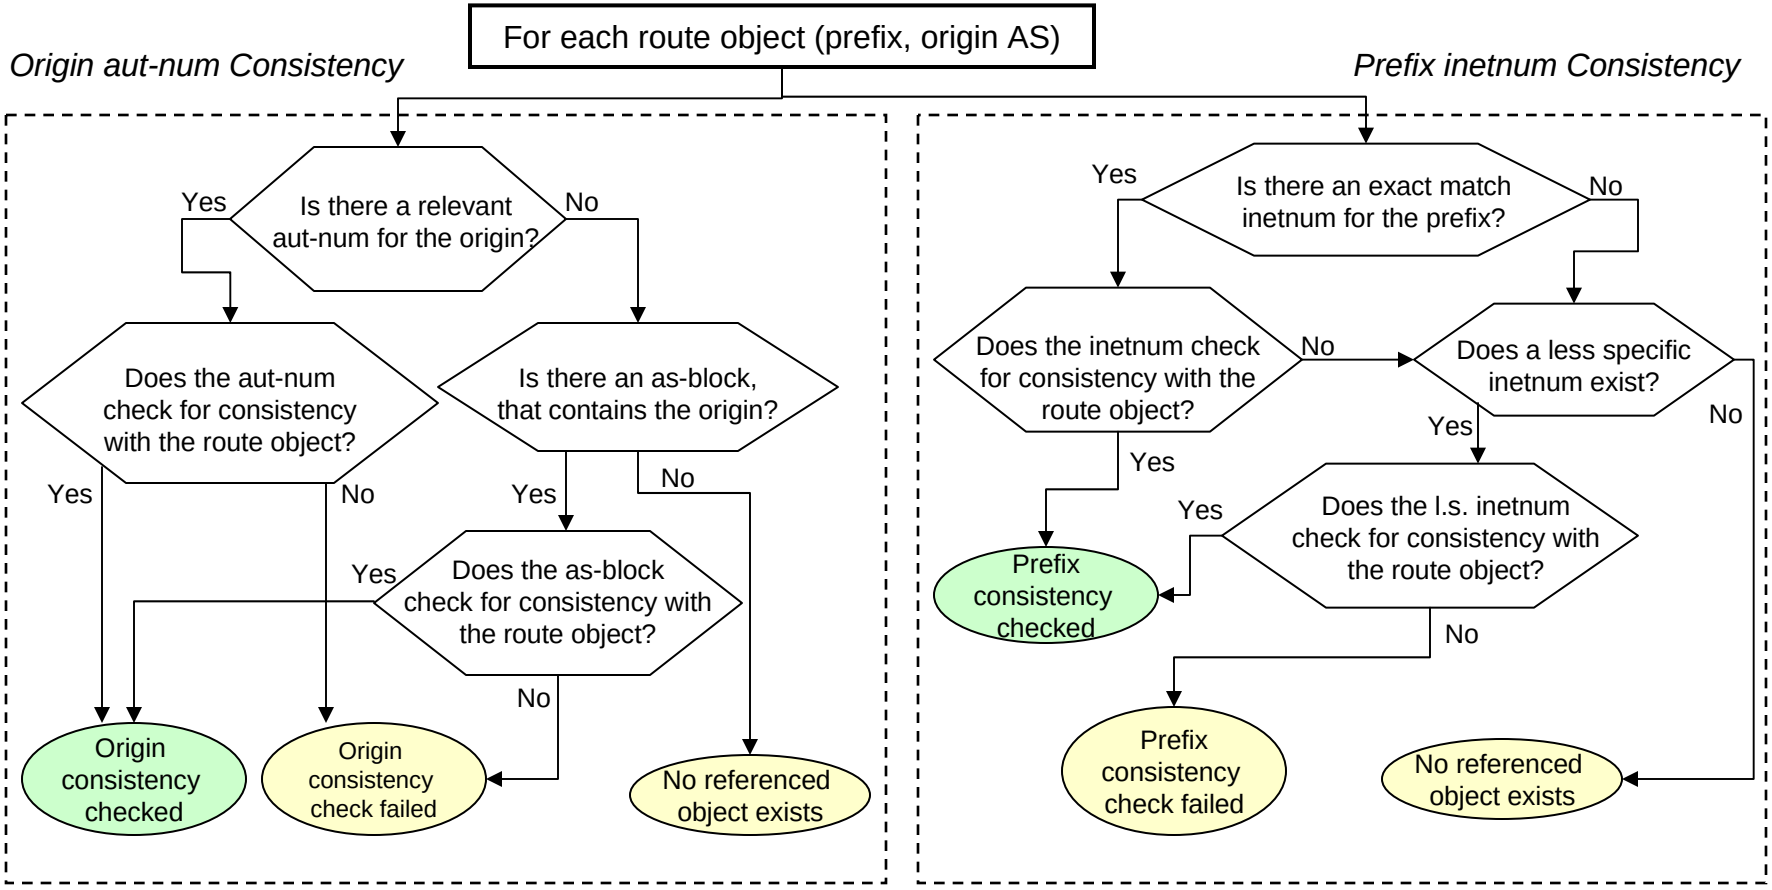

Registry-Based Algorithm for Scoring Routes Observed in Trace Data

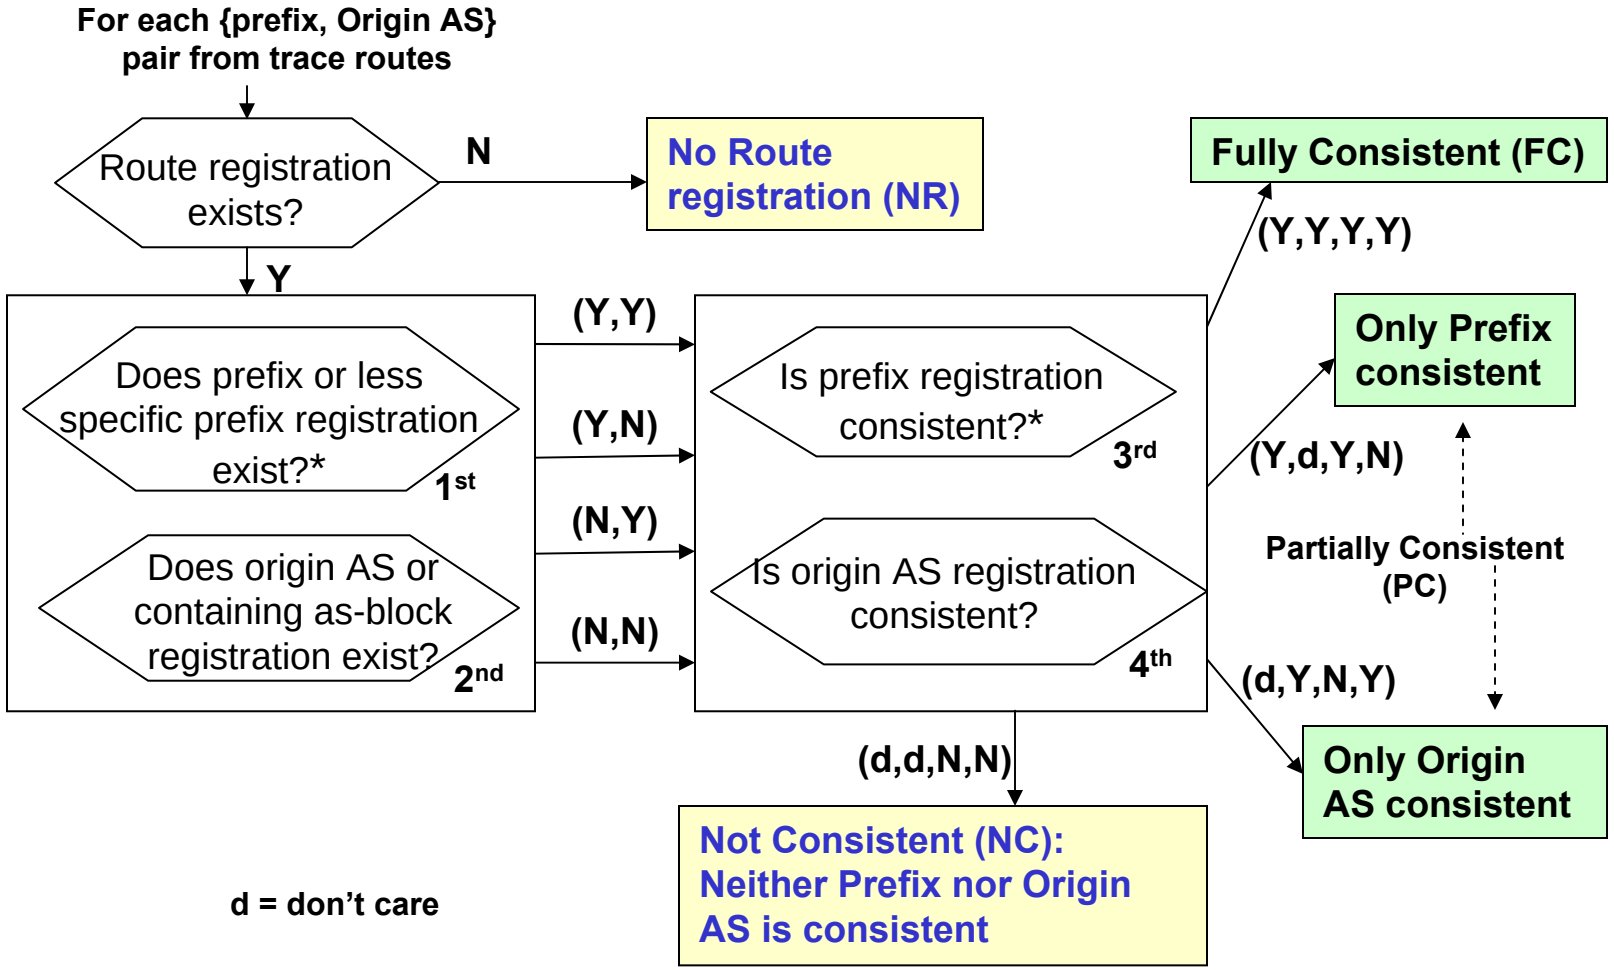

Enhanced History-Based Algorithm for Determining Stability of (p, OAS) in the Trace Data

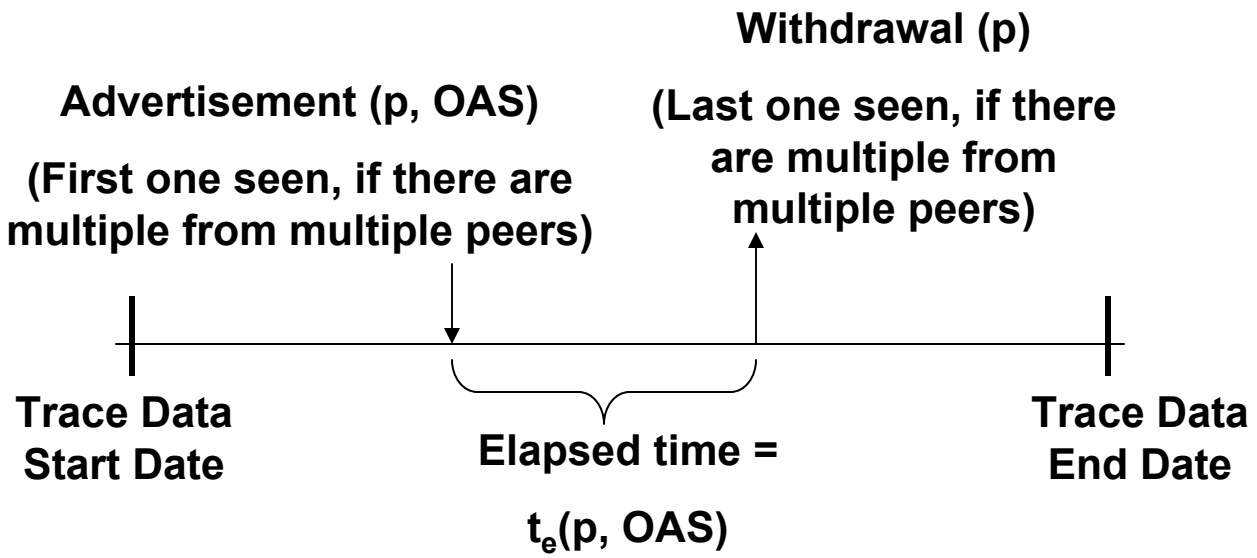

- If (p, OAS) had no withdrawal after the advertisement, set t_e(p, OAS) = 10^6 hours
- If t_e(p, OAS) $\geq$ 48 hours, then (p, OAS) is a stable (prefix, Origin AS) pair
- If t_e(p, OAS) < 48 hours, then (p, OAS) is an unstable (prefix, Origin AS) pair
- Update data is initialized with stable (i.e., persistent for $\geq$ 48 hours) RIB entries

Enhanced Hybrid Algorithm for Validating (p, OAS) in the Trace Data

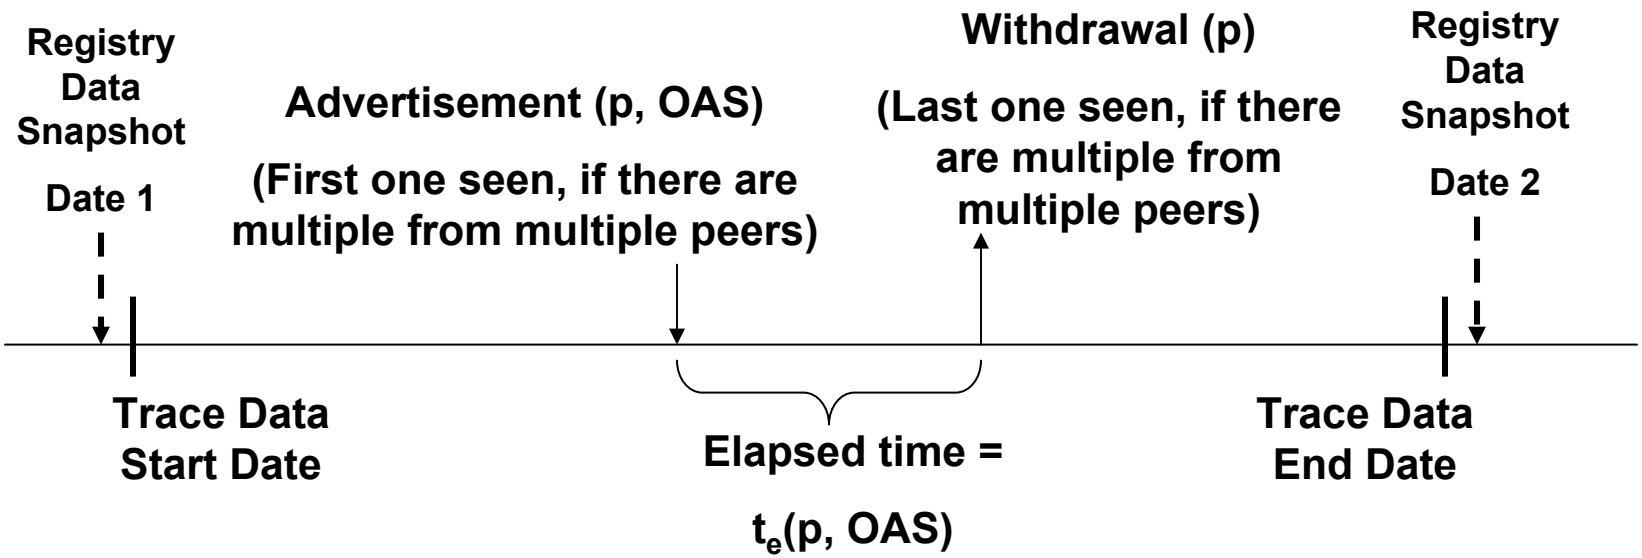

- Use enhanced history-based (i.e., trace-data-based) algorithm as in previous slide
- Complement it with combined results of the registry-based algorithm with data from two dates (close to start and end dates of the history algorithm)
- Registry and historical data can be complementary to each other in enhancing the performance of anomaly detection algorithms

Outline of the Talk

- Known / New BGP robustness schemes
- Evaluation of BGP robustness algorithms
  - Qualitative / comparative analysis of utility
  - Preliminary quantitative results
- Conclusions / Future Work

Origin AS Approval Check List: Comparison

		Which checks are included in each approach?			
	Checks/Questions	Registry based (e.g., Nemecis)	Trace-data based (PGBGP)	Enhanced Trace-data based	Enhanced Hybrid
Q1.	Is prefix registered (same or less specific)?	√			√
Q2.	Is there a route registered (with same or less specific prefix and origin AS)?	√			√
Q3.	Is announced (p, origin AS) fully consistent with corresponding registry objects in RIR/IRR?	√			√
Q4.	Is announced (p, origin AS) partially consistent with corresponding registry objects in RIR/IRR?	√			√
Q5.	Was (p, origin AS) seen in RIB in the last *h* (= 10) days? (Also, if it was suspicious, did it remain in RIB beyond the suspicious period of *s* (= 24) hours?)		√		
Q6.	Would a less specific prefix with the same origin AS pass the test in Q5?		√		
Q7.	Was prefix previously announced by the same origin AS and remained stably (48 hrs or more) in the RIB over the observation period (*d* months)?			√	√
Q8.	Would a less specific prefix with the same origin AS pass the test in Q7?			√	√

Algorithm Robustness Checklist

	Algorithmic Features	Registry based (e.g., Nemecis)	Trace-data based (PGBGP)	Enhanced Trace-data based	Enhanced Hybrid
Data Sets	Utilization of self-consistent registry objects	Yes	No	No	Yes
	Utilization of update history	No	Yes	Yes	Yes
	Utilization of historical RIB entries	No	Yes	Yes	Yes
Situations Handled	Pass a subprefix announcement if a less specific prefix with same origin AS could be passed	Yes	Yes	Yes	Yes
	False Positives: Alert raised when genuine prefix owner announces multi-homing related AS origin change	Moderate probability	High probability	Moderate probability	Low probability
	Alert raised when attacker announces a prefix after sensing it has just been withdrawn	Yes	NO (Path propagates with lower pref)	Yes	Yes
	Pass a subprefix announcement in conjunction with multi-homing related AS origin change	Moderate probability	Low probability	Moderate probability	High probability

*** This is a ballpark qualitative assessment; subject to corroboration using extensive quantitative studies.**

Comparative Analysis of Existing and Enhanced Algorithms

- We have encoded Registry-based, Enhanced Trace-data-based and Enhanced Hybrid algorithms for evaluation
- Enhanced trace-data based algorithm is a variant from PGBGP (see slides 6-9, 14)
- Algorithms are run on top of the NIST TERRAIN* framework
  - Unified database of Registry / Trace data (RIRs, IRRs, RIPE-RIS, Routeviews)
- Tested and compared the algorithms

* TERRAIN: Testing and Evaluation of Routing Robustness in Assurable Inter-domain Networking

Comparative Analysis of Existing and Enhanced Algorithms (Contd.)

For the purpose of this presentation:

- Results focus on Origin AS validation
- Use of Global RIR/IRR/RADB and RIPE RIS data
  - Results are reported broadly for Global RIR/IRR registries as well as specifically for Regional RIR/IRR registries
- Six-month trace-data window (January through June 2007); initialized with stable (persistent for $\geq$ 48 hours) RIB entries
- Registry data – two dates prior to and towards the end of the six-month window (December 12, 2006 and June 18, 2007)
- Results on comparison of algorithms follow

Some Caveats Apply

- This presentation is mainly to demonstrate the capability and to solicit feedback on approach
- Quantitative results are subject to change when the following enhancements to the study are made (ongoing / future work)
  - Consideration of new NetHandle format in ARIN which includes origin AS information
  - Consideration of multiple trace-data collectors (here we considered trace-data from RRC00 only)
  - Use of ROAs/BOAs based on RPKI efforts (in future)

Classification of Observed (p, OAS) Pairs According to Stability / Consistency Scores

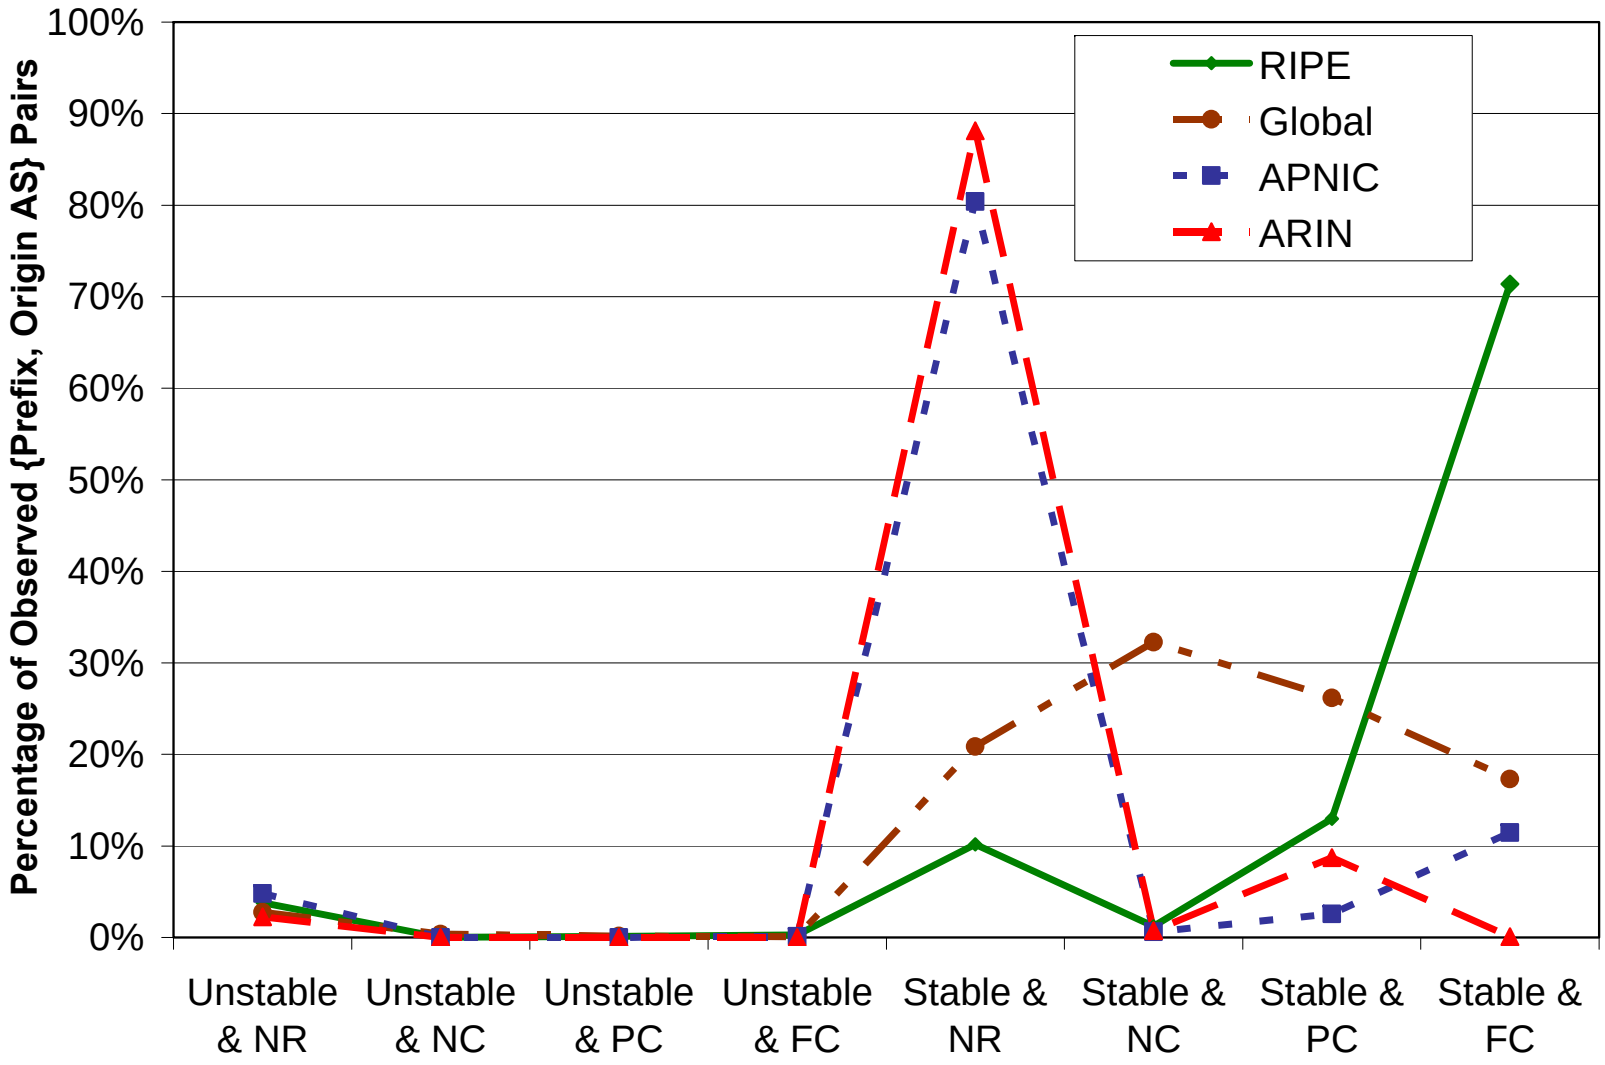

p = prefix; OAS = Origin AS; FC = Fully Consistent; PC = Partially Consistent; NC = Not Consistent; NR = Not Registered

Comparative Performance of Algorithms

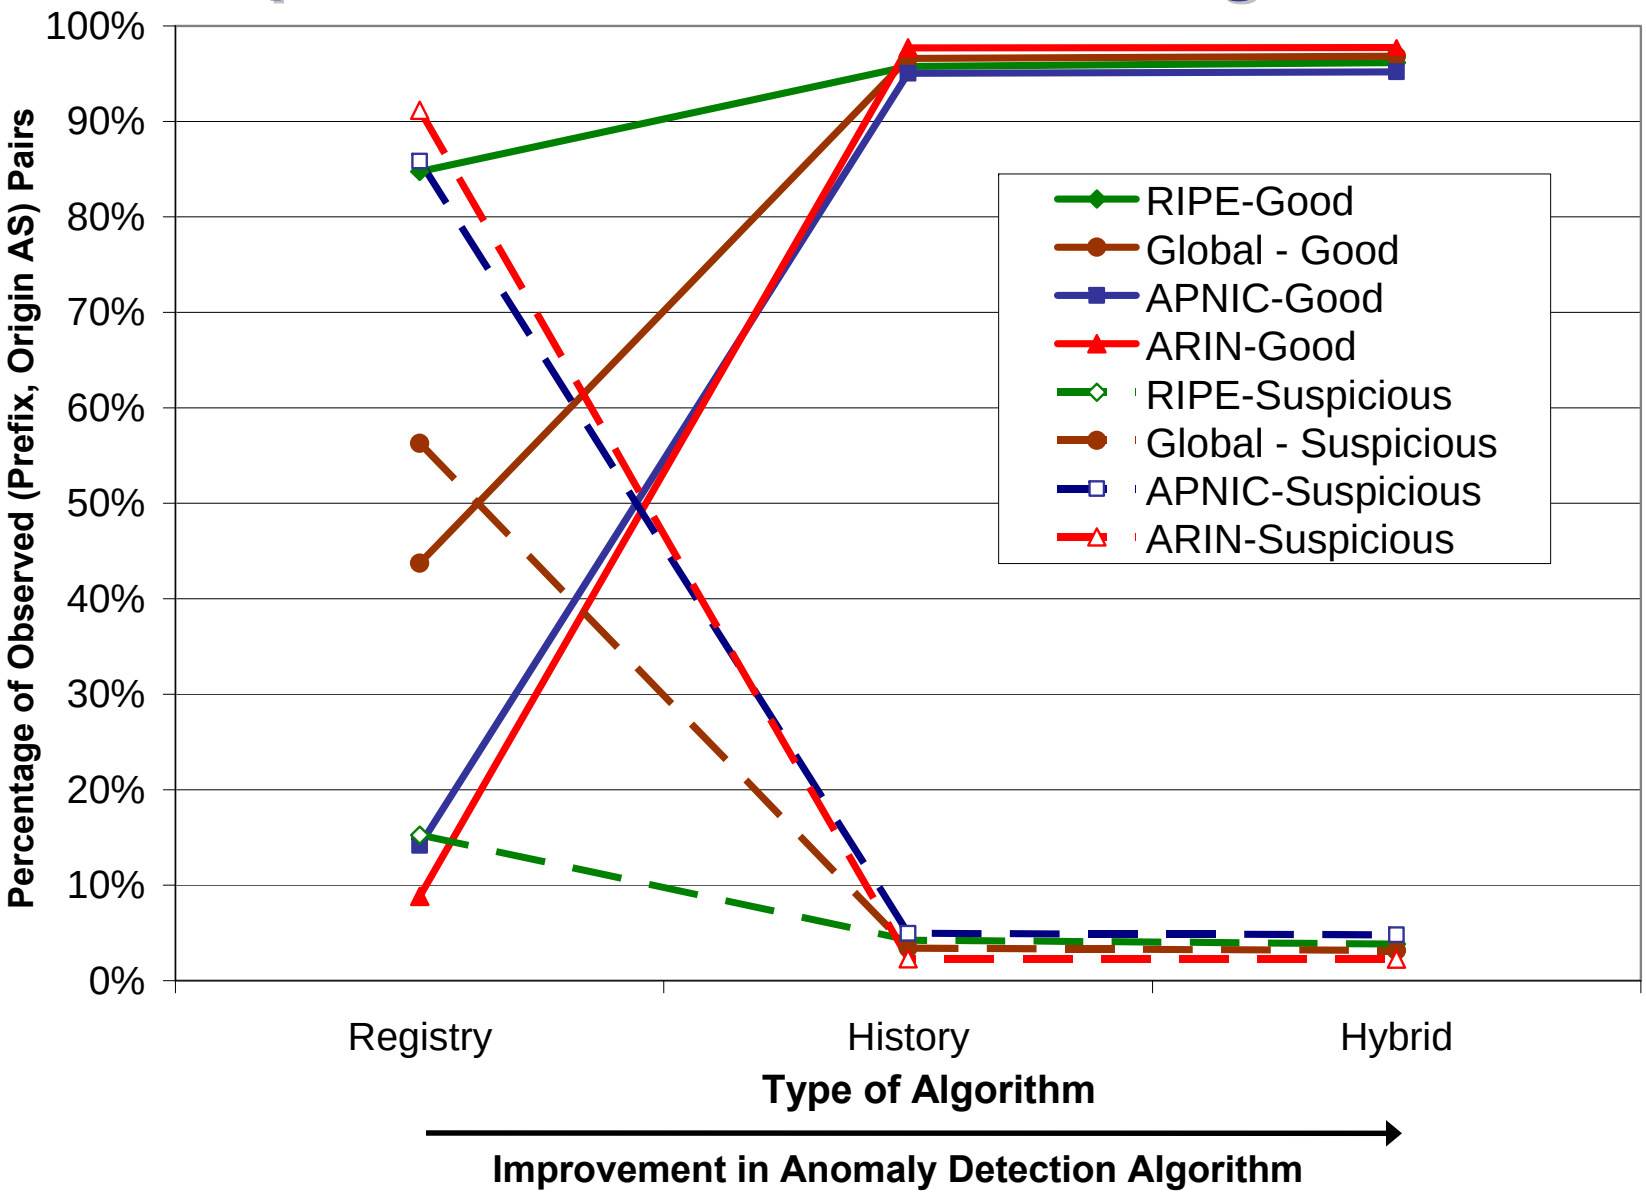

Comparative Performance of Algorithms

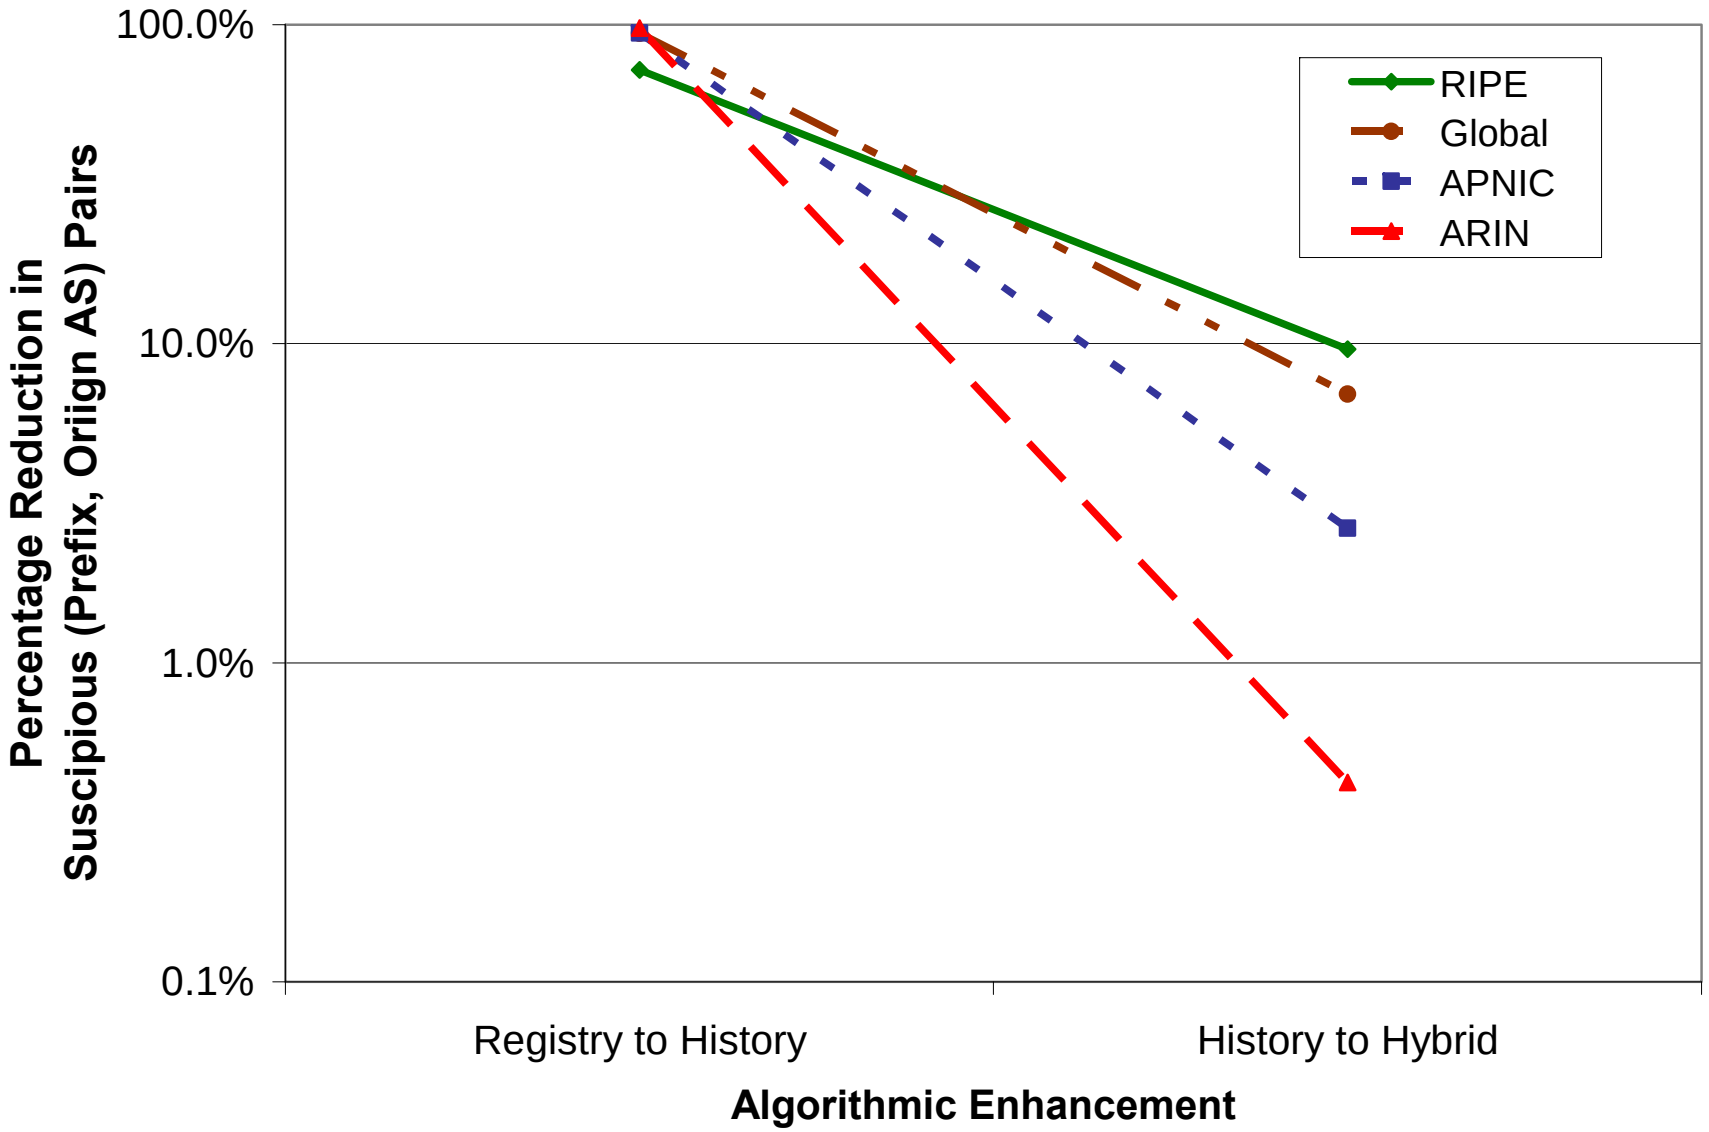

Heatmap Depicting Origin Validation for Announced Prefixes

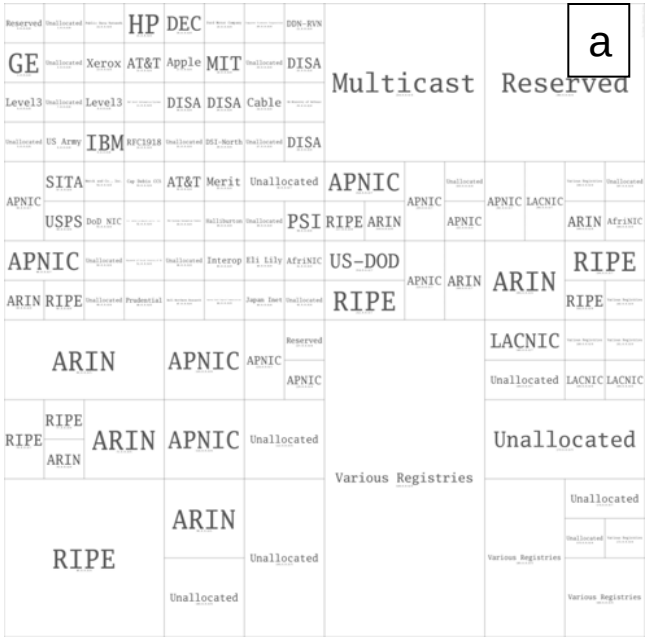

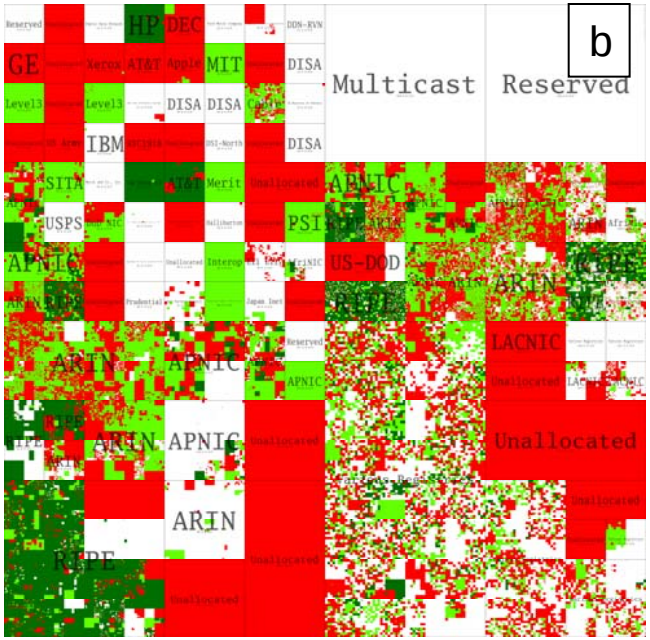

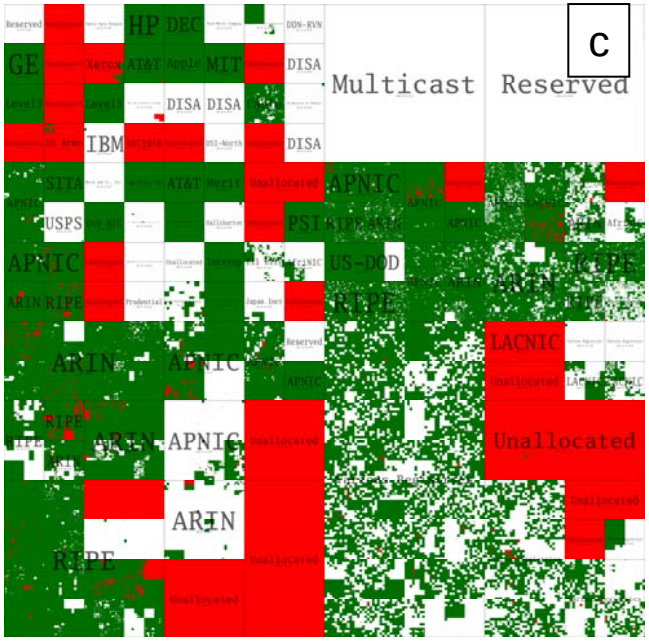

a. **Allocations**
b. **Registry-based Algorithm**
c. **Enhanced Trace-data-based Algorithm**
d. **Enhanced Hybrid Algorithm**

For (b), (c), (d) :

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

White: Not found in trace data

Reference: http://maps.measurement-factory.com/software/ipv4-heatmap.1.html

Checking Origin AS : Comparison of Algorithms

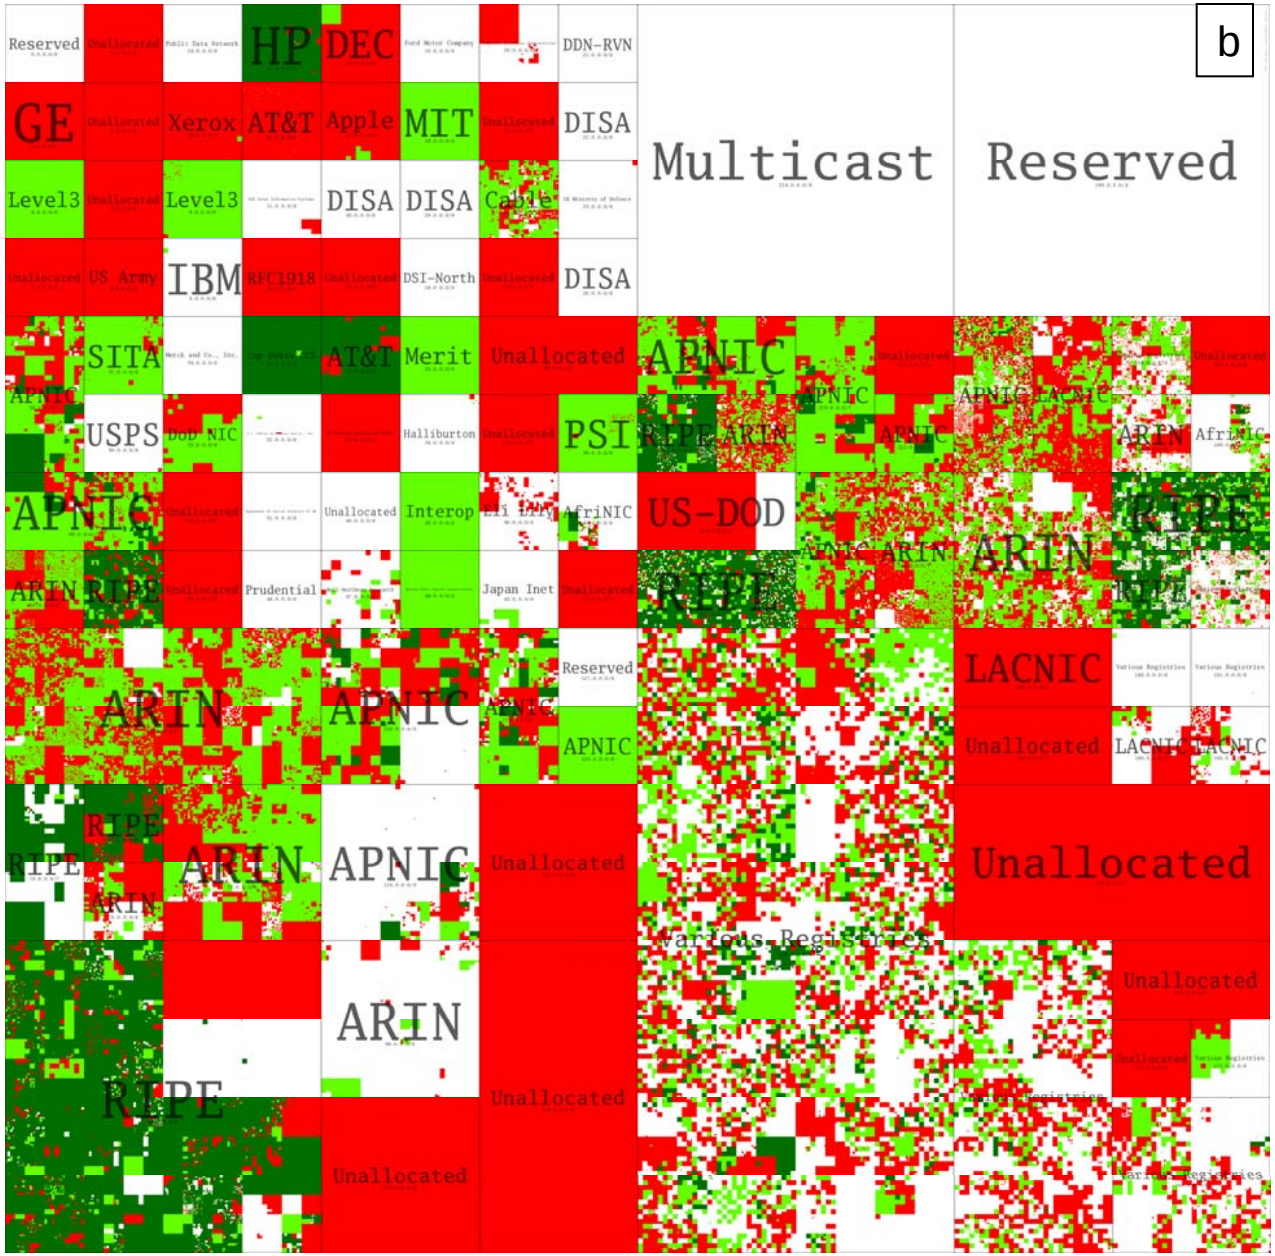

Registry-based Algorithm

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

White: Not found in trace data

Checking Origin AS : Comparison of Algorithms

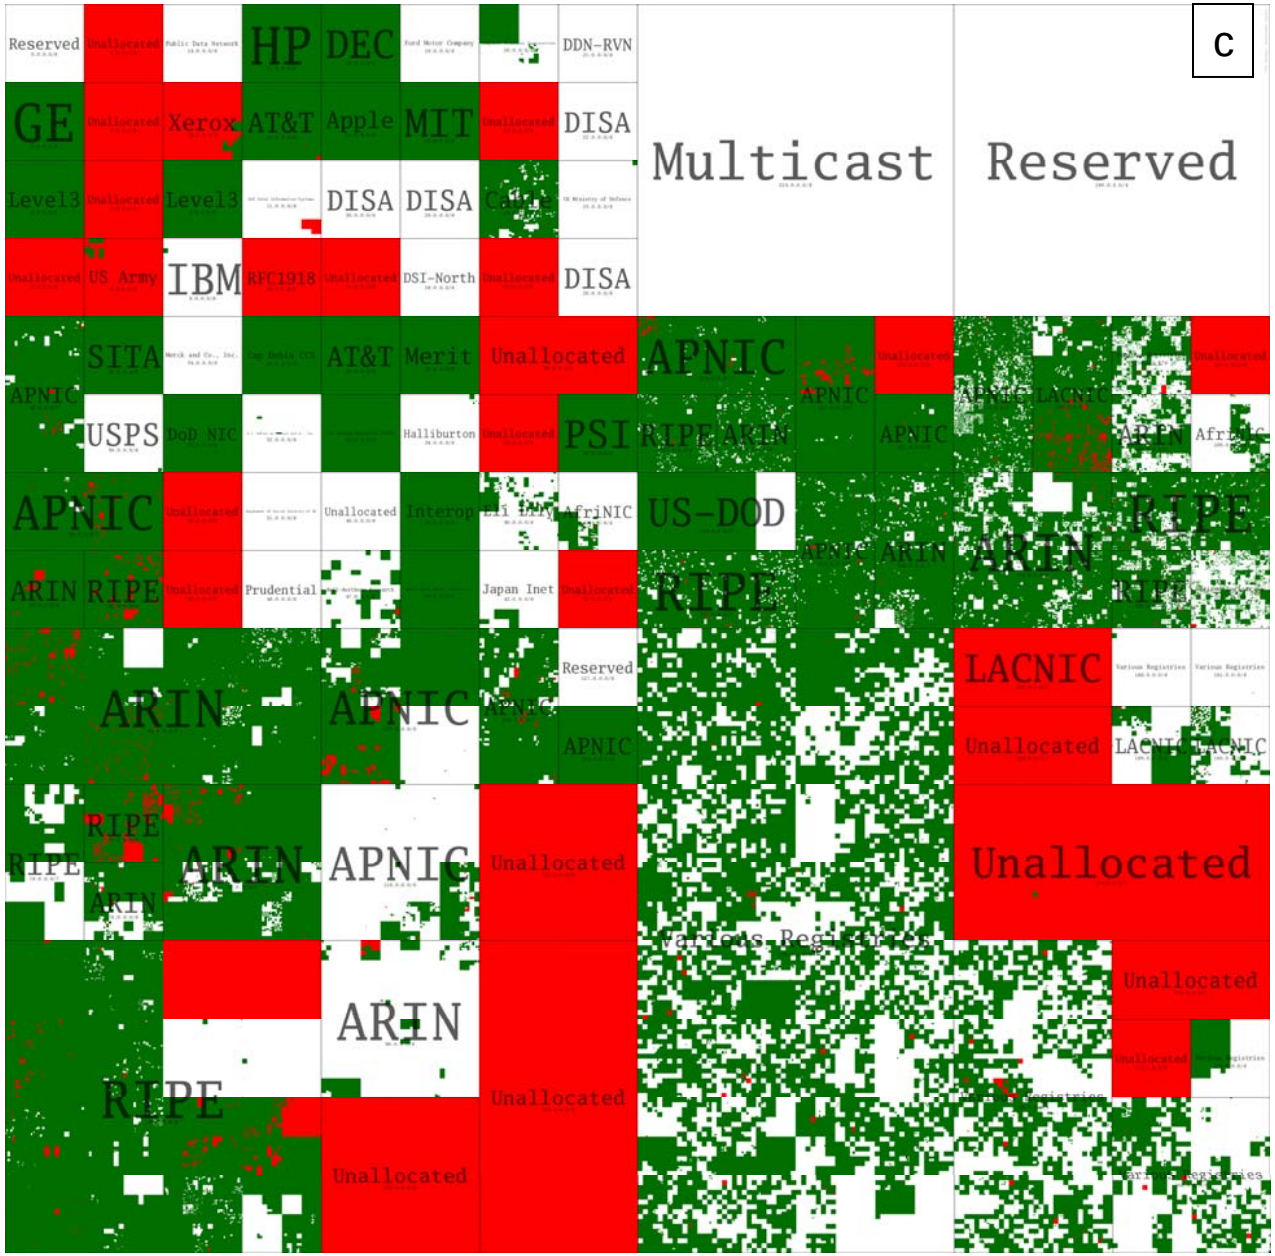

Enhanced trace-data-based Algorithm

Green: Good

Red: Suspicious

White: Not found in trace data

Checking Origin AS : Comparison of Algorithms

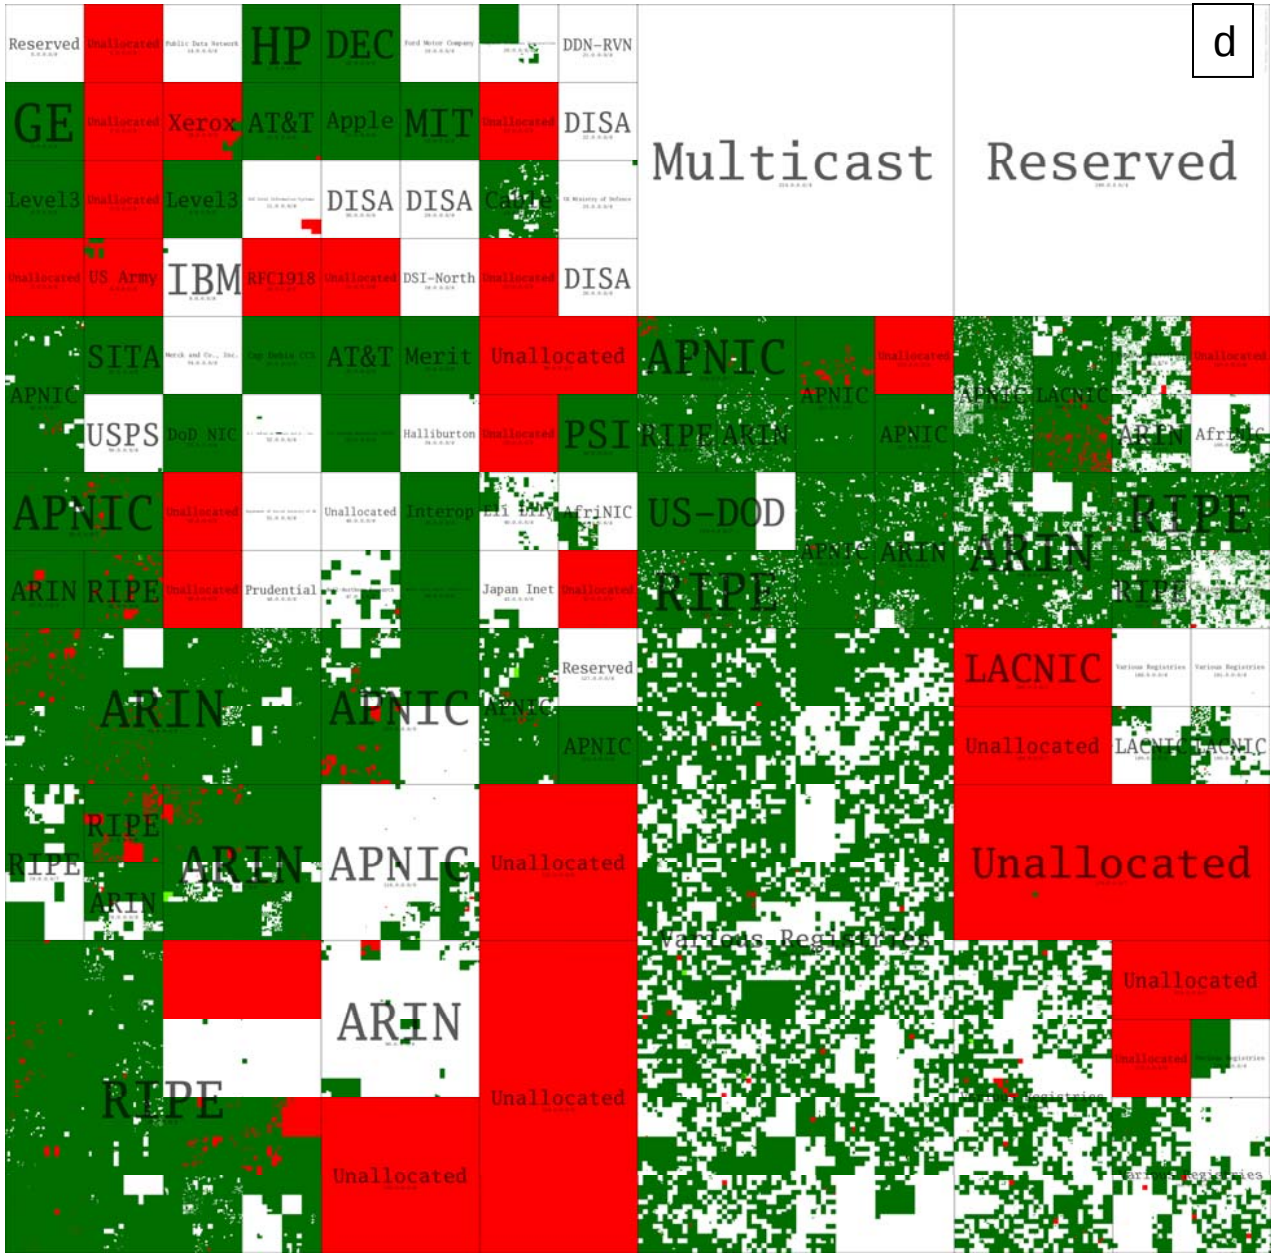

Enhanced Hybrid Algorithm

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

White: Not found in trace data

Prefixes with Multiple Origin ASes

# Origin ASes	# Prefixes
1	476243
2	55673
3	10419
4	2683
5	965

For prefixes with two Origin ASes:

OAS1	OAS2	# Prefixes
FC + Stable	FC/PC + Unstable	23
PC + Stable	FC/PC + Unstable	41
NC + Stable	FC/PC + Unstable	104
NR + Stable	FC/PC + Unstable	0
Total		168

- Statistics of prefixes with two Origin ASes where the primary path is stable (with or without consistency in the registry), while the secondary (failover) path is transient (unstable) but consistent in the registry

Analysis of Registered But Unobserved Routes

- Large number of {prefix, origin} pairs registered but never announced
- In most cases, super-prefixes are announced with the same origin AS (as in registered route) or a different origin AS
- Is it due to aggregation by a higher tier ISP?
- Needs further investigation

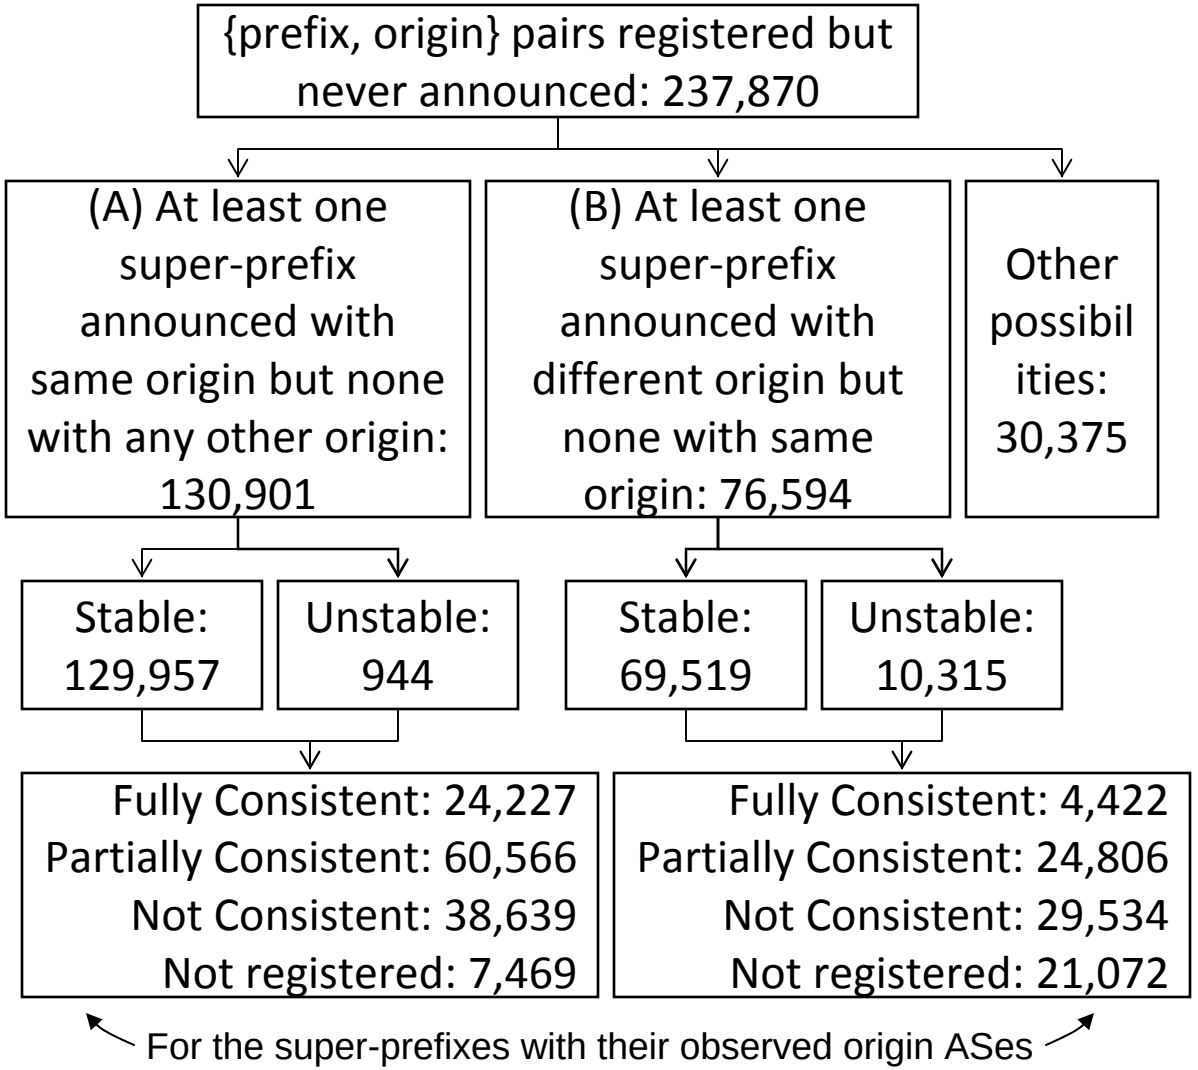

Conclusions and Planned Future Work

- Enhanced hybrid algorithm – history and registry data have complementary influence on improvement in origin validation
- Some ***caveats*** apply in the reported results (To Do list)
  - Consideration of new NetHandle format in ARIN which includes origin AS information
  - Consideration of multiple trace-data collectors (here we considered trace-data from RRC00 only)
- Further testing for robustness of the algorithms will be performed with extensive real and synthetic trace data
- This will lead to numerical results for benchmarking the algorithms
- Help industry understand implications of proposals emerging from various ongoing R&D projects

Thank you!

Questions?

Full paper available including detailed results and discussion:
http://www.antd.nist.gov/~ksriram/NIST_BGP_Robustness.pdf

Backup Slides

Trace-Data Based Algorithm: Differences Relative to PGBGP

- PGBGP considers a moving 10-day window of trace data
- We keep in our stable list any (p, OAS) pair that remained in the RIB for 48 hours or more at least once in our observation period (six months or more)
- The idea is that backup protection paths may be infrequently used
  - An AS may have served as the origin AS a few months ago during failover and is used again now
  - It is better to make that part of “stable” history if the (p, OAS) pair earlier remained in RIB for 48 hours or more
- We augment the above with consideration of route registration and registry consistency in our enhanced hybrid algorithm

YouTube Hijack: Background Information

Prefix normally advertised by YouTube: 208.65.152.0/22 via AS 36561

Related (overlapping) prefixes seen historically and stayed stable for 48-hour or more:

Prefix	Origin AS	AS name	Time
208.65.152.0/22	AS 36561	YOUTUBE: YouTube, Inc.	02-20-08 15:43:50 (RIPE RIS) 02-20-08 15:37:46 (rrc02)

YouTube Hijack: Sequence of Events

Prefix normally advertised by YouTube: 208.65.152.0/22 via AS 36561

Date: 2/20/08 15:43:50	Normal announcement of 208.65.152.0/22 by AS 36561
15:37:46	rrc02: Prefix: 208.65.152.0/22, Origin: 36561, AS path: 14361 36561
Date: 2/24/08	
18:47:45	first evidence of hijacked route propagating in Asia, AS path 3491 17557 (208.65.153.0/24)
18:37:46	rrc02: Prefix: 208.65.153.0/24, Origin: 17557, AS path: 2497 3491 17557
18:49:00	most of the DFZ now carrying the bad route (and 93 ASNs)
18:49:30	all providers who will carry the hijacked route have it (total 97 ASNs)
20:07:25	YouTube, AS 36561 advertises the /24 that has been hijacked
20:07:25	rrc02: Prefix: 208.65.153.0/24, Origin: 36561, AS path:19089 3549 36561
20:08:30	a total of 40 some-odd providers have stopped using the hijacked route

Notes: rrc02 update data (yellow rows) is from TERRAIN database

Event timeline (white rows) obtained from Martin A. Brown's blog at Renesys:

http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml

Continued on next page …

YouTube Hijack: Sequence of Events (Contd.)

Prefix normally advertised by YouTube: 208.65.152.0/22 via AS 36561

Date: 2/24/08	
20:18:43	and now, two more specific /25 routes are first seen from 36561
20:18:43 20:18:43	rrc02: Prefix: 208.65.153.0/25, Origin: 36561, AS path:19089 3549 36561 rrc02: Prefix: 208.65.153.128/25, Origin: 36561, AS path: 19089 3549 36561
20:19:37	25 more providers prefer the /25 routes from 36561
20:50:59	evidence of attempted prepending, AS path was 3491 17557 17557
20:59:39	hijacked prefix is withdrawn by 3491, who disconnect 17557

Notes: rrc02 update data (yellow rows) is from TERRAIN database

Event timeline (white rows) obtained from Martin A. Brown's blog at Renesys:

http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml

How Effective in a YouTube Like Incident:

Detecting and Alerting the Attack by Pakistan Telecom

		Results of checks included in each approach			
	Checks/Questions	Registry-based (e.g., Nemecis)	Trace-data based (PGBGP)	Enhanced Trace-data-based	Enhanced Hybrid
Q1.	Is prefix registered (same or less specific)?	**No**			**No**
Q2.	Is there a route registered (with same or less specific prefix and origin AS)?	**No**			**No**
Q3.	Is announced (p, origin AS) fully consistent with corresponding registry objects in RIR/IRR?	**No**			**No**
Q4.	Is announced (p, origin AS) partially consistent with corresponding registry objects RIR/IRR?	**No**			**No**
Q5.	Was (p, origin AS) seen in RIB in the last *h* (= 10) days? (Also, if it was suspicious, did it remain in RIB beyond the suspicious period of *s* (= 24) hours?)		**No**		
Q6.	Would a less specific prefix with the same origin AS pass the test in Q5?		**No**		
Q7.	Was prefix previously announced by the same origin AS and remained stably (48 hrs or more) in the RIB over the observation period (*d* months)?			**No**	**No**
Q8.	Would a less specific prefix with the same origin AS pass the test in Q7?			**No**	**No**

How Effective in a YouTube Like Incident:

Detecting and Allowing Recovery Using Sub-prefixes by YouTube

		Results of checks included in each approach			
	Checks/Questions	Registry-based approach (e.g., Nemecis)	Trace-data based approach (PGBGP)	Enhanced Trace-Data-Based Approach	Enhanced Hybrid
Q1.	Is prefix registered (same or less specific)?	**Yes**			**Yes**
Q2.	Is there a route registered (with same or less specific prefix and origin AS)?	**Yes**			**Yes**
Q3.	Is announced (p, origin AS) fully consistent with corresponding registry objects in RIR/IRR?	**Yes**			**Yes**
Q4.	Is announced (p, origin AS) partially consistent with corresponding registry objects RIR/IRR?	**Yes**			**Yes**
Q5.	Was (p, origin AS) seen in RIB in the last *h* (= 10) days? (Also, if it was suspicious, did it remain in RIB beyond the suspicious period of *s* (= 24) hours?)		**No**		
Q6.	Would a less specific prefix with the same origin AS pass the test in Q5?		**Yes**		
Q7.	Was prefix previously announced by the same origin AS and remained stably (48 hrs or more) in the RIB over the observation period (*d* months)?			**No**	**No**
Q8.	Would a less specific prefix with the same origin AS pass the test in Q7?			**Yes**	**Yes**

YouTube Hijack: Actions by Different Algorithms

Time	Event	Registry -based	PHAS	PGBGP	Enhanced Hybrid Algorithm
Date: 2/20/08 15:43:50Z	Normal /22 Re-Advertisement	No alert	No alert	Propagate update	Propagate update
Date: 2/24/08 15:37:46	Hijack attempt with /24 subprefix	Alert	Alert: new origin	Quarantine update	Quarantine update
18:37:46	Recovery attempt with /24 subprefix	No alert	Alert: Notify subprefix	Propagate update	Propagate update
20:07:25	Recovery attempt with /25 subprefix	No alert	Alert: Notify subprefix	Propagate update	Propagate update

- The proposed enhanced hybrid algorithm would effectively deal with certain special situations that did not manifest in this set of events.