NIST National Institute of Standards and Technology

INFORMATION TECHNOLOGY LABORATORY

Trustworthy Networking Program

Routing Data Quality and Its Impact on BGP Anomaly Detection 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)

Project website:

http://www.antd.nist.gov/bgp_security/

ISOC Routing Resiliency Measurements Workshop, Atlanta

November 2012

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.

Talk Based on Previous Publications

- DHS CATCH Conference, Washington DC, March 2009

http://www.nist.gov/itl/antd/upload/NIST_BGP_Robustness-2.pdf

- NANOG-45, Santo Domingo, January 2009

http://www.nanog.org/meetings/nanog45/abstracts.php?pt=MTE5NSZuYW5vZzQ1&nm=nanog45

- ARIN-23, San Antonio, TX, April 2009

https://www.arin.net/participate/meetings/reports/ARIN_XXIII/pdf/monday/nethandles.pdf

Theme of the Talk

- Registered data (RIR, IRR, RADB, RPKI, etc.) as well as historical BGP trace data are and/or will likely be the basis for implementing routing robustness and security
- Characterization of correctness and completeness of the data
- Data pruning to improve its reliability
- What implications does the data quality have on BGP robustness/security algorithms?
  - Focus: Reduce probability of false alarms & false negatives

One Aspect of BGP Robustness Problem Space

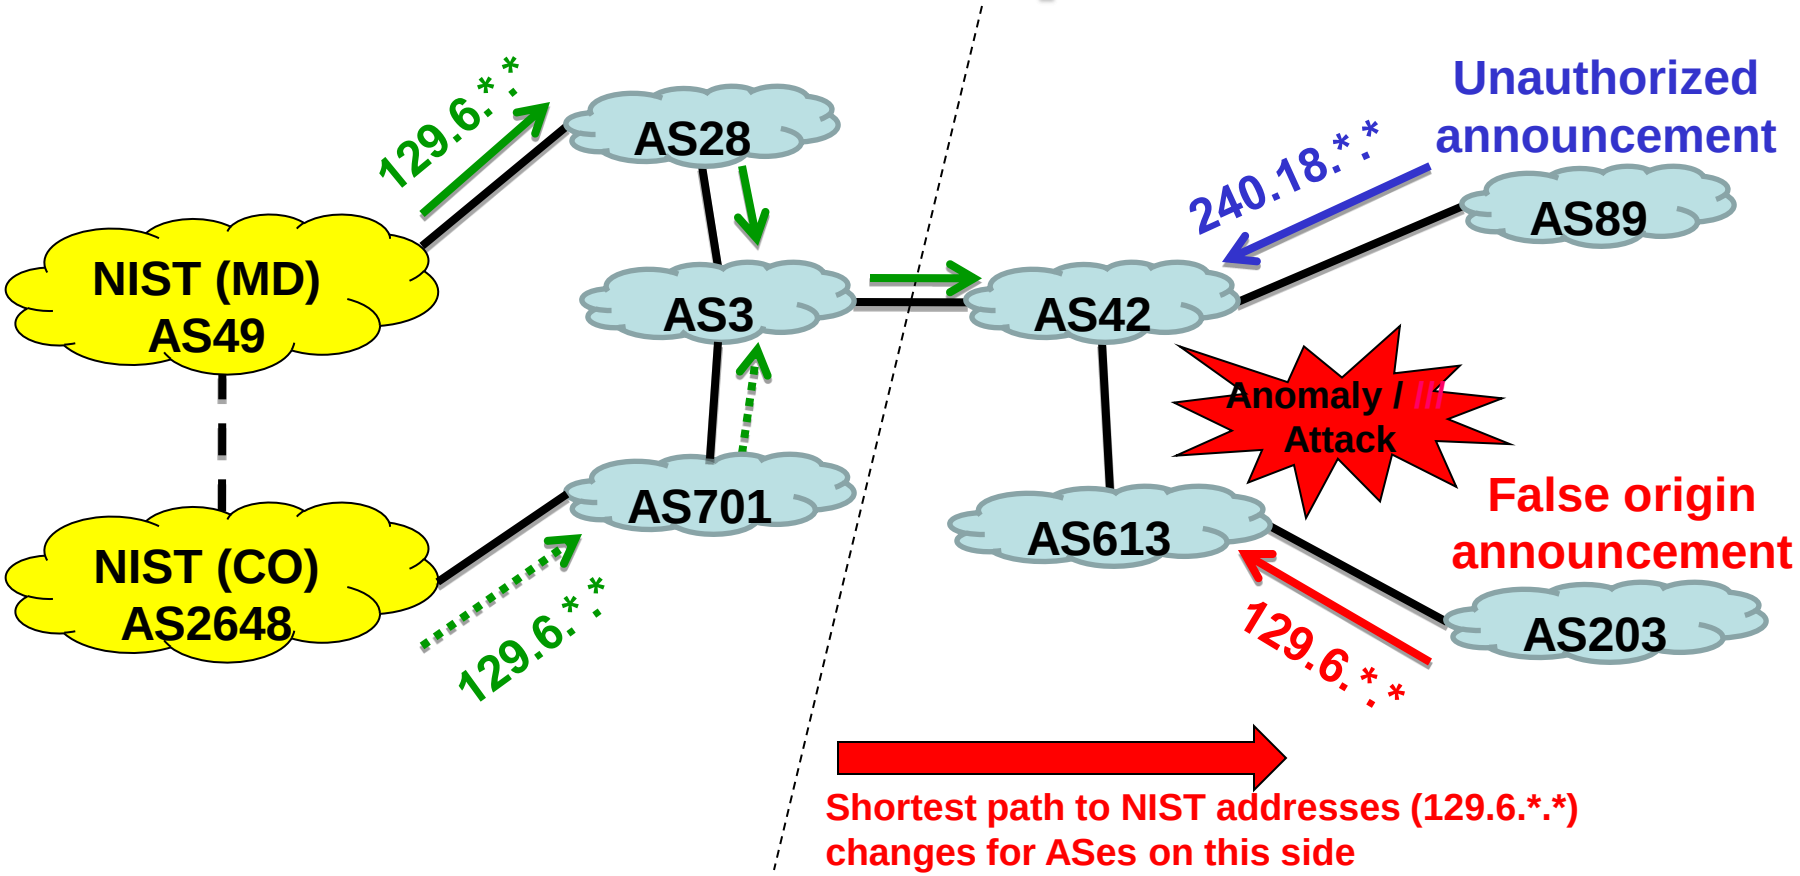

- **Other aspects: Route leaks, Path modification**

Data Driven BGP Robustness

What are the Data Sources?

- Addressing Registries
  - global databases of address block and autonomous system number assignments.
- Routing Registries
  - loosely maintained global databases of contractual relationships for routing services.
- Monitoring Data
  - public BGP monitoring and measurement projects that collect BGP protocol exchanges at various spots around the Internet.

Why is this hard?

- Registries
  - known to be incomplete and inaccurate, and are maintained in differing formats, by differing processes in different regions of the world.
- Robustness Algorithms
  - to be effective, must make precise policy decisions from imperfect data.
- Needle in a Hay Stack
  - millions of BGP update messages per day; millions of registry entries; rare but potent threats.

Solution Components / Players

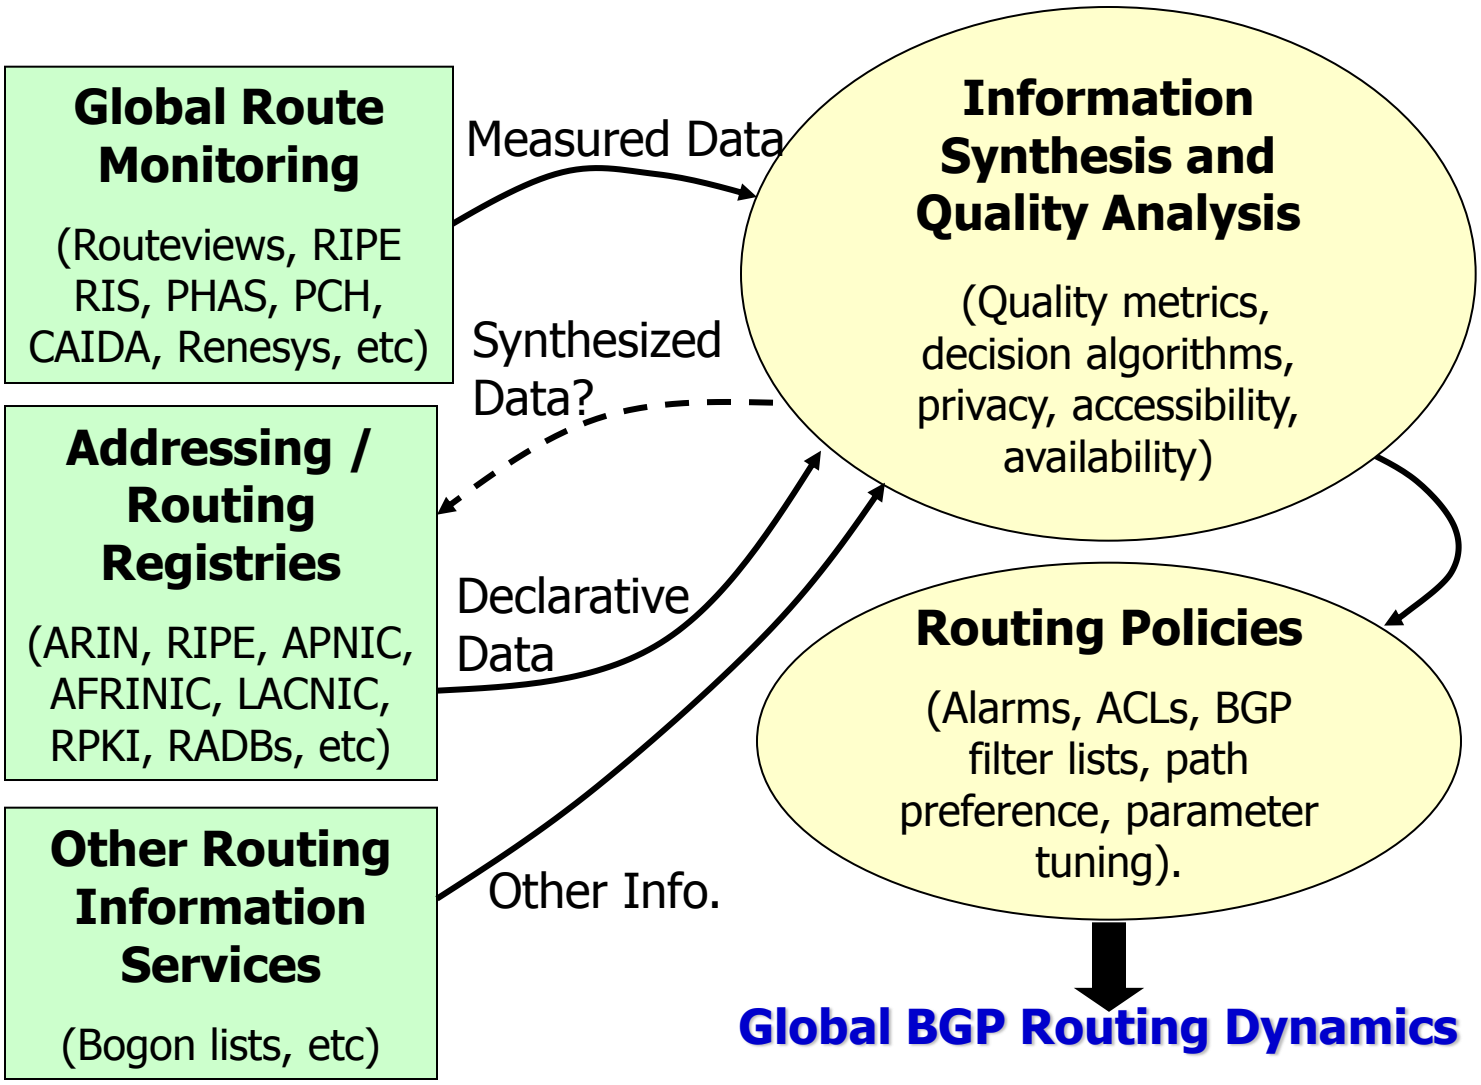

Registry Data and Analysis of Its Completeness and Correctness

Registry Data Object Counts by Source

	route			inetnum (ARIN NetHandle)			aut-num (ARIN ASHandle)		
RIR/IRR	06/18/2007	10/18/2008	Incr	06/18/2007	10/18/2008	Incr	06/18/2007	10/18/2008	Incr
ARIN	7,330	8,201	12%	338 (1,618,197)	434 (1,924,454)	28% 19%	758 (18,050)	890 (19,678)	17% 9%
RIPENCC	71,569	89,957	26%	2,044,536	2,458,119	20%	14,106	16,969	20%
APNIC*	23,616	35,515	50%	822,891	1,080,999	31%	4,559	5,347	17%
AFRINIC	0	0		13,948	22,706	63%	342	445	30%
LACNIC**	0	0		45,346	83,036	83%	1,219	1,339	10%
Standalone IRRs+	345,129	497,124	44%	1	1		3,785	4,643	23%
Total:	447,644	630,797	41%	2,927,060 (1,618,197)	3,645,295 (1,924,454)	25% 19%	24,769 (18,050)	29,633 (19,678)	20% 9%

*** Includes TWNIC, JPIRR, JPNIC and APNIC**
**** RIR only**
+ Independent IRR databases that are mirrored via the RADB website including RADB, but *EXCLUDING ARIN, APNIC, JPIRR and RIPE*
Note that route objects can be registered at any IRR regardless of where the address spaces are allocated.

Distribution of Prefix Length of inetnum (RPSL) and NetHandle (SWIP)

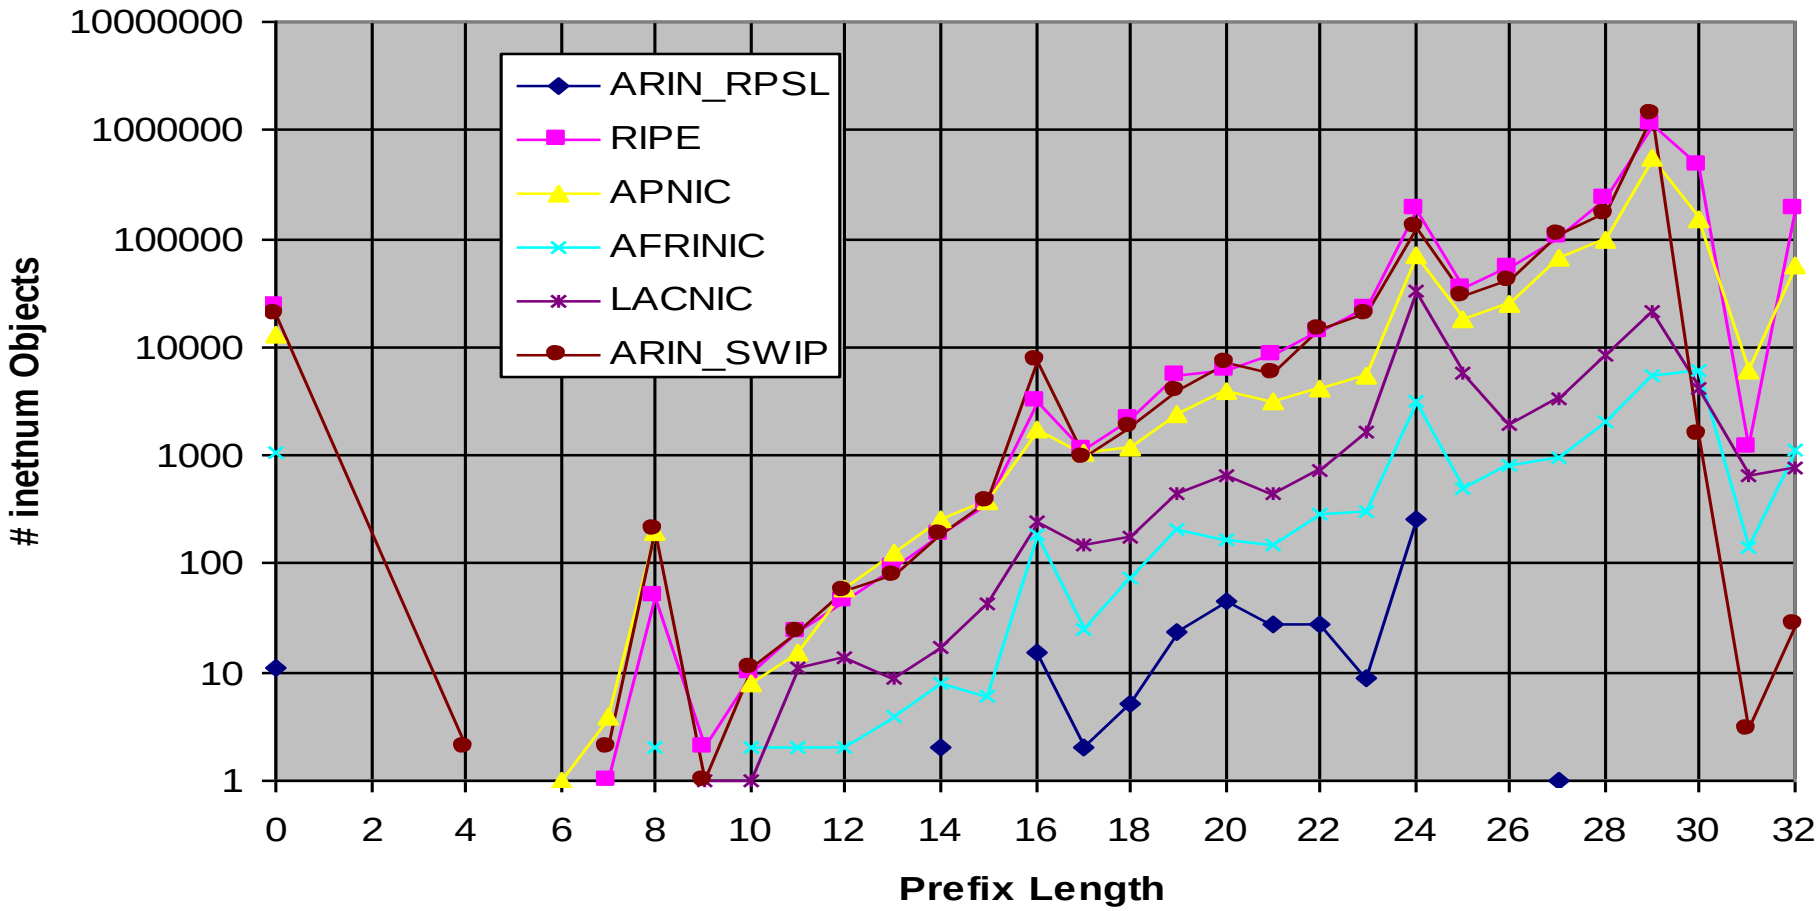

- Length 0 indicates that an address block cannot be represented by a single CIDR
- Length 4 specifies Multicast and Reserved Future Use blocks
- Some Legacy and ERX blocks may be included in one or more RIRs

Distribution of Prefix Length of Route Objects in IRR

Registry Data Date: 2008-10-18

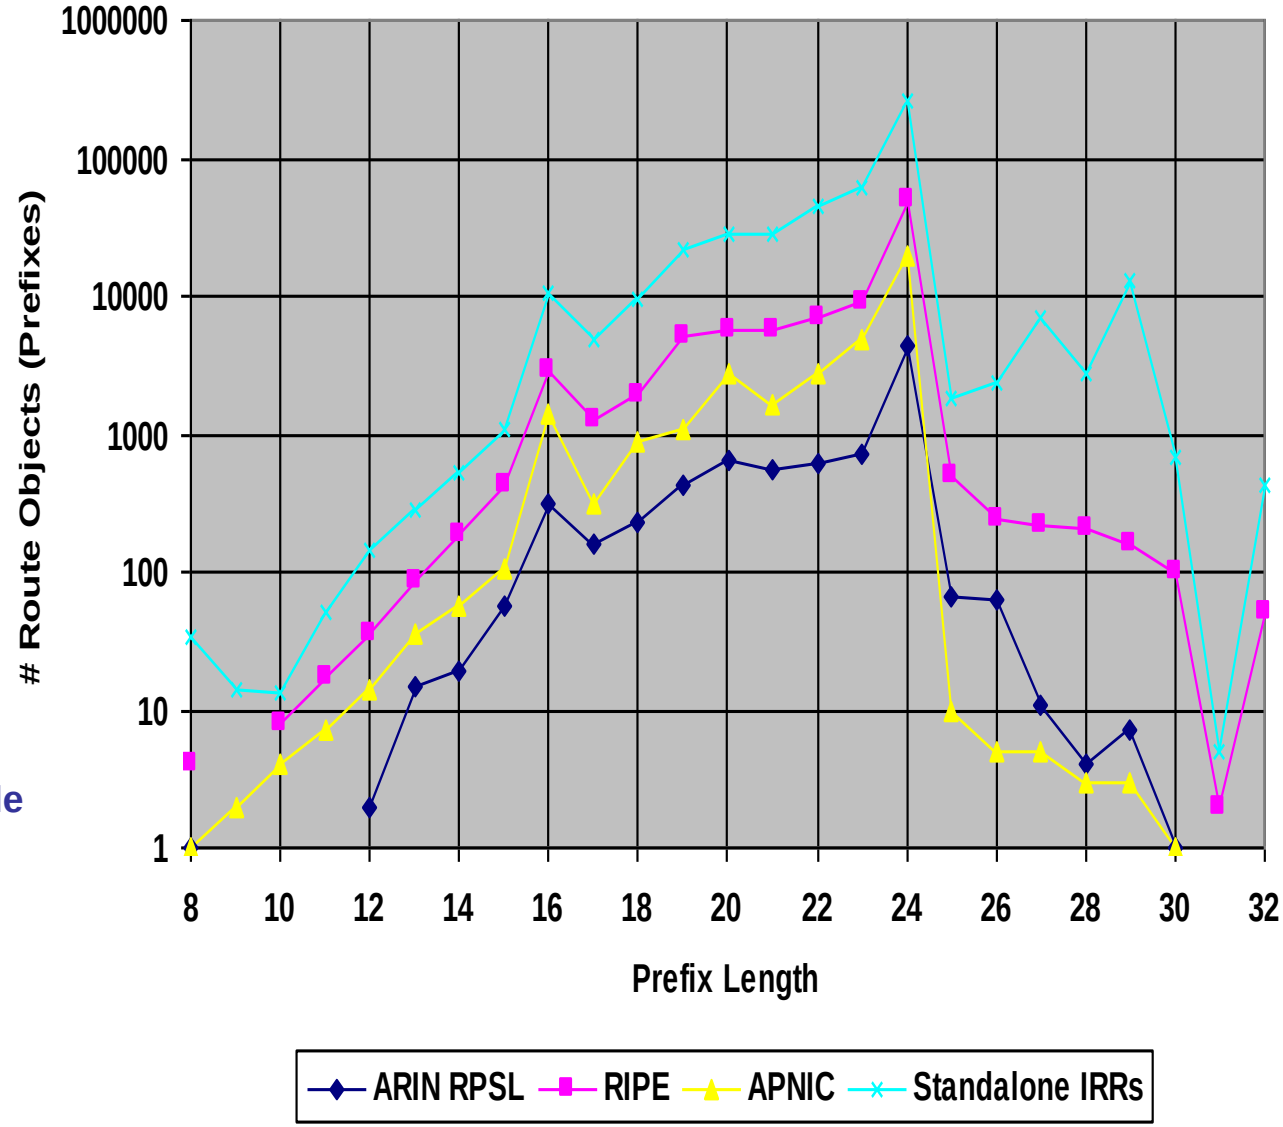

Log scale

Distribution of Sources of Prefix Allocations of Route Objects Registered to Standalone IRRs

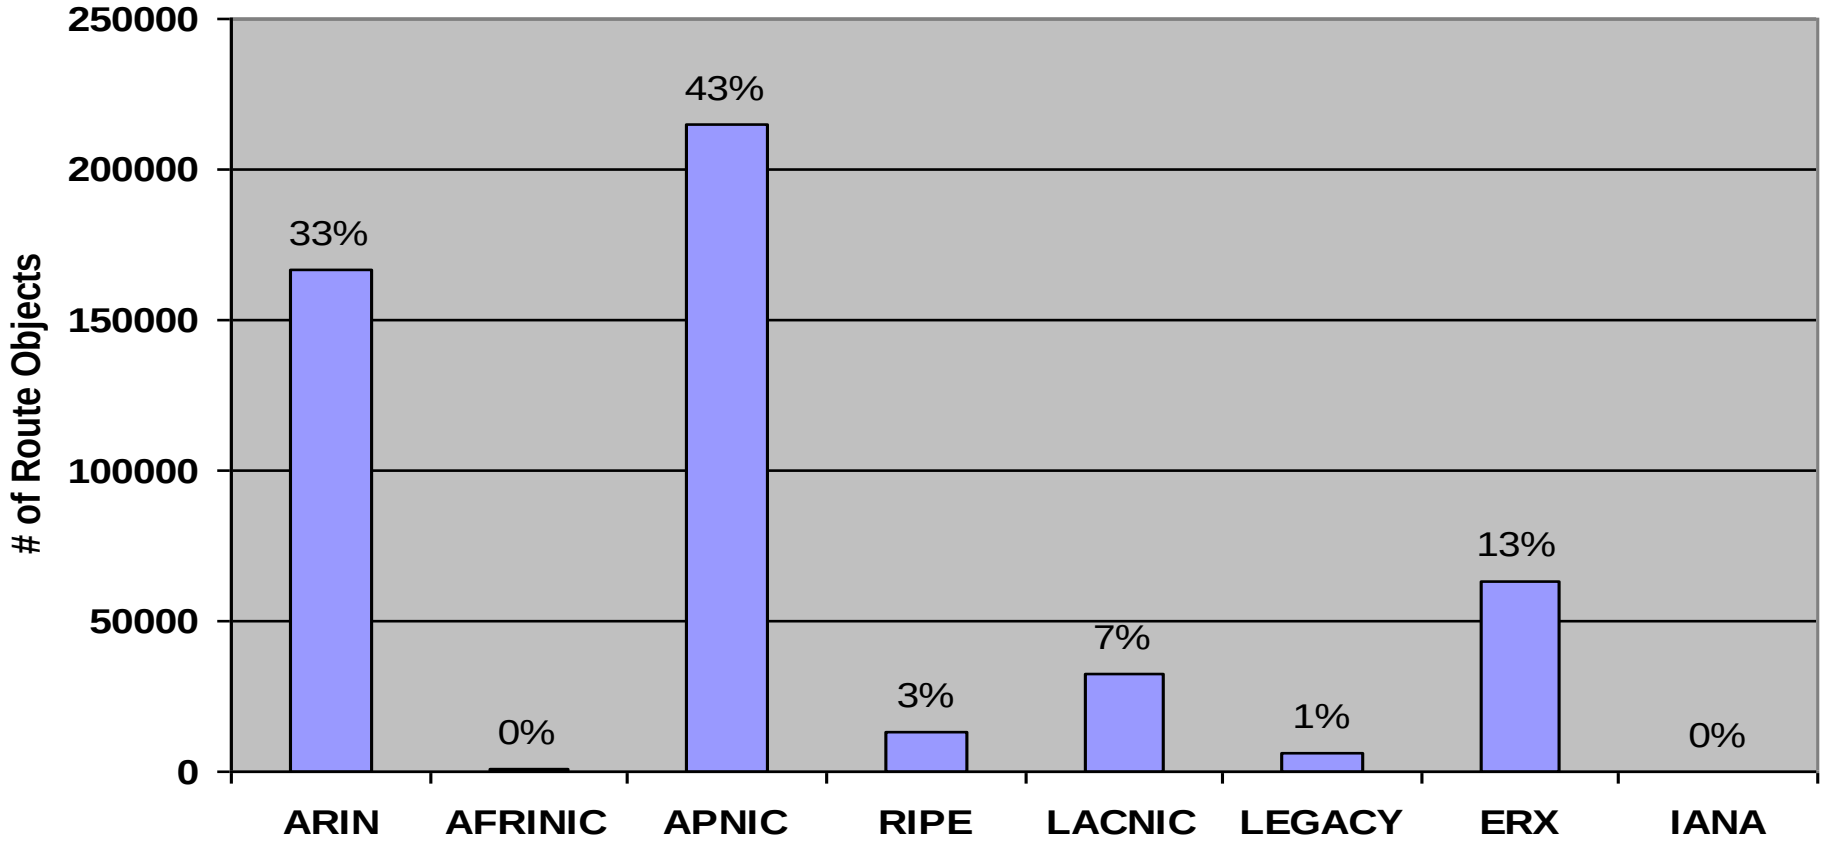

All route objects registered in standalone IRRs on 2008-10-18: 497,124

Growth of NetHandles with OriginAS

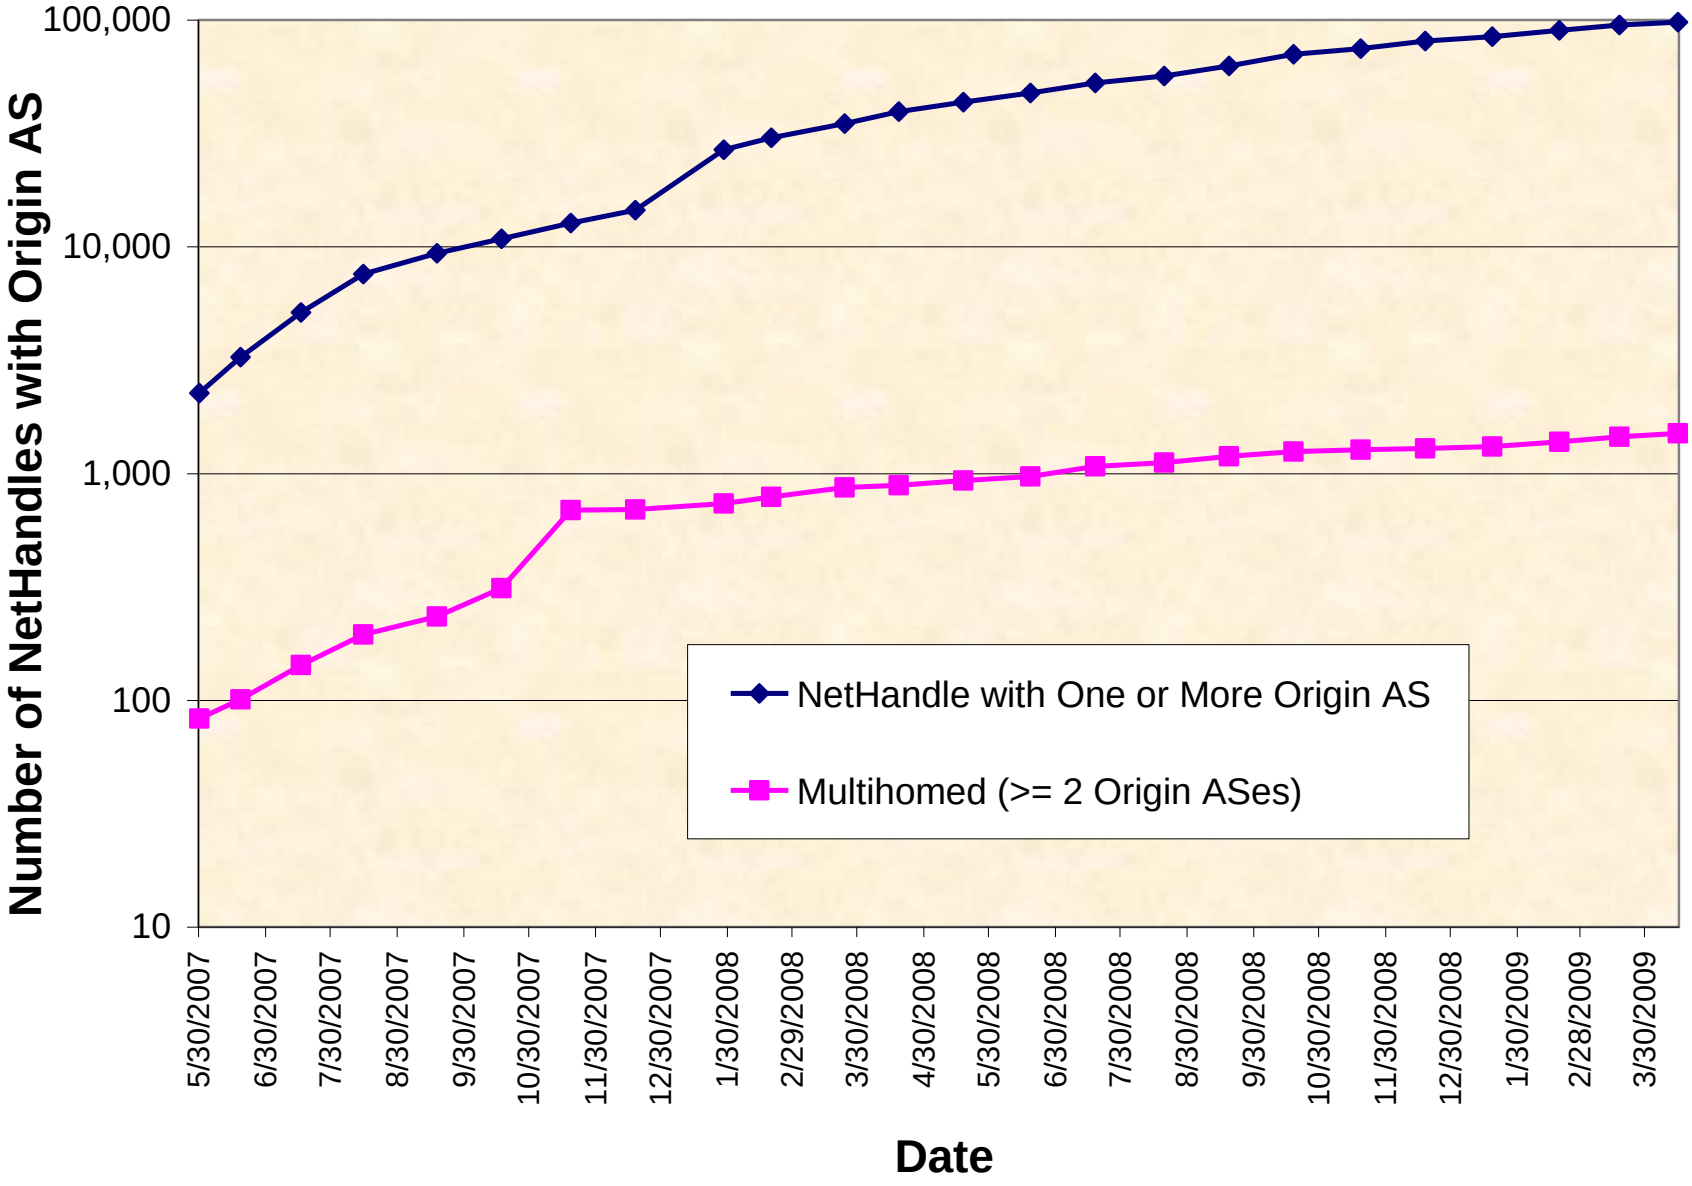

Checking Consistency of a Registered Route with Corresponding Inetnum and Aut-Num

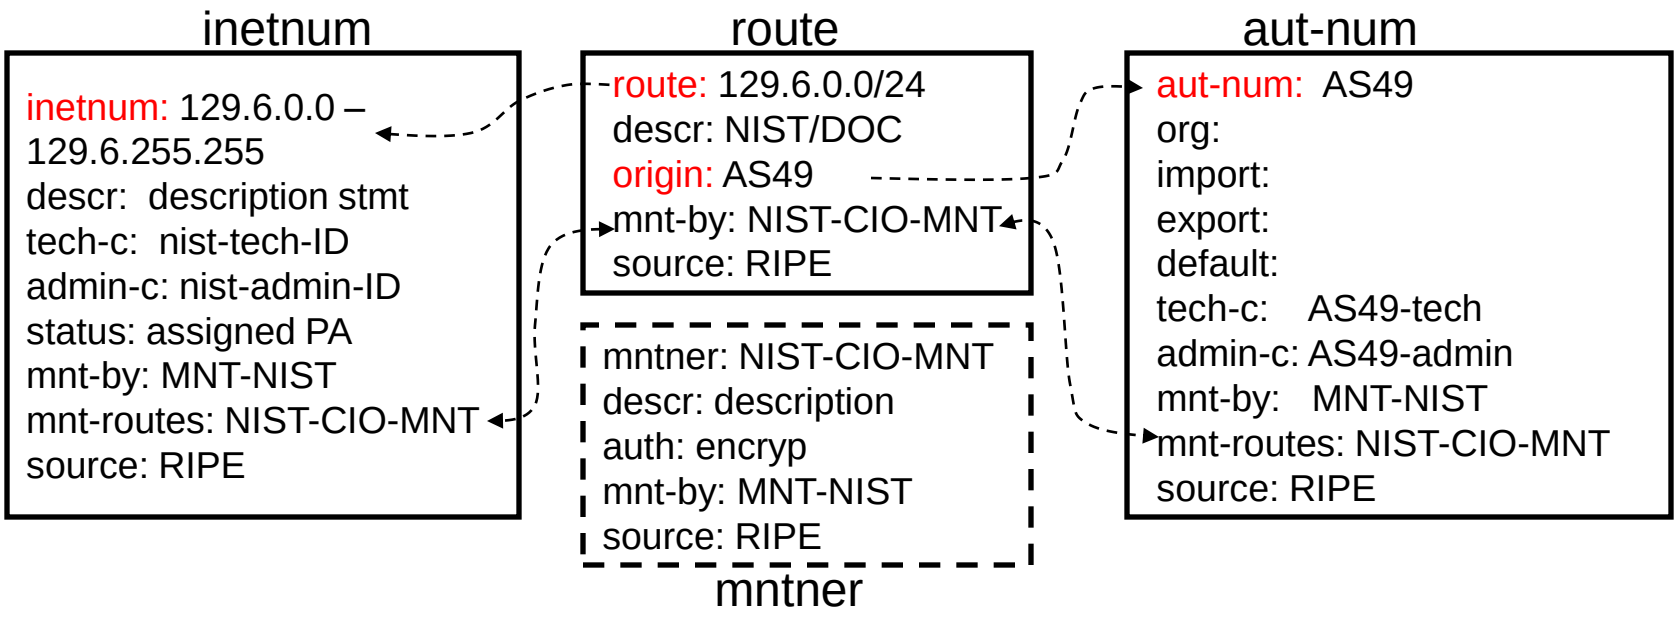

Registry-Based Algorithm for Scoring Routes Observed in Trace Data

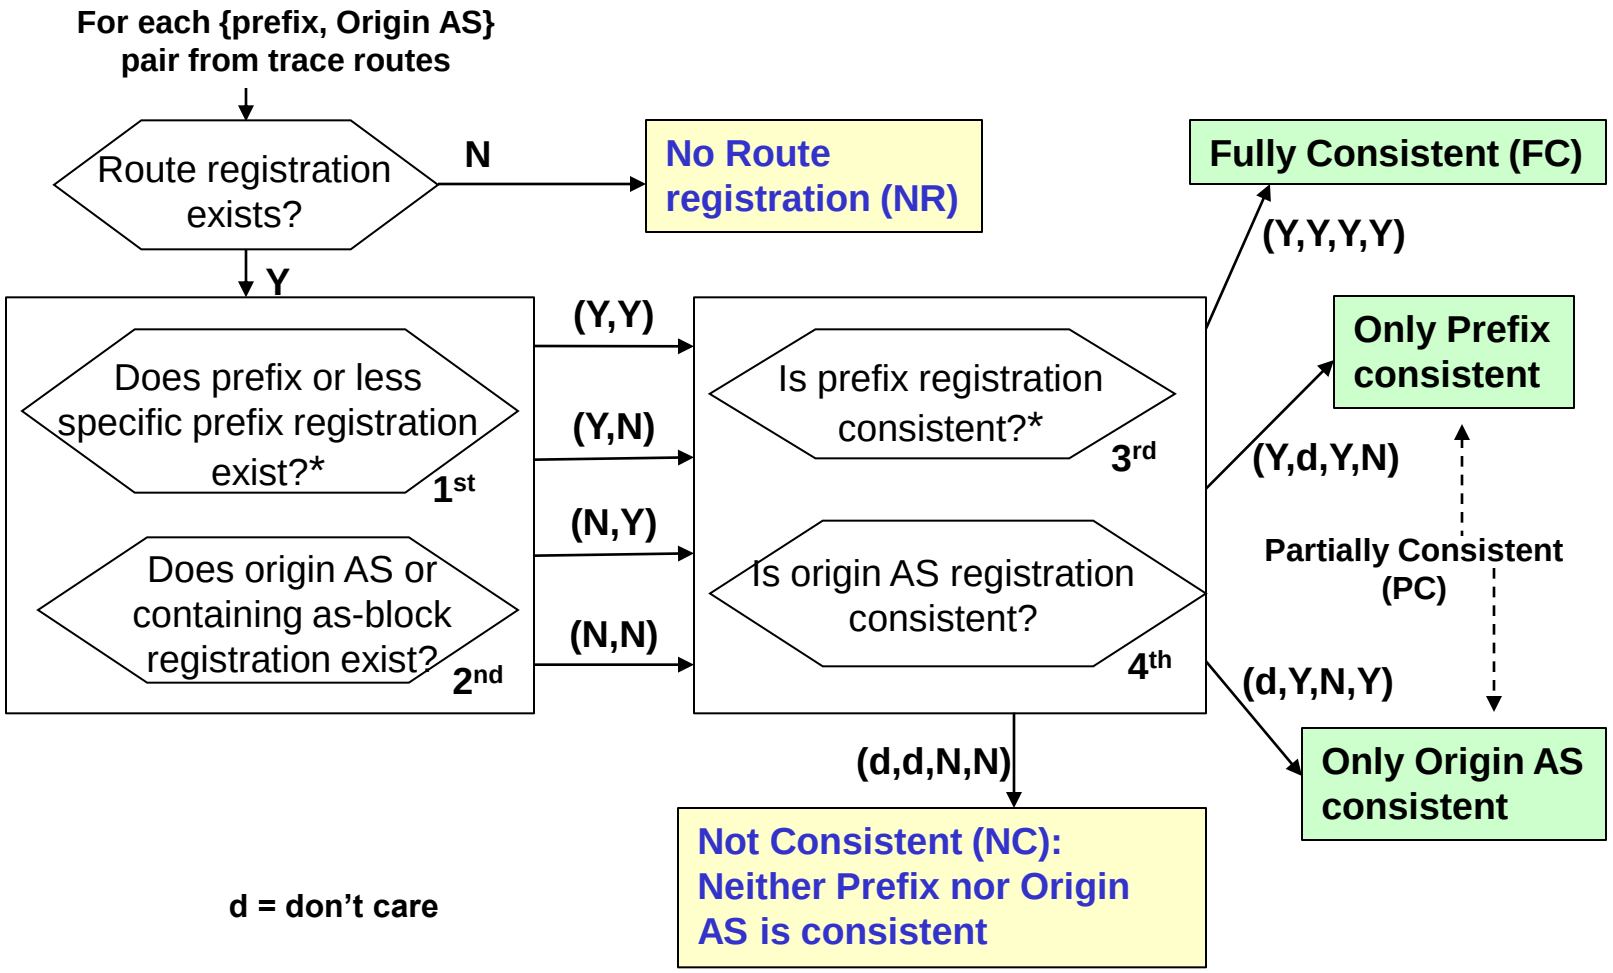

Characterization of IRR Consistency Based on Route Object Registrations

Registry Data Date: 2008-10-18

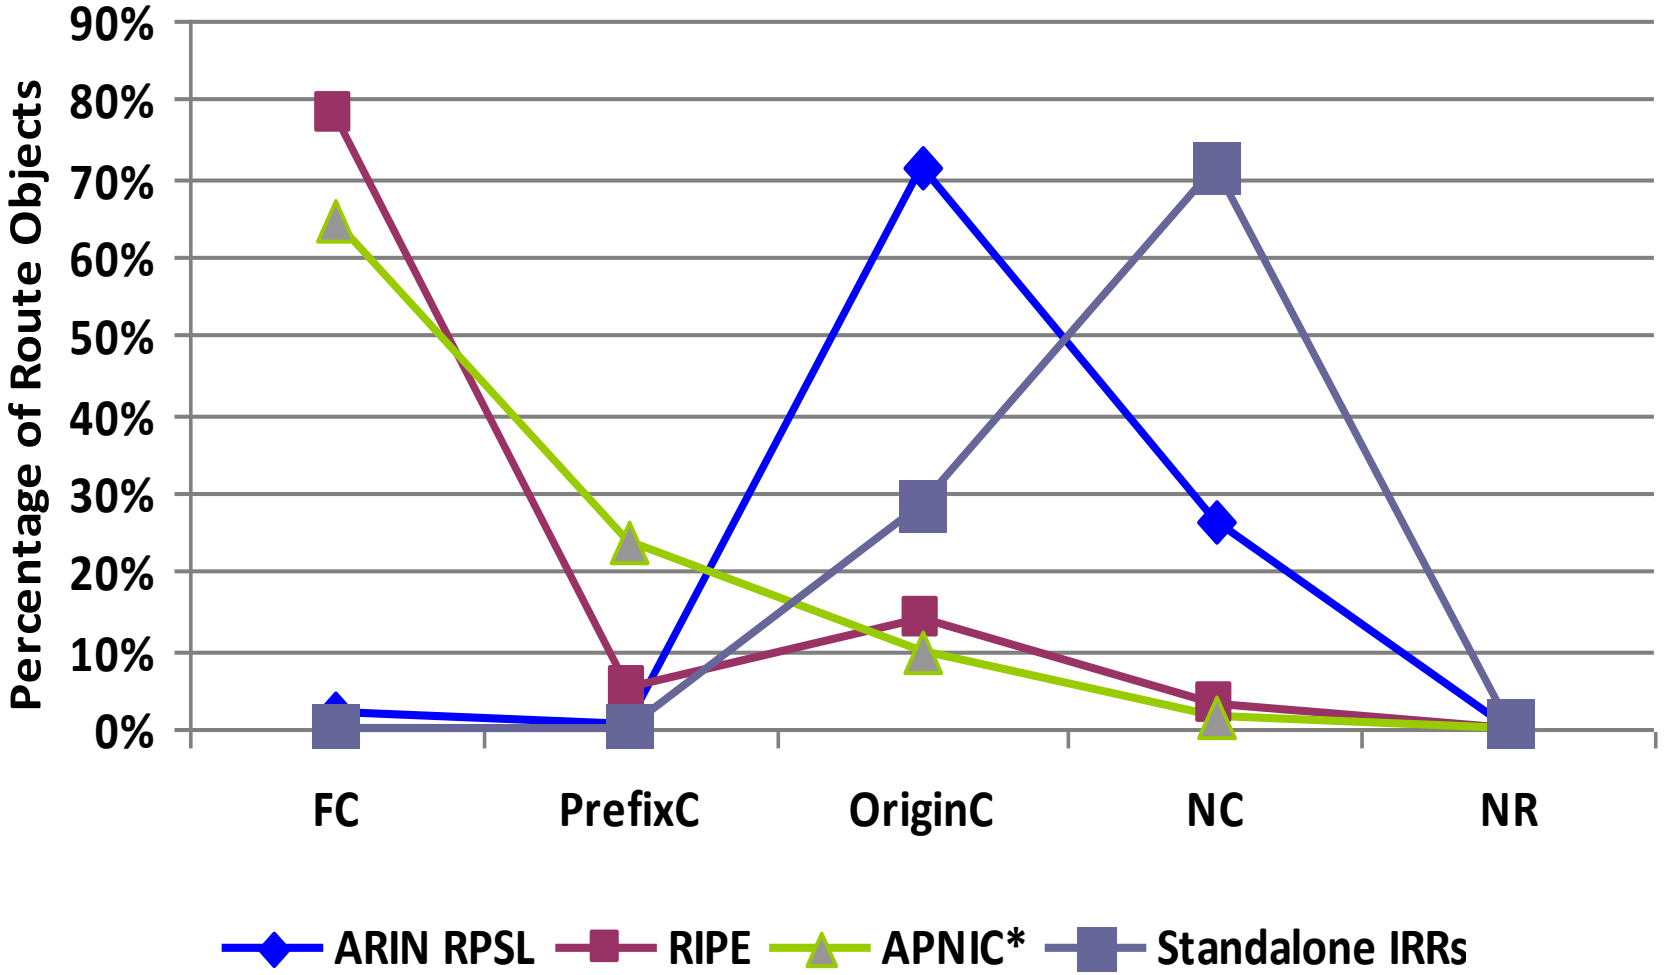

- Note: This does not reflect the total number of routes registered at each IRR. For example ARIN has only 8K whereas RIPE has 90K as of 2008-10-18.

BGP Robustness Algorithms

Known BGP Robustness Algorithms

- General goal: Validate an observed (p, Origin AS) pair
- Nemecis: Compare with registered objects (route, inetnum, autnum)
- PHAS: Compare with historically observed (p, Origin AS) pairs, AS-paths:
  - ➢ Identify origin changes, subprefix announcements; generate alerts
- Pretty Good BGP (PGBGP): Compare with historically observed (p, Origin AS) pairs
  - ➢ Influence forwarding or holding back of updates in real-time in BGP processing

New Integrated Approach

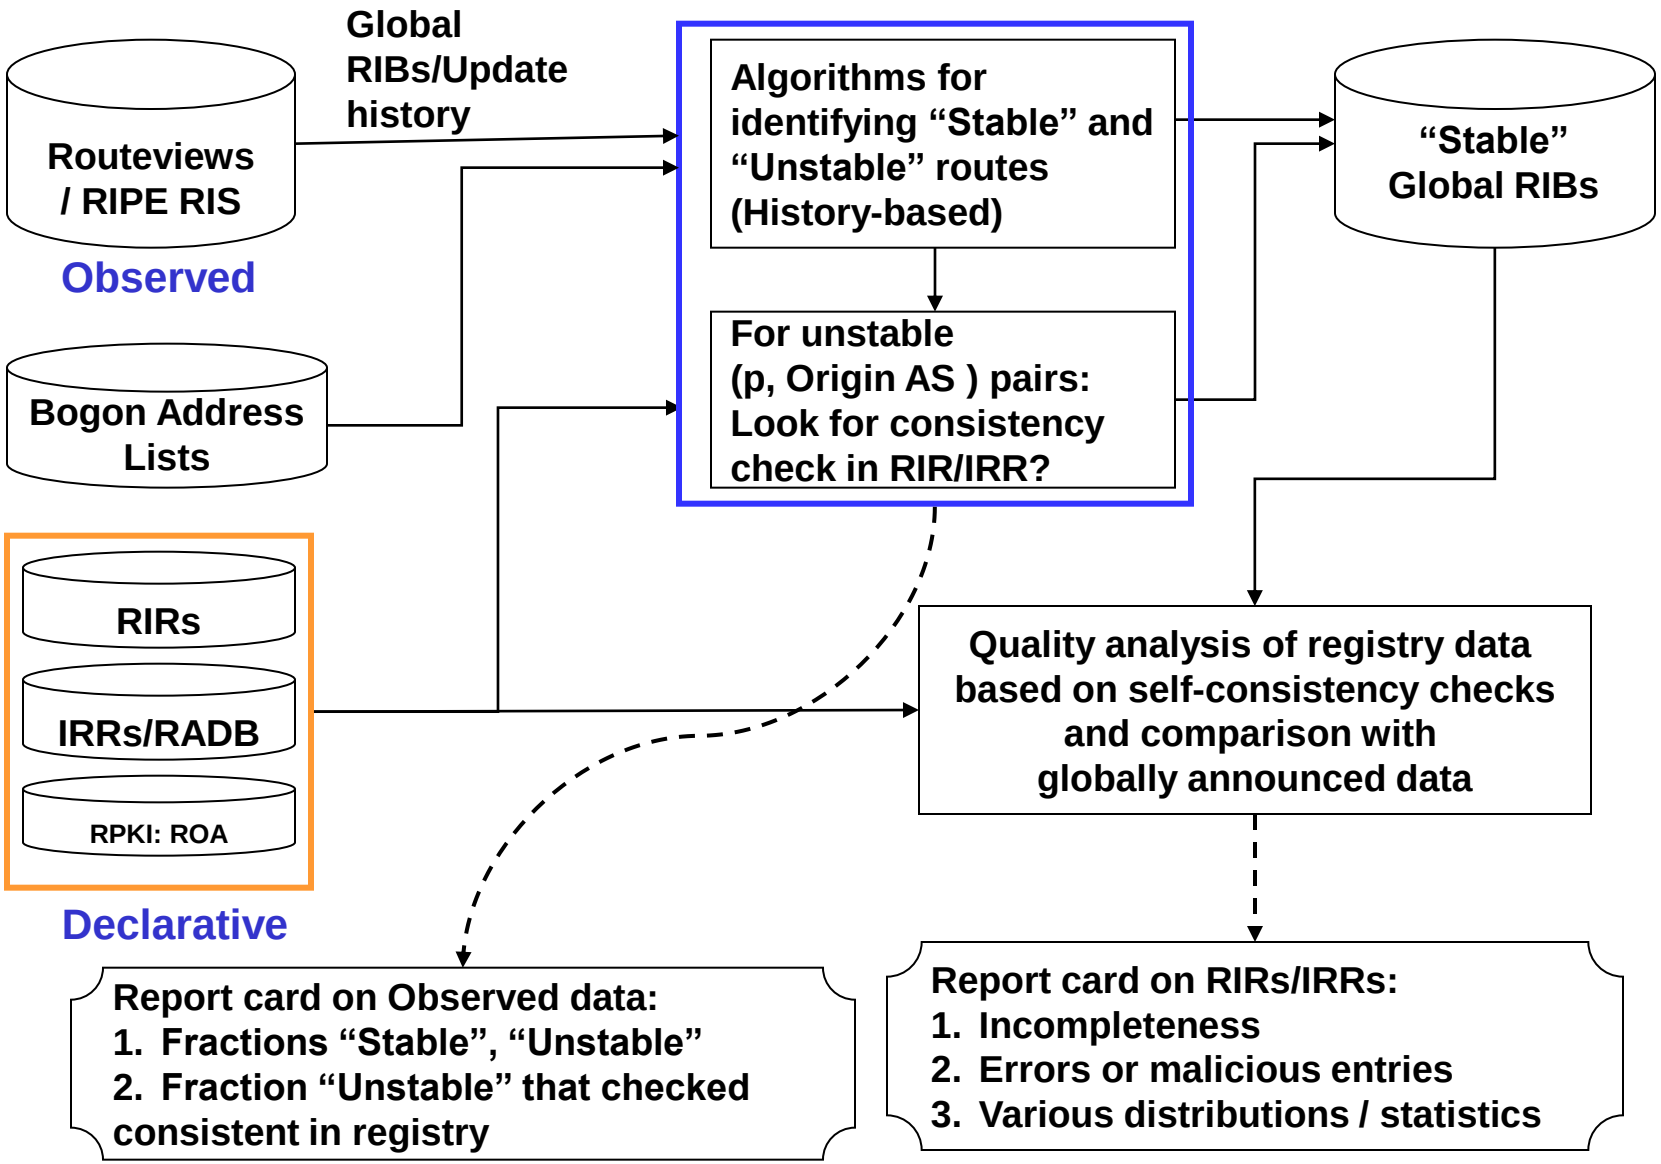

ROA: Route Origin Attestation

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

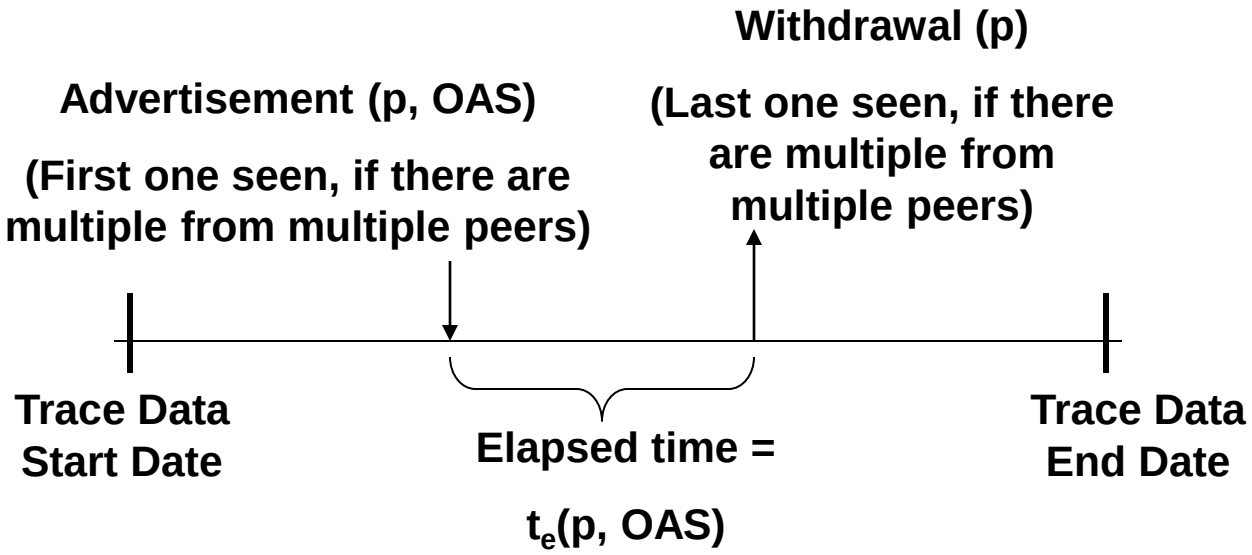

- If t_e(p, OAS) $\geq$ 48 hours, then (p, OAS) is a stable (prefix, Origin AS) pair
- If t_e(p, OAS) < 48 hours, then (p, OAS) is an unstable (prefix, Origin AS) pair
- Update data is initialized with stable (i.e., persistent for $\geq$ 48 hours) RIB entries

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

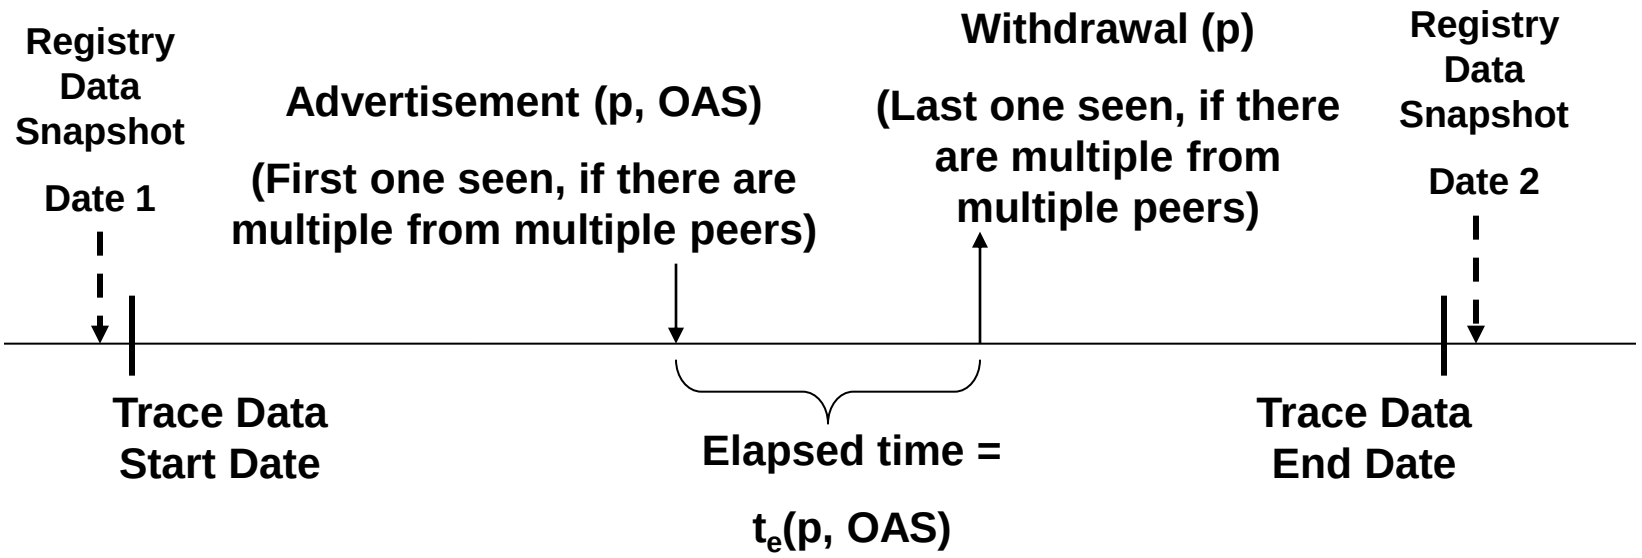

- Use enhanced history-based (i.e., trace-data-based) algorithm as in previous slide
- Complement it with combined results of the registry-based algorithm with data from two dates (close to start and end dates of the history algorithm)
- Result: Better performance of anomaly detection algorithms

Comparative Analysis of Existing and Enhanced Algorithms

- We have encoded Registry-based, Enhanced Trace-data-based and Enhanced Hybrid algorithms for evaluation
- 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
- Results are reported globally for all prefixes as well as selectively for regional (RIPE, ARIN, …) prefixes
- Six-month trace-data window (January through June 2007); initialized with stable RIB entries
- Registry data – two dates prior to and towards the end of the six-month window (December 12, 2006 and June 18, 2007)

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

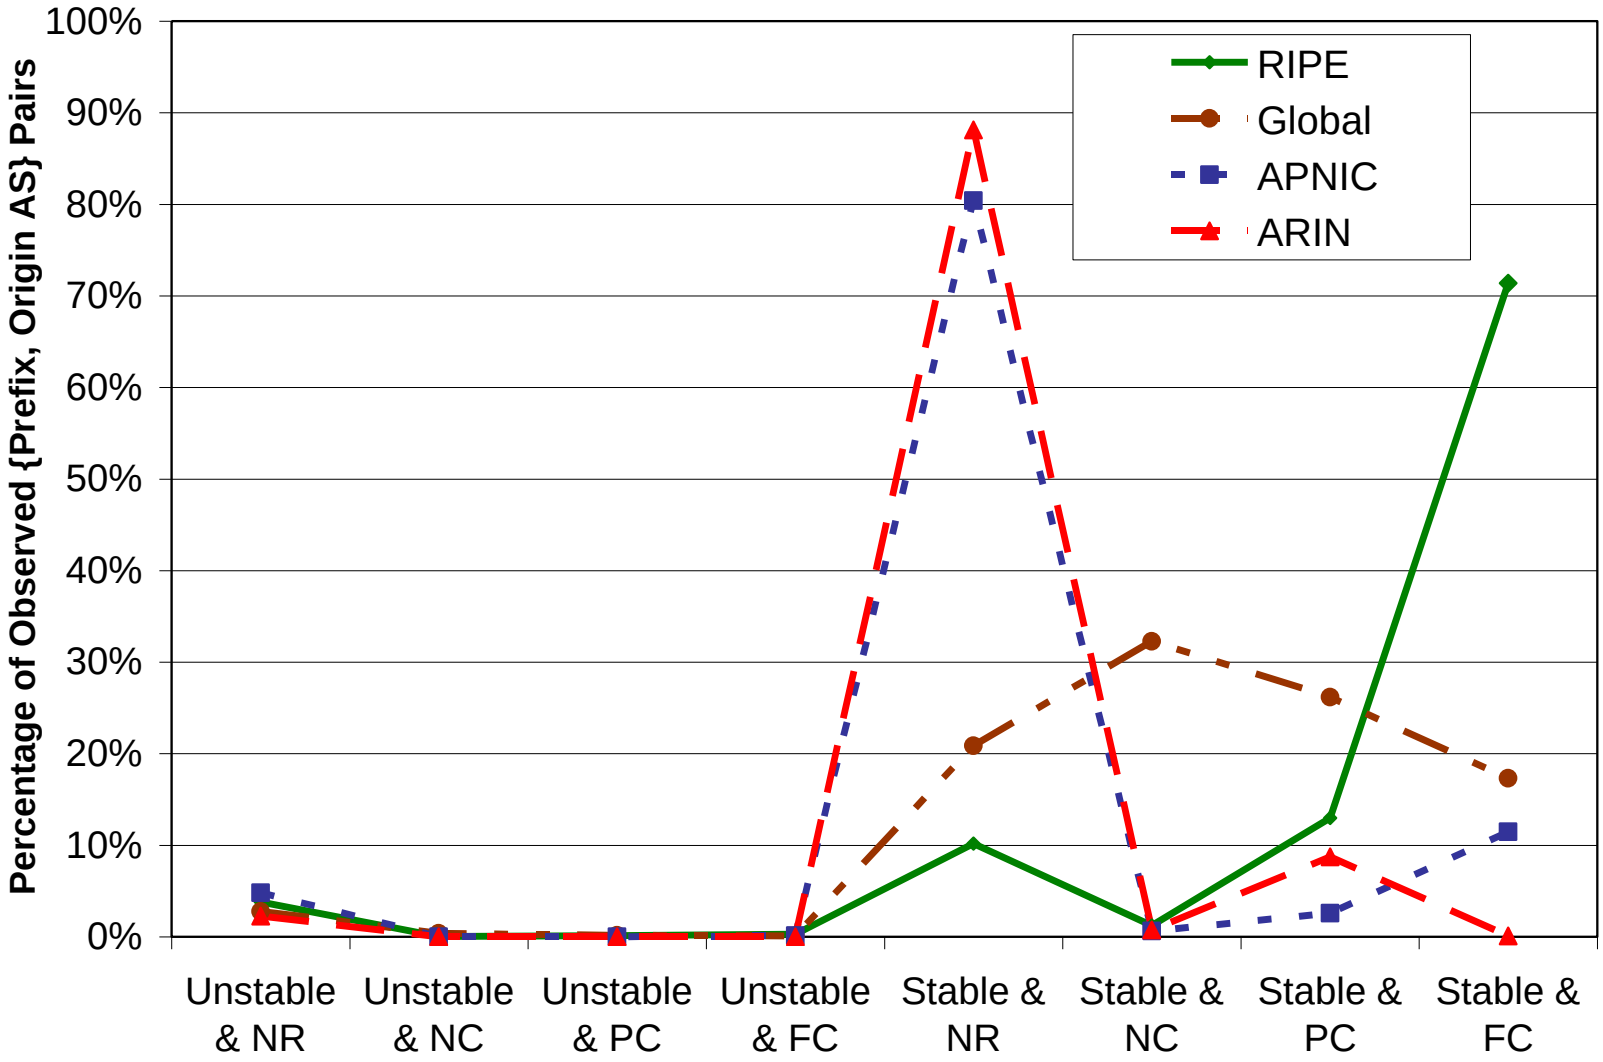

p = prefix; OAS = Origin AS; FC = Fully Consistent; PC = Partially Consistent; NC = Not Consistent; NR = Not Registered

Comparative Performance of Algorithms

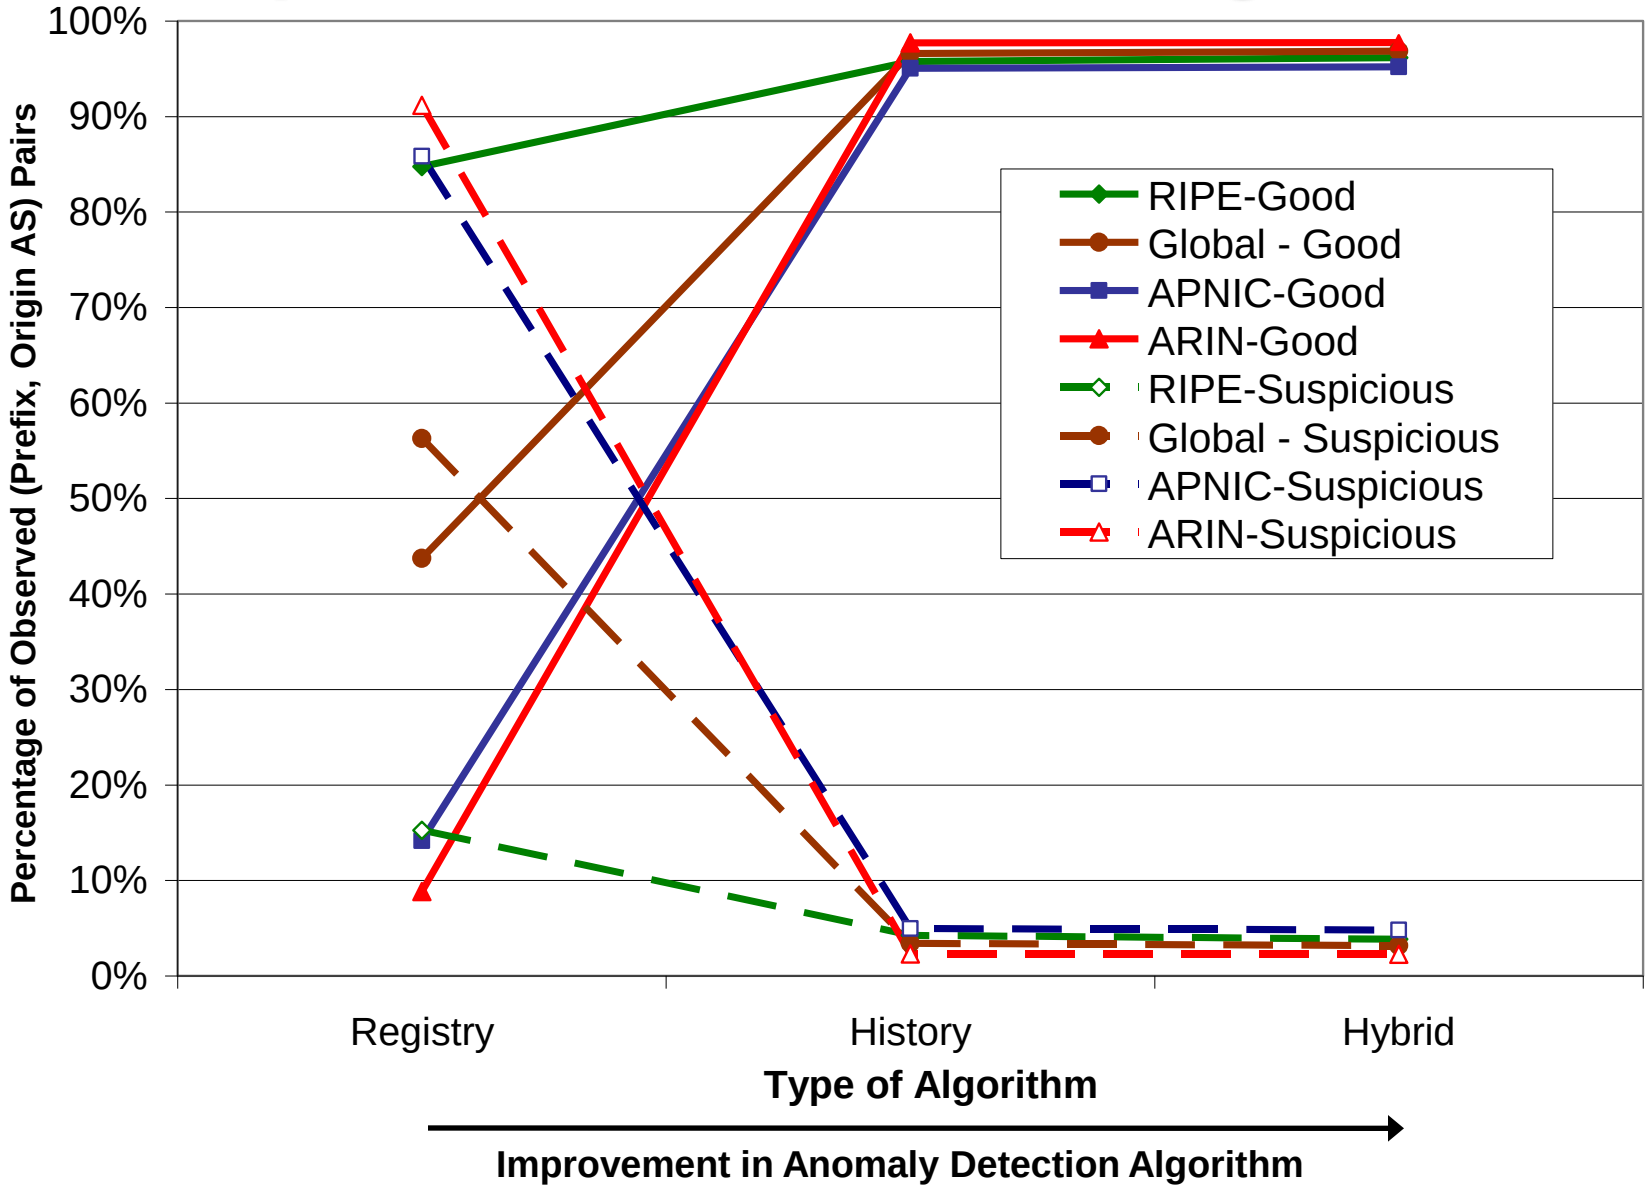

Comparative Performance of Algorithms

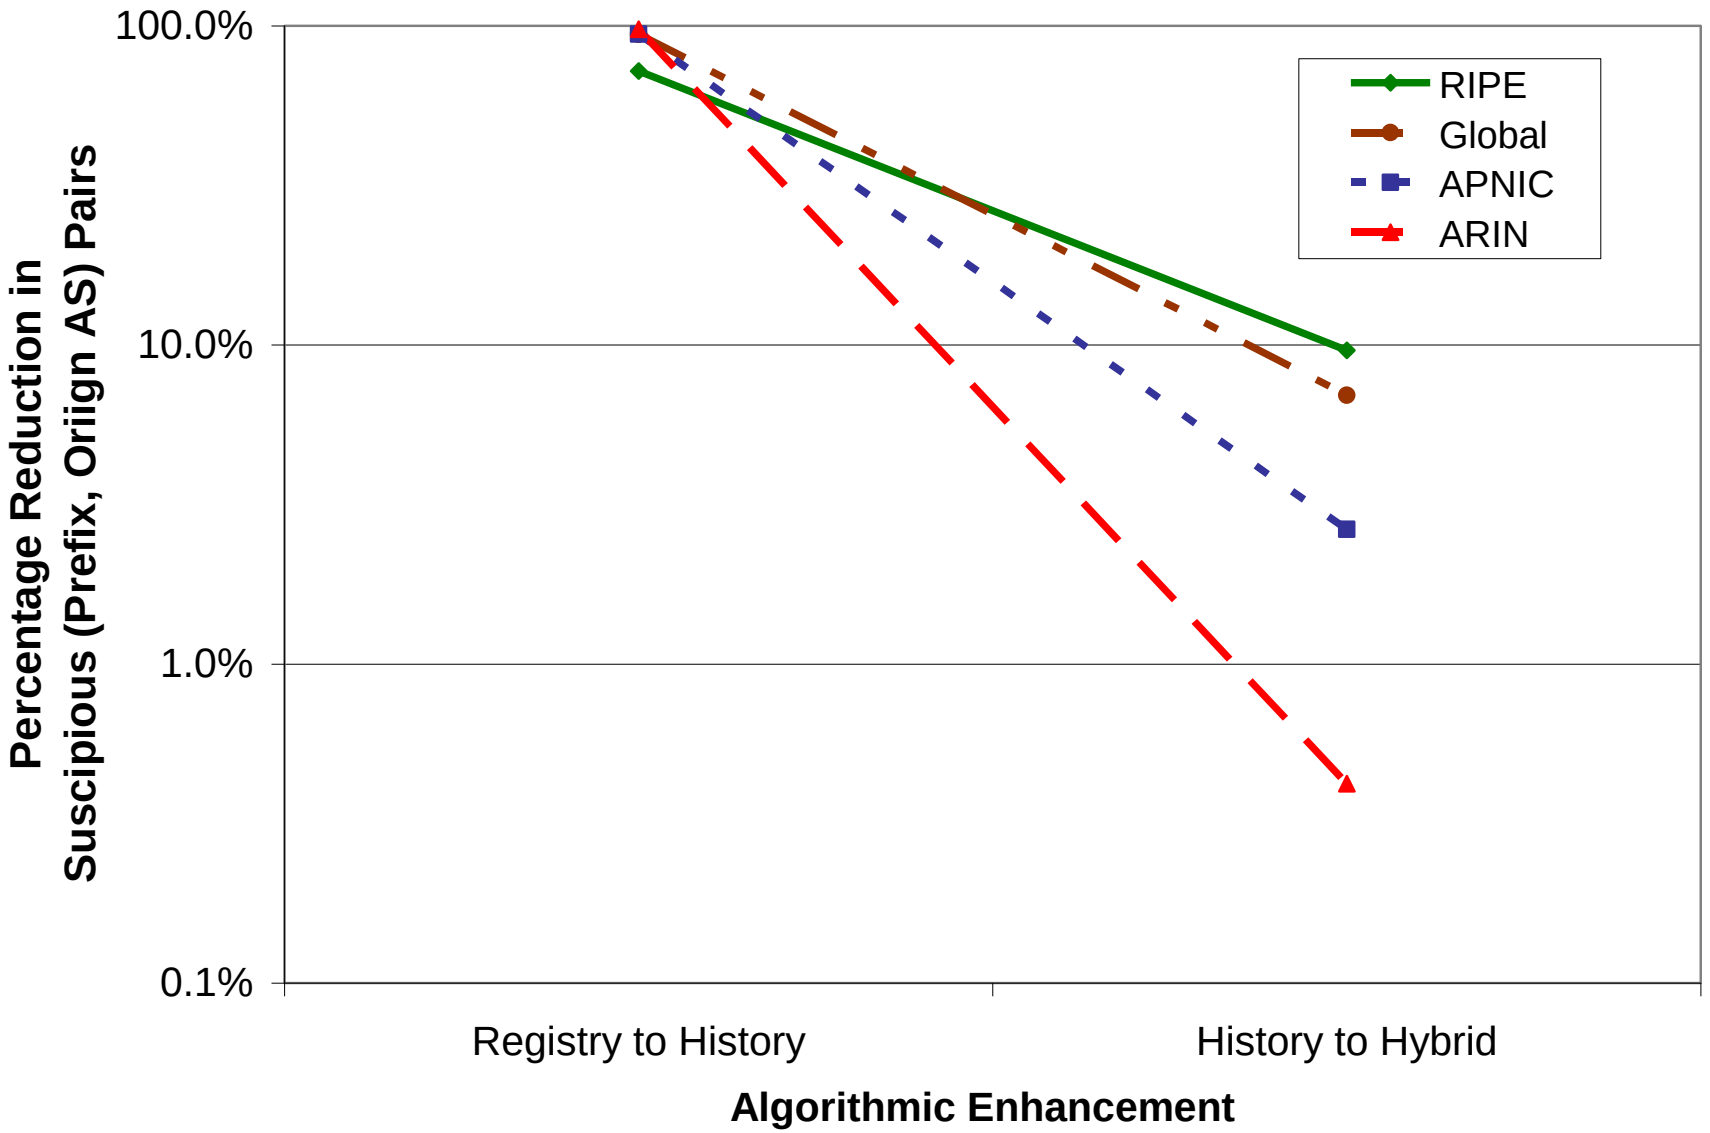

Checking Origin AS : Comparison of Algorithms

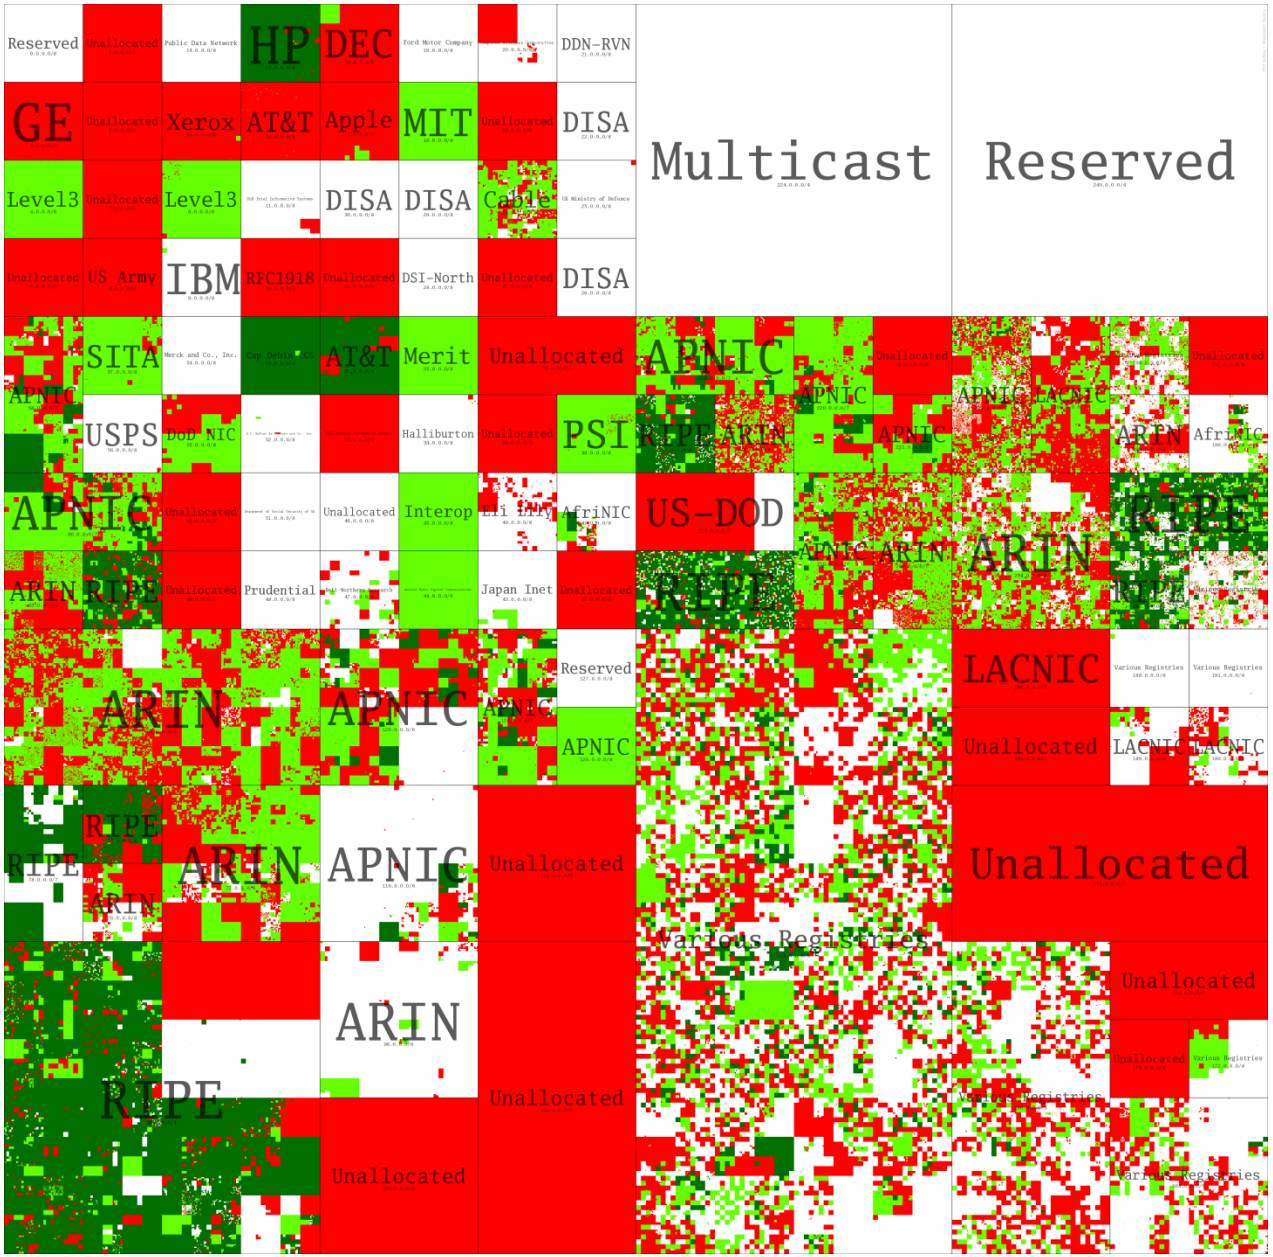

Registry-based Algorithm

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

White: Not found in trace data

Checking Origin AS : Comparison of Algorithms

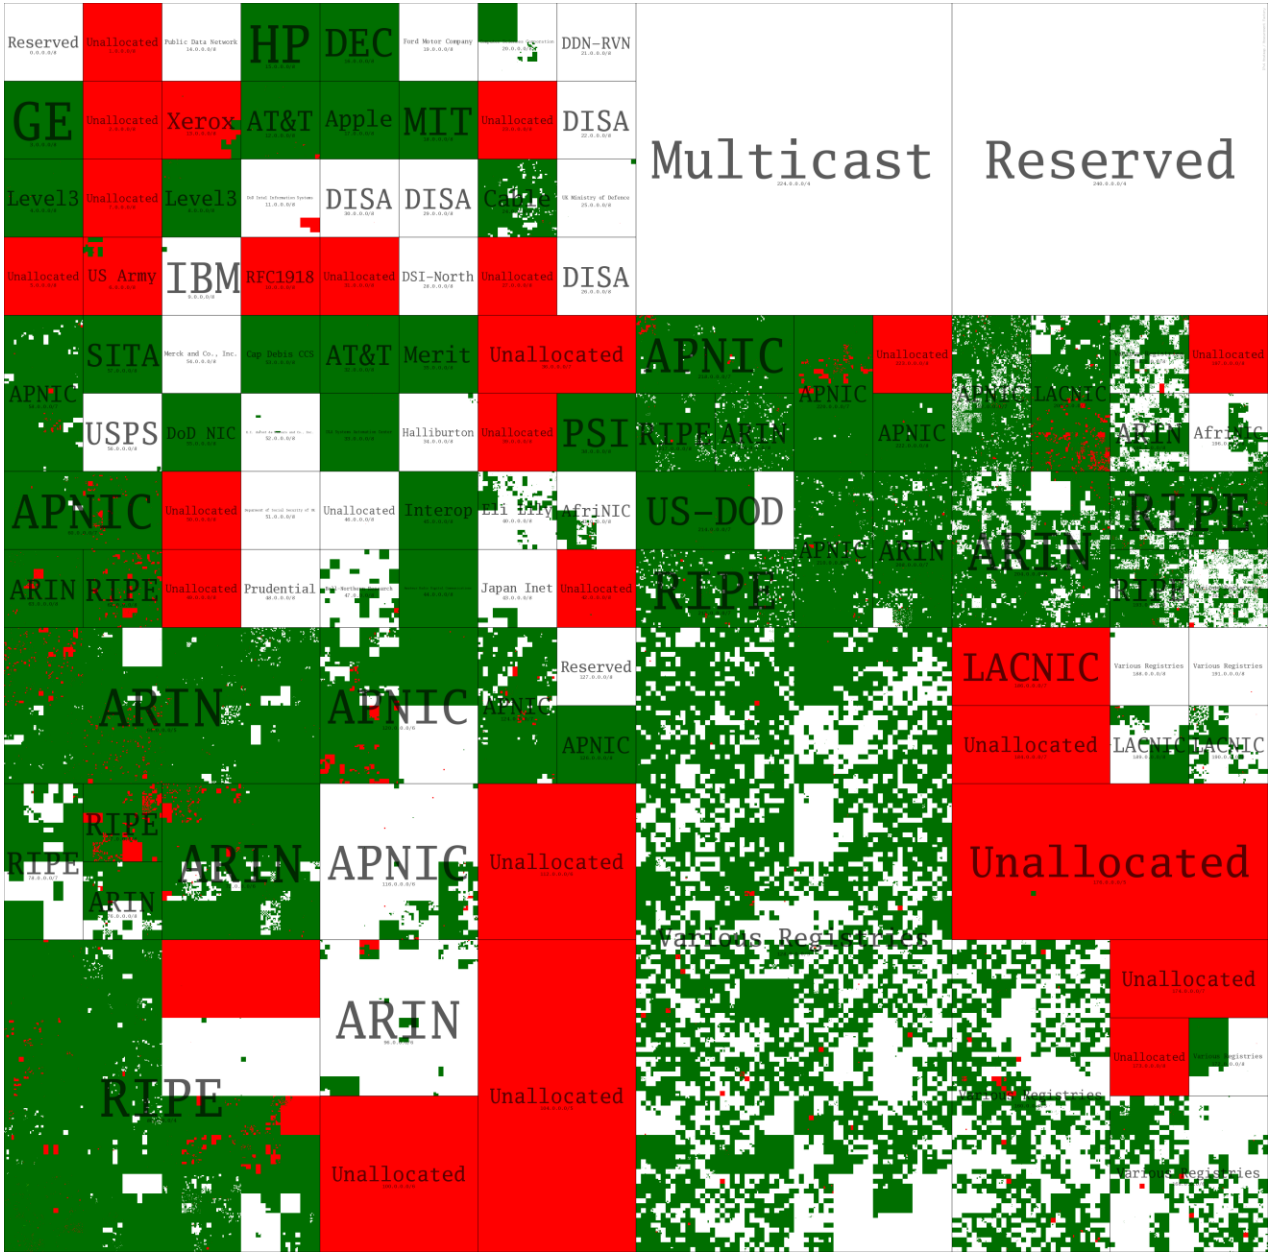

Enhanced trace-data-based Algorithm

Green: Good

Red: Suspicious

White: Not found in trace data

Checking Origin AS : Comparison of Algorithms

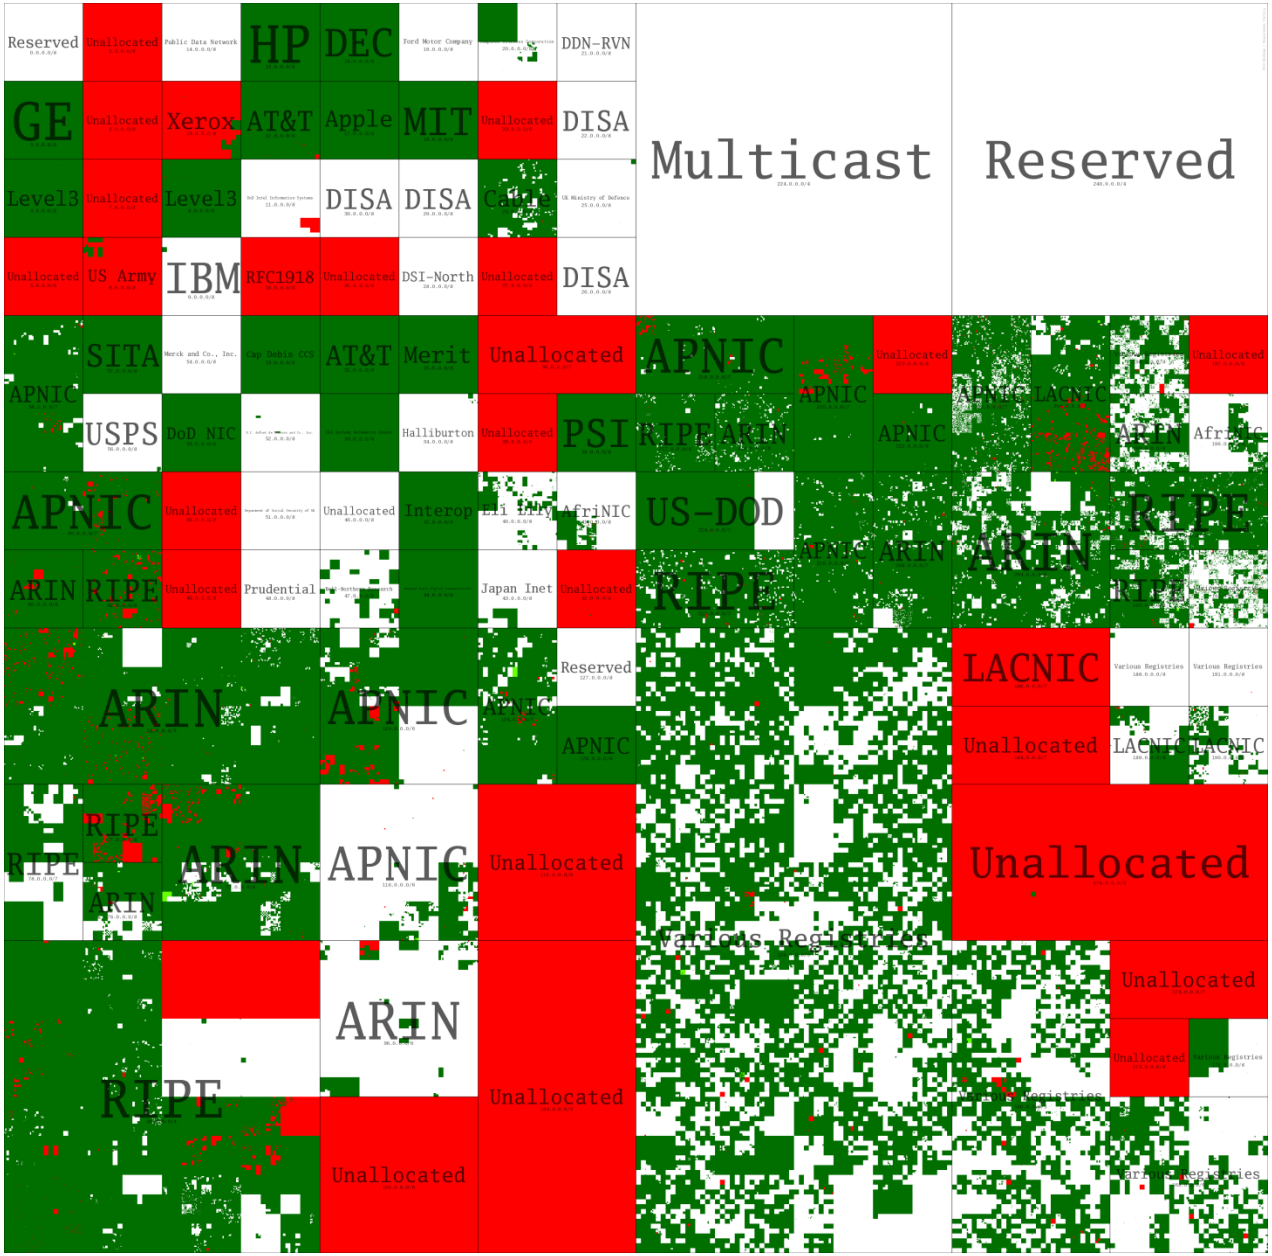

Enhanced Hybrid Algorithm

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

White: Not found in trace data

Summary

- Examined and quantified the quality (completeness, correctness) of registry data
- Enhanced hybrid algorithm – history and registry data have complementary influence on improvement in origin validation
- Further testing for robustness of the algorithms needs to be performed with extensive real and synthetic trace data
- NIST has begun to monitor and quantify the growth and quality of the RPKI data

Backup slides

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?

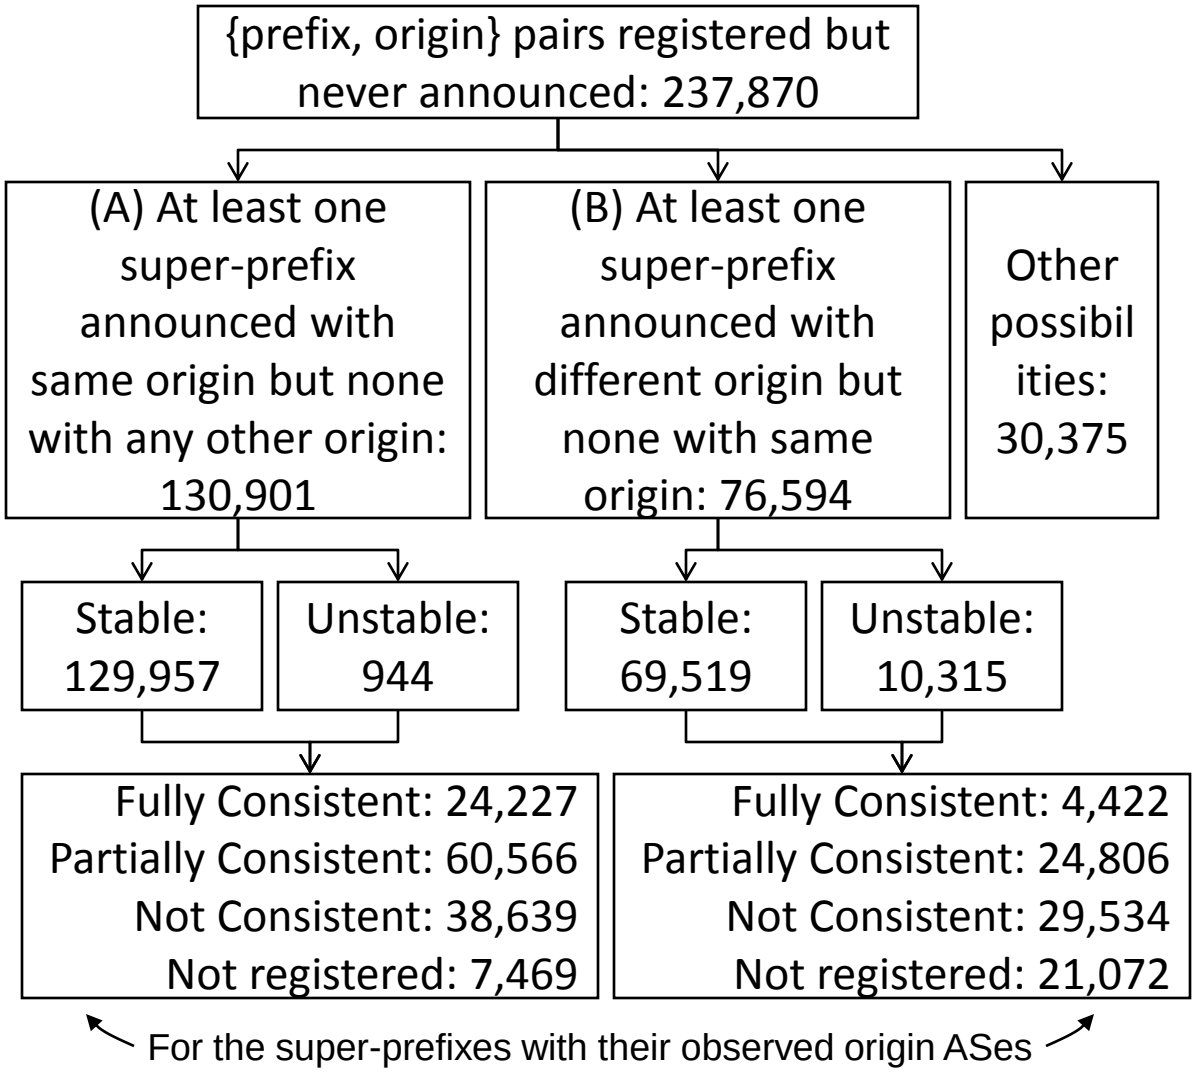

Nemecis: Registry Based Algorithm

- 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

- Make use of BGP trace data
- 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.

Checking Registry Consistency of Registered Routes (Algorithm)

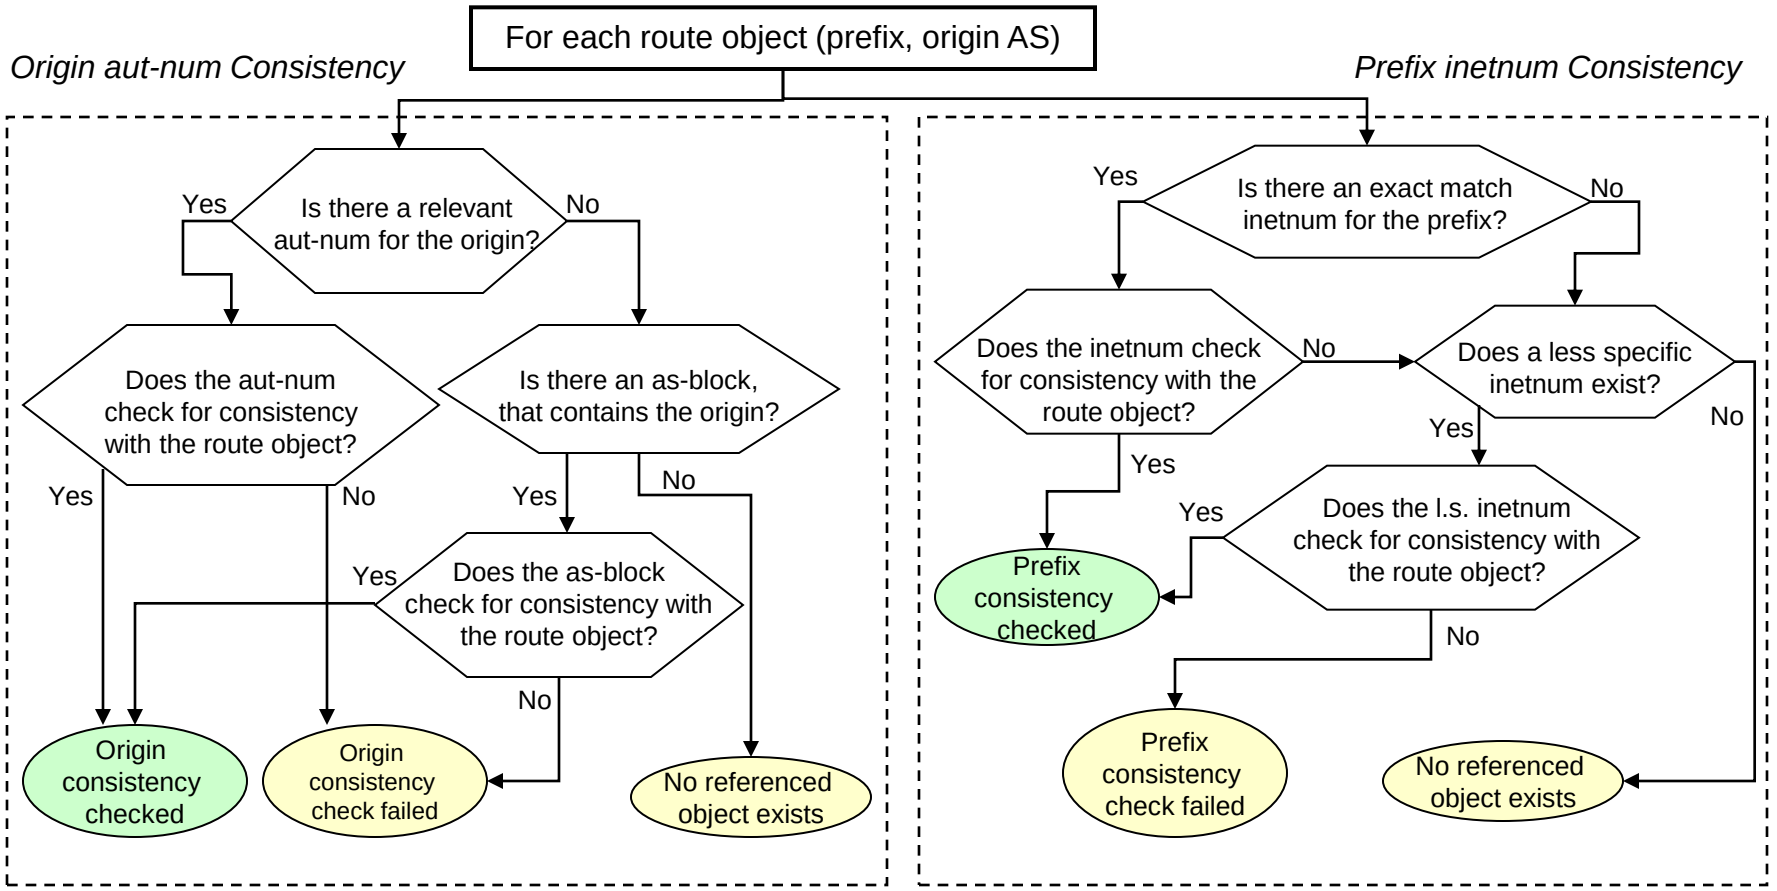

Origin AS Approval Check List: Comparison

		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.**

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 based on RPKI efforts (in future)

Heatmap Depicting Origin Validation for Announced Prefixes

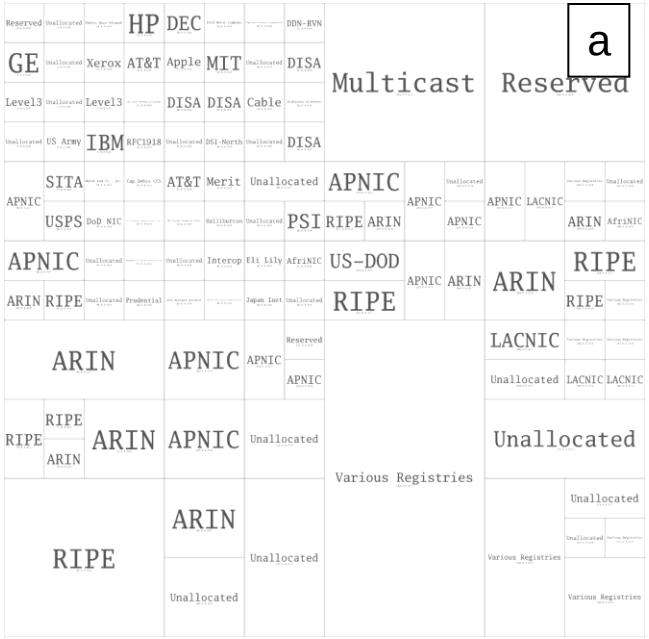

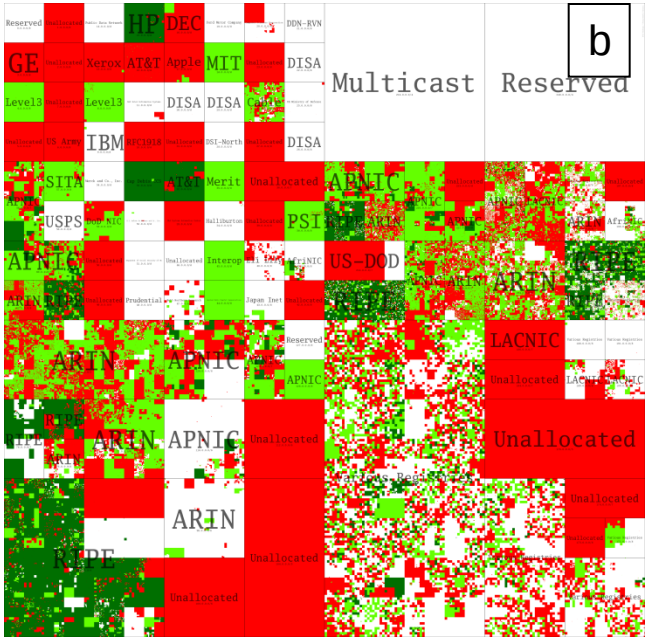

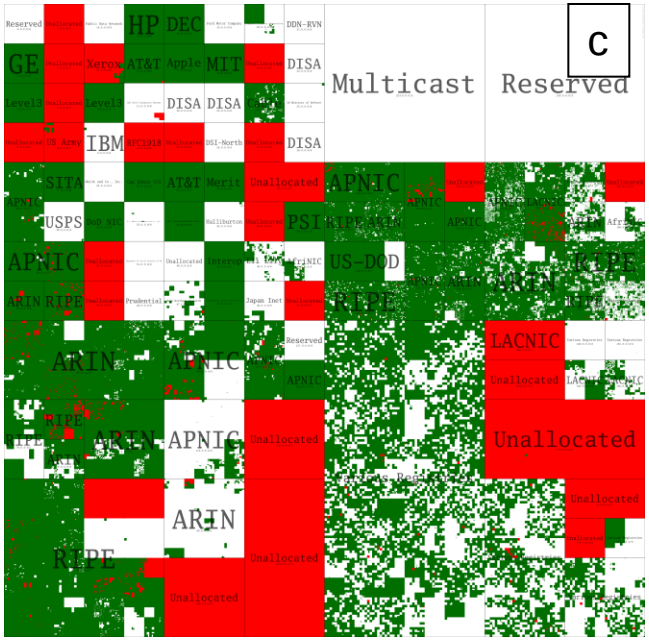

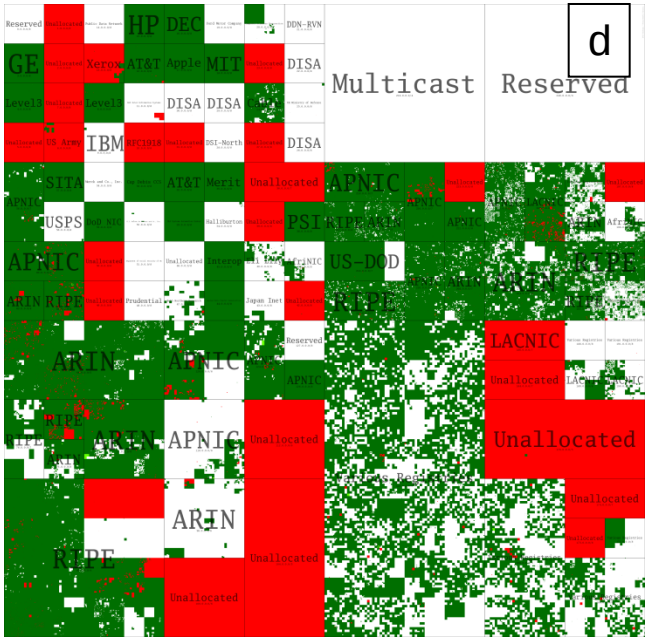

a. **Allocations**
b. **Registry-based Algorithm**
c. **Enhanced Trace-data-based Algorithm**
d. **Enhanced Hybrid Algorithm**

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