

Analyzing the effectiveness of the CIDR Report

Stephen Woodrow

MIT CSAIL

Outline

- History of the CIDR Report
- Analytical approach
- Results
- Observations about the Report
- Conclusions


What this talk is and is not

- This is specifically about the CIDR Report in the context of deaggregation
- Deaggregation and its causes have been considered before:
 - RAS' "Inconvenient Prefix", NANOG 50
 - Cittadini et al. "Evolution of Internet Address Space Deaggregation: Myths and Reality", IEEE JSAC October 2010
 - ripe-399 RIPE WG on Route Aggregation

History of the CIDR Report

- Initially a tool to promote CIDR aggregation following BGP4 deployment in early-mid 1990s
- Earliest “top 10” list on IETF CIDRD list in 1994
- Appeared on NANOG list in September 1996, and soon evolved into its current format
- Initially launched by Tony Bates & Philip Smith, later transitioning to Geoff Huston in 2002

Initially successful? Down and to the right



Why did it (presumably) work?

- Providing information to the clueless
- Social forces
 - Shame and reputation
 - Peer pressure
 - The CIDR Police (NANOG 27)

Is the CIDR Report still effective?

- “Not anymore, but it worked \$N years ago ”
from several people I spoke to
- “Doesn’t matter; Routing table growth isn’t a
problem anymore, thanks to Moore’s law”
- No longer needed to promote classful-to-
classless transition
- **No empirical study?**

Defining “effective”

- Hypothesis: If the CIDR Report “works”, it should encourage aggregation in order to reduce one’s ranking. Thus:

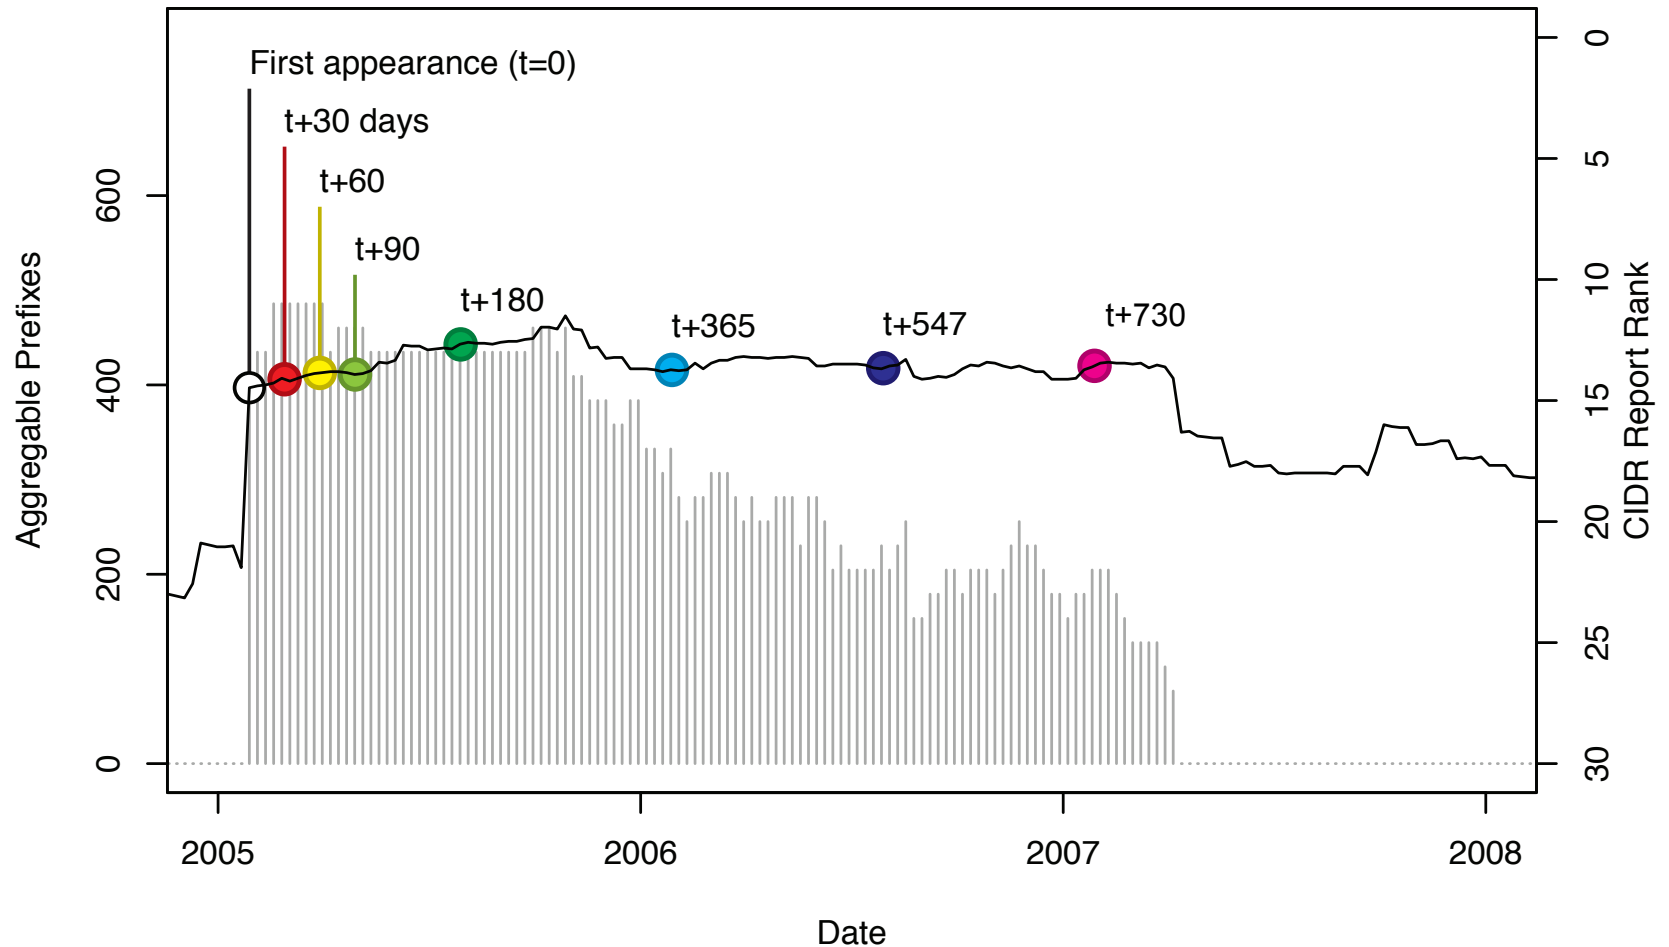
Does an AS improve its aggregation behavior after appearing on the CIDR Report?

- Aggregation behavior in terms of:
 - Deaggregation factor (netsnow/netsaggr)
 - Aggregable prefixes (netgain)

Measuring AS behavior

- Archived CIDR Reports are not sufficient
 - Relative ranking—dropping off top 30 is ambiguous: did AS improve, or did others become worse?
- Re-implementing the CIDR Report
 - Data: Route Views RIBs
 - Preprocess to canonicalize AS_PATH
 - Mark prefix aggregable if any RV peer can aggregate
 - Some differences, but generally consistent

Measuring change in AS behavior



Measuring change in AS behavior

- Locate each AS' first appearance on the Report
- Sample at various periods after appearance
- Establish control group of ASes
 - Same number as in treatment group
 - Randomly selected from ASes never appearing
 - Must advertise at least 10 prefixes
 - Must be visible for full sampling period


Measures of deaggregation

- Deaggregation factor (DF)

$$\frac{\text{actual number of prefixes advertised by an AS}}{\text{minimum number of prefixes required for AS' routing policy}}$$


- netgain
 - the number of prefixes advertised by an AS that could be withdrawn without altering routing policy

Deaggregation factor over time




Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 1998 – 31 Dec 2001

Deaggregation factor over time




Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 1998 – 31 Dec 2001

Deaggregation factor over time




Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 1998 – 31 Dec 2001

Deaggregation factor over time




Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 1998 – 31 Dec 2001

Deaggregation factor over time




Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 1998 – 31 Dec 2001

Deaggregation factor over time




Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 1998 – 31 Dec 2001

At first, the top 30 mostly improve, while the control group deaggregates




Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 1998 – 31 Dec 2001

Over time, aggregation improvement by top 30 decreases




Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 2001 – 31 Dec 2003

Over time, aggregation improvement by top 30 decreases



Deaggregation factor ($\text{netsnow} / (\text{netsnow} - \text{netgain})$)
1 Jan 2004 – 31 Dec 2006

Over time, aggregation improvement by top 30 decreases



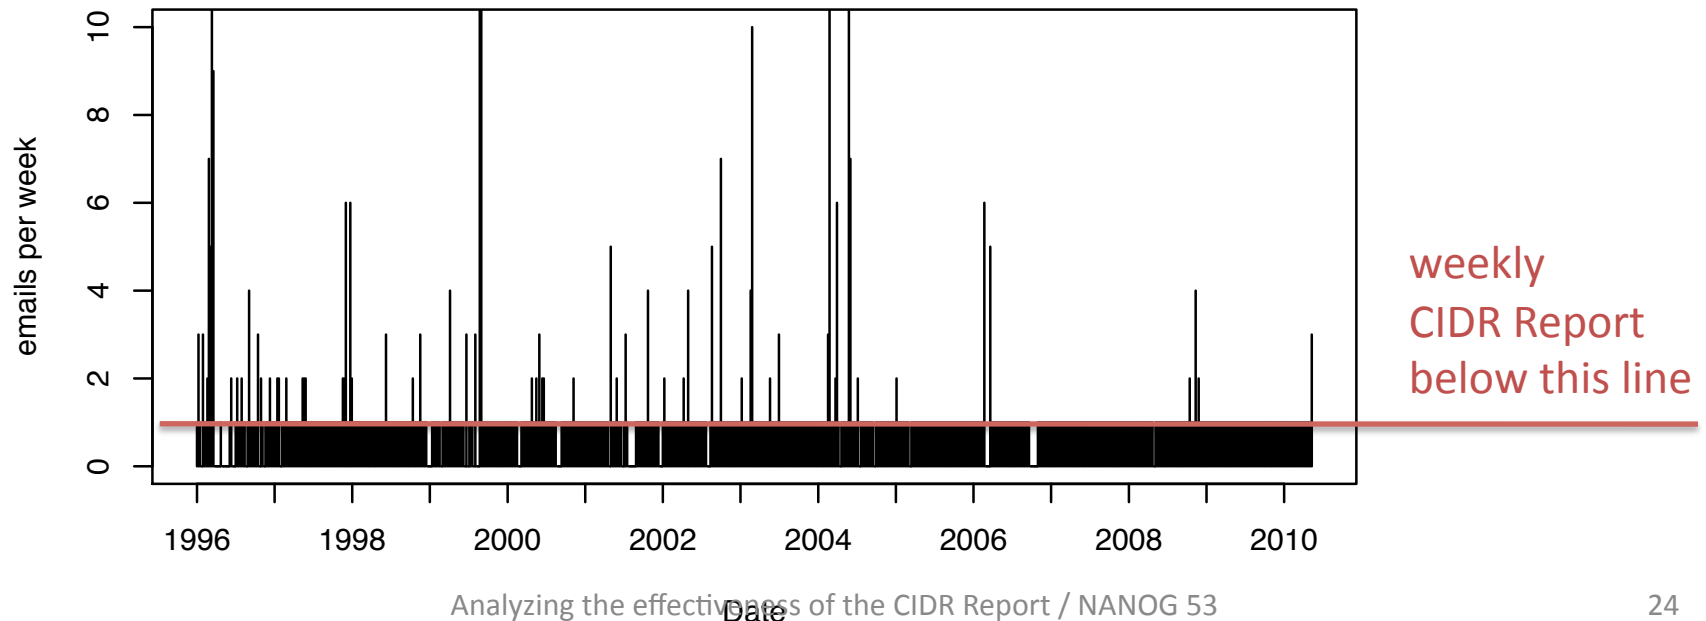
Deaggregation factor (netsnow / (netsnow – netgain))
1 Jan 2007 – 31 Dec 2009

Interpreting these data

- Control groups are consistent over time
- Improvement in deaggregation factor has decreased over time
- **If CIDR Report was source of previous behavior change, it is no longer effective**
- Potential confounding factors
 - Selection Bias: Top 30 ASes may be intrinsically different than other networks
 - Regression to the mean

Why the change? Theories:


- RIB/FIB slots are plentiful — less pain/concern
- Increased need for deaggregation (TE, etc.)
- Change in community response



Observations about the Report


- CIDR Report was initially developed to encourage CIDR-ization
- **It has remained essentially unchanged since, while the Internet has grown and changed**

With fixed “top 30”, the Report focuses on outlier behavior as Internet grows



Threshold lines indicating the cut-off point for appearing on the CIDR Report in 1998 and 2011 are indicated. Note that the graph is rescaled; approximately 70% of ASes in the routing table do not advertise any aggregable prefixes.

The top of the Report contains ASes with extreme netgains: hard to displace



Should we do something about the CIDR Report?

- Unclear whether routing table growth is a pain point right now or in the near future
- CIDR Report would probably not be the best way to deal with painful growth
 - FIB aggregation/compression
 - New interdomain routing architecture (RFC 6115)
- But, if it's something operators care about, it could probably be improved
 - e.g. focus on more actionable behavior: “low-hanging fruit” or recent offenders

Conclusions

- The CIDR Report appeared effective in its early days (though it's difficult to establish causality)
- There are a number of potential contributing factors for its decreased efficacy over time
- The CIDR Report has not evolved with the Internet, but it probably could be improved
- Social forces in the operator community may yet have value in solving problems



More detail: <http://people.csail.mit.edu/woodrow/sm-thesis.pdf>

woodrow@csail.mit.edu