"SASUKE" Traffic Monitoring Tool
Traffic Shift Monitoring Based on Correlation between BGP Messages and Flow Data

Atsushi Kobayashi
Yutaka Hirokawa
Hiroshi Kurakami
NTT Information Sharing Laboratories

Outline

- **Introduction**
  - Background
  - Motivation
  - Challenge
- **System Architecture**
  - BGP Collection
  - Flow Collection
  - Visualization
- **Traffic Change Detection Method**
- **Conclusion**

Background

- **Announcement of unwanted or invalid BGP route suddenly leads to traffic diversions.**
  - Cutting of submarine cable, route hijacking, misconfiguration, …
- **Moreover, it disrupts traffic or causes congestion on other backbone links.**

Motivation

- **Our goal is to reduce the load for troubleshooting.**
- **Our tool detects a traffic change and then identifies BGP route announcements involved.**
  - Monitors traffic volume for BGP attributes that have an impact on the traffic change:
    - ✓ Origin ASN
    - ✓ Neighbor ASN (peer ASN)
    - ✓ AS Path
    - ✓ BGP Next Hop
    - ✓ Community.
  - Identifies route changes that have an impact on the traffic change.

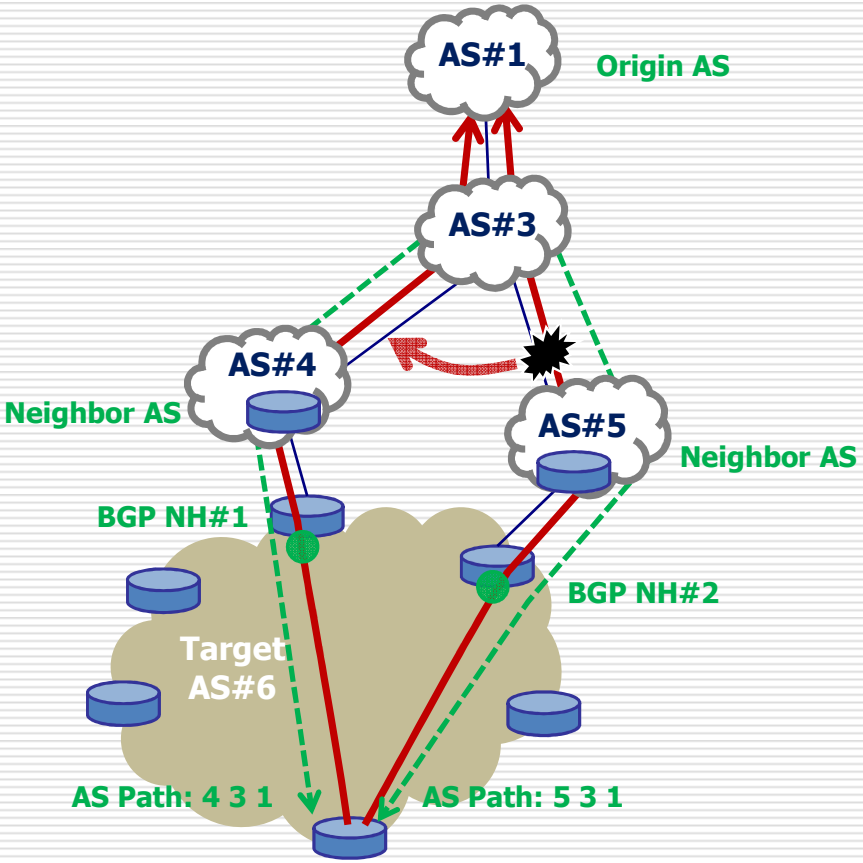

Related Work

- **Flow records from border routers can be utilized for origin or neighbor ASN traffic analysis.**
  - However, border gateway router cannot export both origin and neighbor ASNs.
  - Difficult to collect BGP Next Hop and AS Path info.
- **Some commercial collectors with BGP sessions can sum up traffic on the basis of BGP attributes.**
  - There are few tools for analyzing the interrelation of BGP and Flow data.
- **BGP and Flow analysis system have been proposed by several groups[1)].**
  - Simpler method and its visualization are required.

1) For example, J. Wu, Z. M. Mao, J. Rexford, and J. Wang, "Finding a needle in a haystack: Pinpointing significant BGP routing changes in an IP network," in Proc. NSDI, May 2005.

Challenge

- **The challenge is to identify route changes from a huge number of BGP route announcements.**
  - Hundreds of thousands of route announcements per day
- **Handle the huge load of flow records.**
  - Thousands of flow records per second
- **Explore a simple detection method and its real-time visualization.**

Data Source

- **Captures BGP data from BGP sessions to border routers or route reflectors as a BGP route reflector client.**
  - Border router feeds best routes to SASUKE tool.
- **Sets NetFlow/sFlow observation points at the periphery of the target AS.**

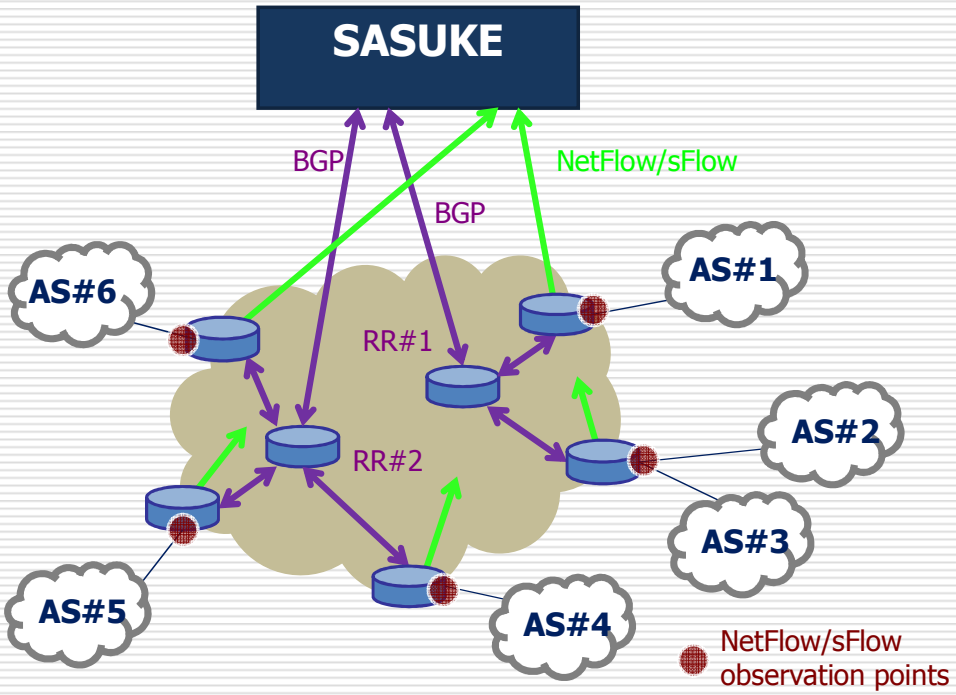

System Architecture

- 3 system components:
  - BGP Collection
  - Flow Collection
  - Data Analysis: correlation between BGP and Flow data

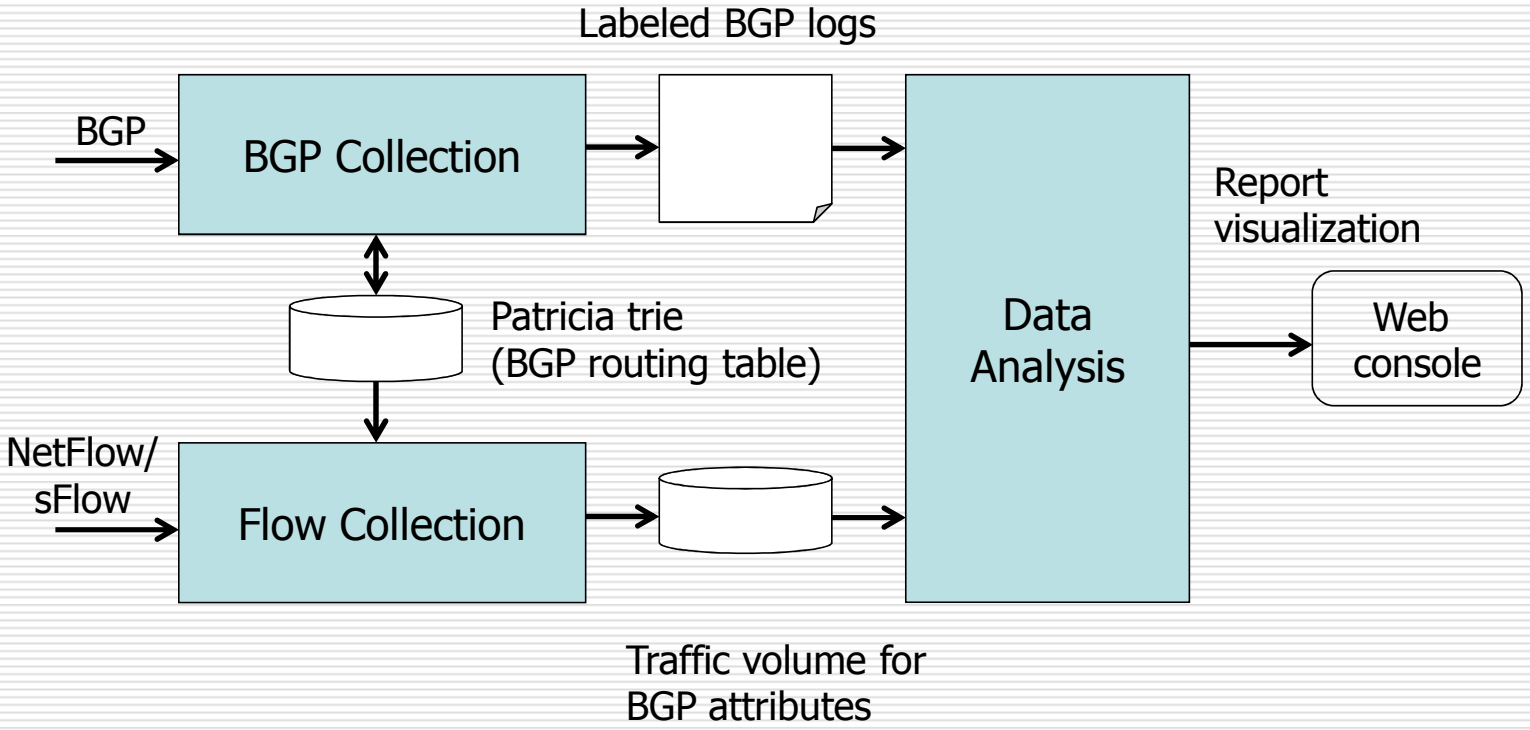

BGP Collection

- **Builds BGP routing tables as Patricia trie.**
  - Maintains tables for each BGP peering session.
- **Creates a BGP log report B to identify BGP messages that may cause a traffic change by comparing against the Patricia trie.**
  - Identifies BGP message type and BGP attributes that have changed from the old ones in the Patricia trie.

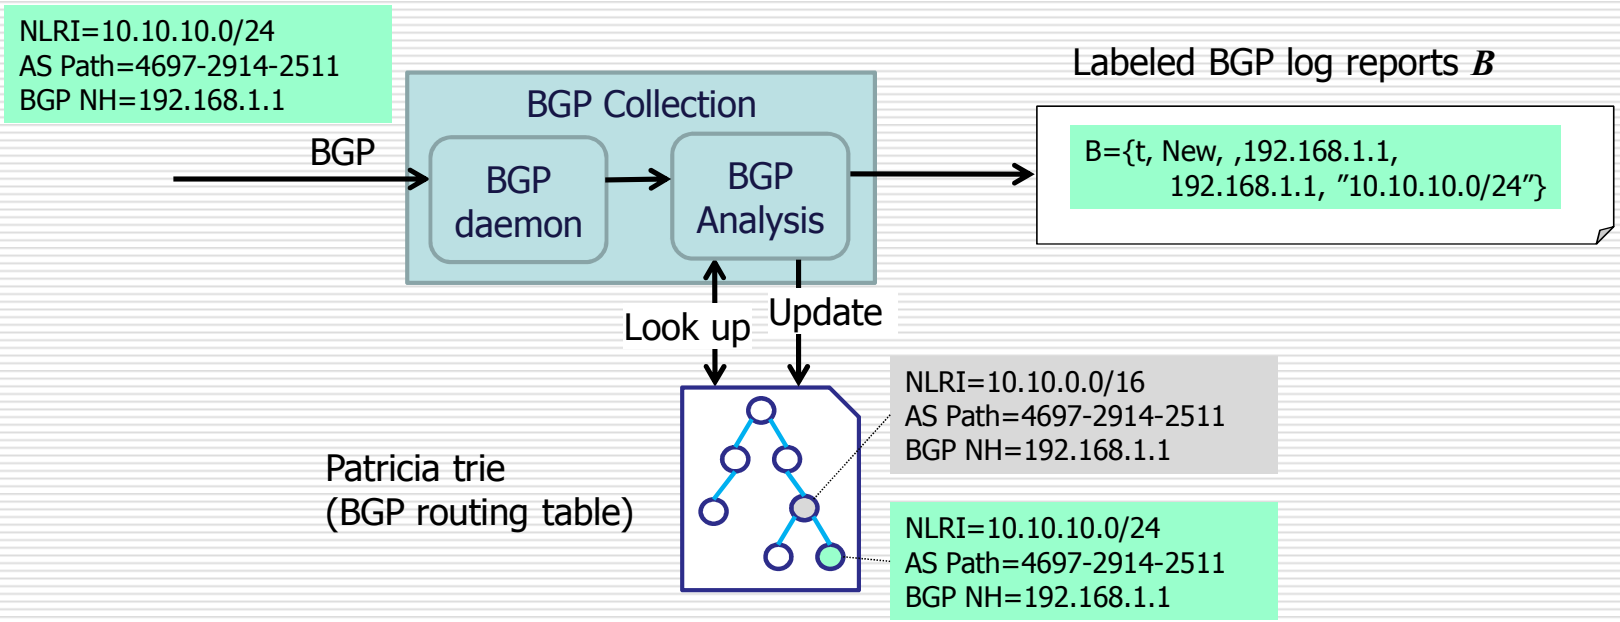

BGP Log Reports

- **BGP log report is represented as follows.**

$$B=\{t, c_{type}, a_{type}, a_{new}, a_{old}, prefix, id\}$$

  - t is timestamp of when the BGP message arrived.
  - c_{type} is change type:
    - "*New*", "*Withdraw*", "*Change*", or "*Duplicate*".
  - a_{type} indicates the changed BGP attribute type:
    - *"Origin ASN", "Neighbor ASN", "AS Path", "BGP NH"*, or *"Community"*
      - BGP community often gives route categories: region, peering type.
  - a_{new}/a_{old} are new/old BGP attribute values.
    - When c_{type} is *"New"* or *"Withdraw"*, a_{old} or a_{new} is a "null" value.
  - ***Prefix*** is network address in NLRI.
  - ***Id*** is an identifier to correlate with traffic data.
    - At this stage, the value is "null".

BGP Log Reports

- **Creates multiple BGP log reports when multiple BGP attributes are changed.**

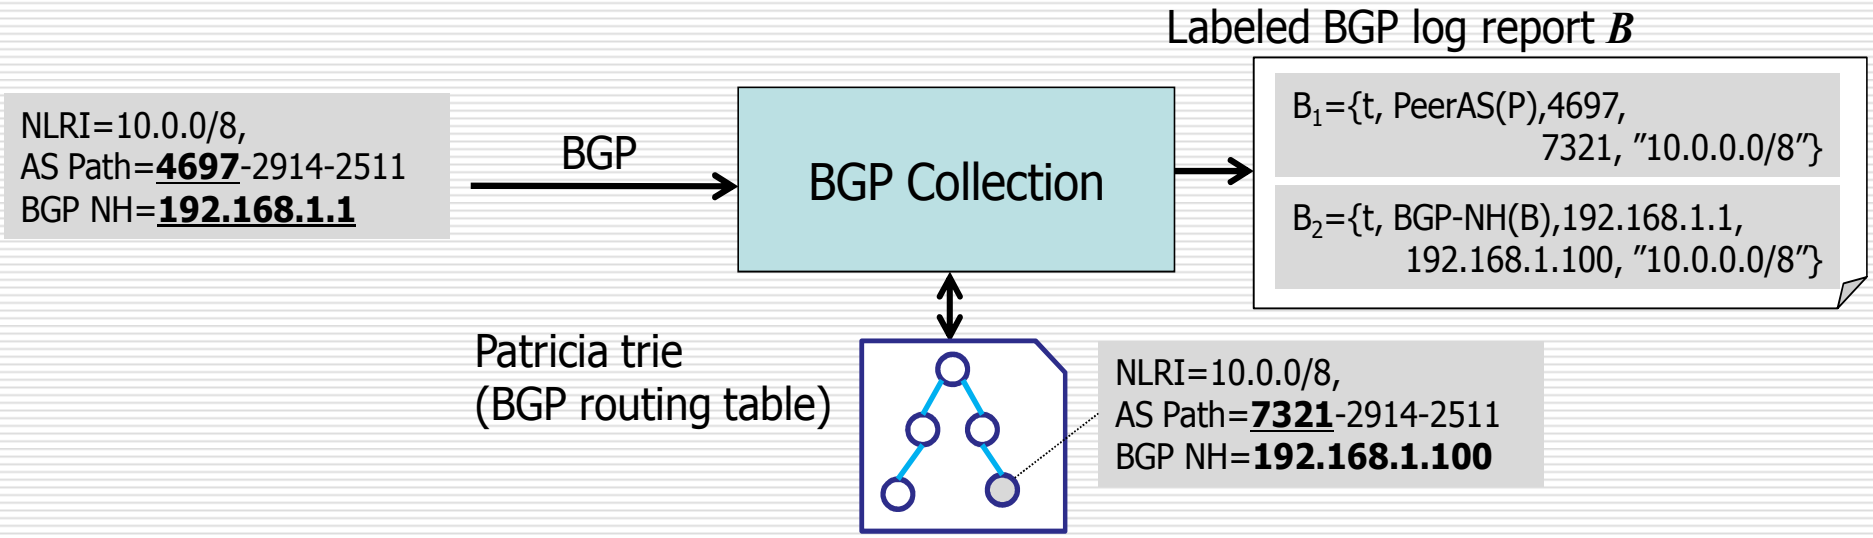

Visualization for BGP Log Reports

- **Labeled BGP logs are presented in time-series.**
  - Top-N origin/neighbor ASNs involved in the most BGP messages are represented when the spike happens.

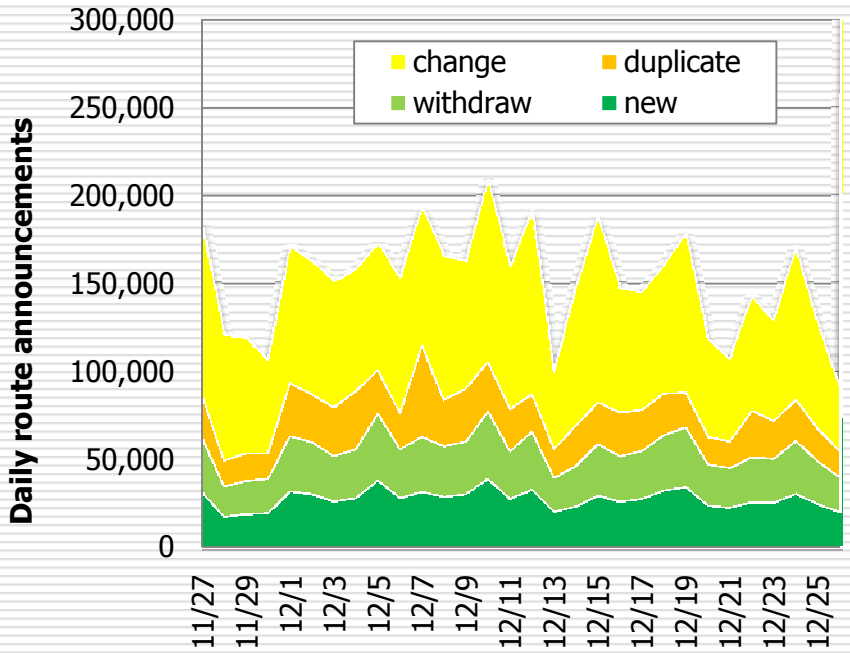

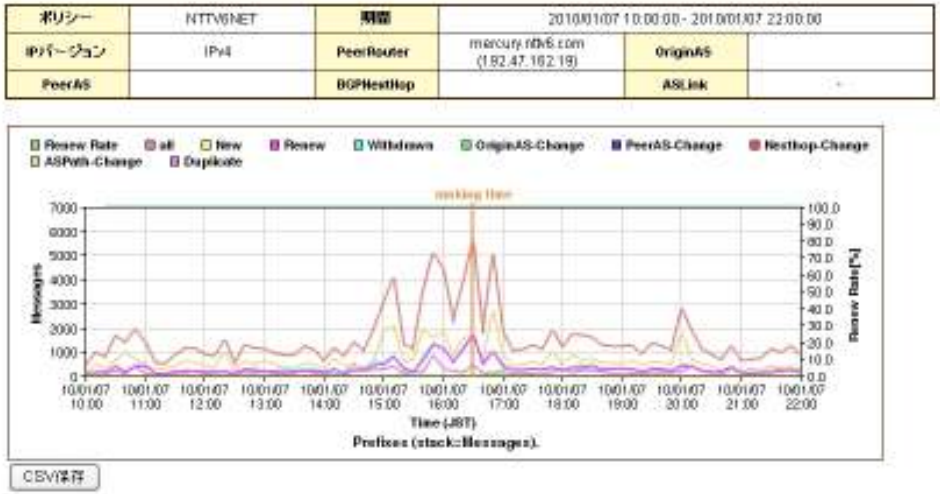

Top20 All Ranking

No.	OriginAS	Count	OriginAS		
			Country	Region	Organization
1	7545	3532		apnic	TPG-INTERNET-AP TPG Internet Pty Ltd
2	14522	425		lacnic	Satnet
3	5800	371		arin	DNIC-ASBLK-05800-06055 - DoD Network Information Center
4	20858	200		afrinic	EGYNET-AS
5	8474	102		ripencc	AS8474 Riscom AS
6	5536	87		afrinic	Internet-Egypt
7	47869	60		ripencc	SMARTLINK-AS Smart Links of Telecommunication Services
8	12524	35		ripencc	INDEX Autonomous System
9	6822	35		ripencc	SUPERONLINE-AS SuperOnline autonomous system
10	3912	34		arin	NMSU-AS - Cherry-net

Flow Collection

- **Selects an appropriate Patricia trie**.
  - Compares the Flow Record and peering data.
    - Are Exporter and BGP router the same device or not?
    - Are Exporter and BGP router located in the same region or not?
- **Sums up traffic on the basis of BGP attributes:**
  - Origin ASN, Neighbor ASN, AS Path, BGP NH, Community, and Prefix
  - These BGP attributes are retrieved from the Patricia trie by a longest match based on source/destination IP addresses.

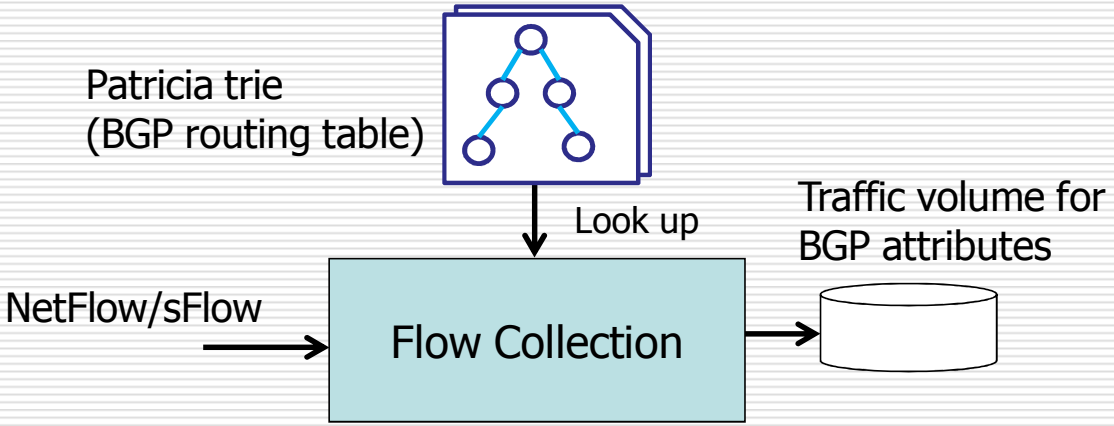

Visualization for Traffic Data

- Drill down into the detailed traffic data step by step from stacked area chart.

Origin ASNs stacked

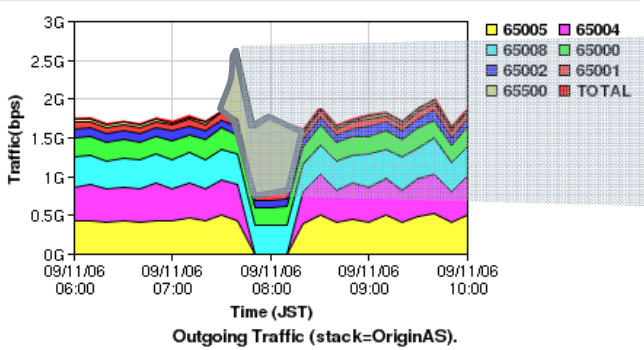

Prefix stacked

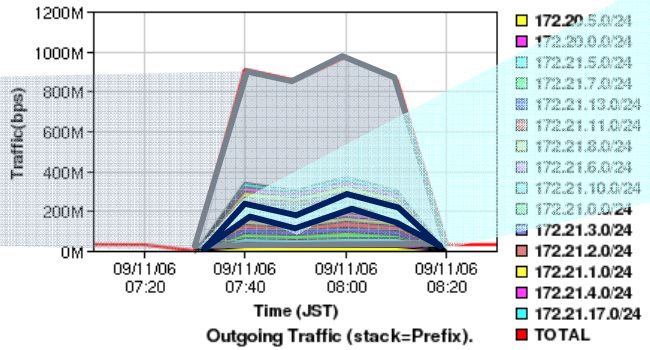

Detailed Flow record list

Neighbor ASNs stacked

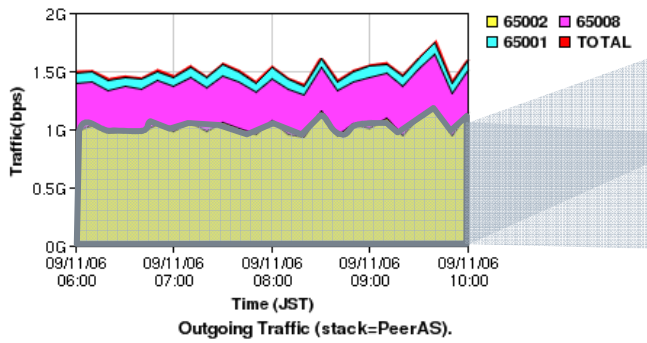

Origin AS stacked

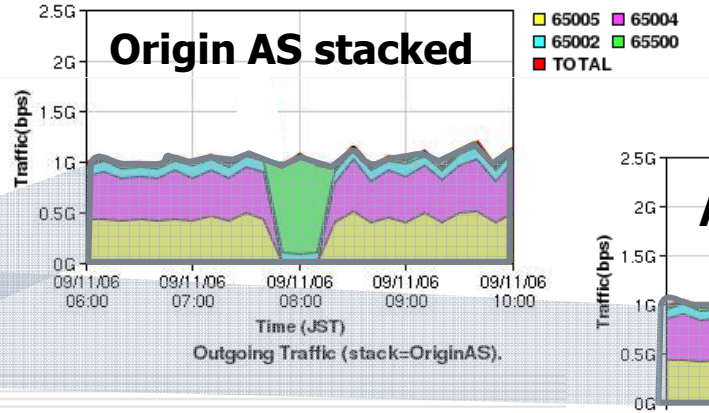

AS Path stacked

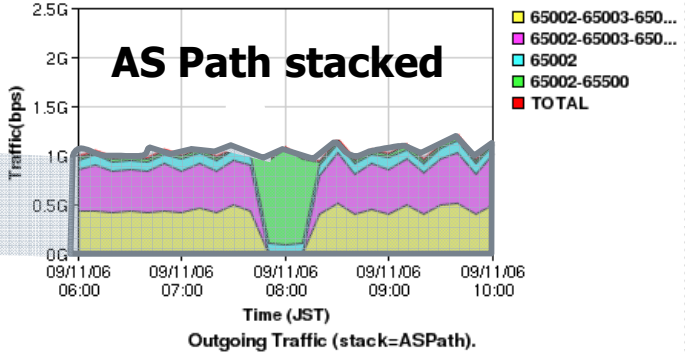

Traffic Change Detection Method

- **Focuses on a Top-N ranked by traffic volume on BGP attributes:**
  - Origin ASN, Neighbor ASN, AS Path, BGP NH, and Community.
- **Evaluates similarities of Top-N ranks between time slot t and t-1.**
  - Traffic volume weights the evaluation results.

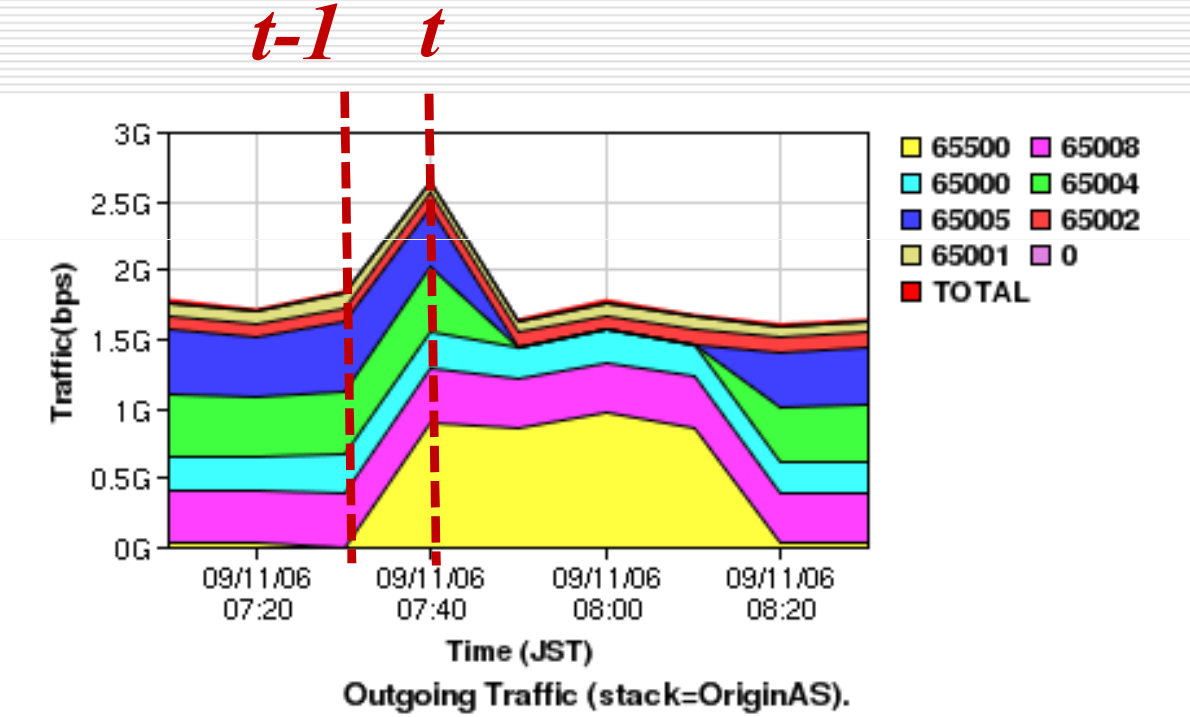

Traffic Change Detection Method

- **Calculates the correlation coefficient *r(t,t-1)* between the ranks of time slots *t* and *t-1*.**

$$r(t,t-1)=\frac{\sum_{i=1}^{N}\left(c(f(i,t),t)-c_{avg}(t)\right)\left(c(f(i,t),t-1)-c_{avg}(t-1)\right)}{\sqrt{\sum_{i=1}^{N}\left(c(f(i,t),t)-c_{avg}(t)\right)^2}\sqrt{\sum_{i=1}^{N}\left(c(f(i,t),t-1)-c_{avg}(t-1)\right)^2}}$$

  - ***f(i,t)*** is defined as a BGP attribute value ranked ***i*** by traffic volume of time slot ***t*.**
  - ***c(f(i,t),t)*** is traffic volume of ***f(i.t)*.**
  - Top-***N*** statistics data set $C(t)_i$ is presented as an array:
    - $C(t)_i=\{c(f(1,t),t), c(f(2,t),t), \ldots, c(f(N,t),t)\}$ $(i=1, 2, 3, \ldots, N)$
  - $C_{avg}(t)$ gives the average of $C(t)_i$.
- **Evaluates whether *r(t,t-1)* exceeds a threshold.**

Identify Most Affected Traffic

- **Investigates which $f(i,k)$ has the greatest impact on traffic change, as follows.**
  - Top-N rank $C(k)(t)_i$ indicates $C(t)_i$ except for $c(f(k,t),t)$ ranked k.

    $C(k)(t)_i=\{c(f(1,t),t), \ldots, c(f(k-1,t),t), c(f(k+1,t),t), \ldots, c(f(N,t),t)\}$
  - Calculates the correlation coefficient $r(k)(t,t-1)$.
  - Selects the greatest values $r(k)(t,t-1)$ from

    $\{r(1)(t,t-1), r(2)(t,t-1), \ldots, r(N)(t,t-1)\}$.
- **Then, we recognize that $f(k,t)$ is the most affected BGP attribute.**

Traffic Change Reports

- **Finally, it creates a traffic change report *R*.**

$$R=\{t, type, f(i,k), r(t,t-1), \delta, id\}$$

  - ***type*** gives traffic volume type:
    - ***"Origin ASN", "Neighbor ASN", "AS Path", "BGP NH"***, or ***"Community"***
  - δ gives the traffic volume difference between ***t*** and ***t-1***.
  - ***id*** is an identifier to correlate with BGP log reports.
    - At this stage, the value is set to a unique value.

Correlation between BGP and Flow

- **Correlates between BGP log report B and traffic change report R.**
  - Looks for the BGP log report B involved with traffic change report R.
  - Then, R and B are given the same id value to link them.

$R=\{t,\ type,\ f(i,k),\ r(t,t\text{-}1)\ ,\ \delta,\ id\}$
$B=\{t,\ c_{type},\ a_{type},\ a_{new},\ a_{old},\ prefix,\ id\}$

for all BGP log reports B where $t\text{-}Tw < B.t < t+Tw$ do

if $R.\delta > 0$ and $B.a_{type} = R.type$ and $B.a_{new} = R.f(i,k)$ *then*
$B.id = R.id$;
else if $R.\delta < 0$ and $B.a_{type} = R.type$ and $B.a_{old} = R.f(i,k)$ *then*
$B.id = R.id$;
end if

end for

Visualization for BGP and Flow

- **Creates traffic-change-related alert for operators.**
  - Alert links the graphs of traffic volume area chart and of BGP log reports involved.

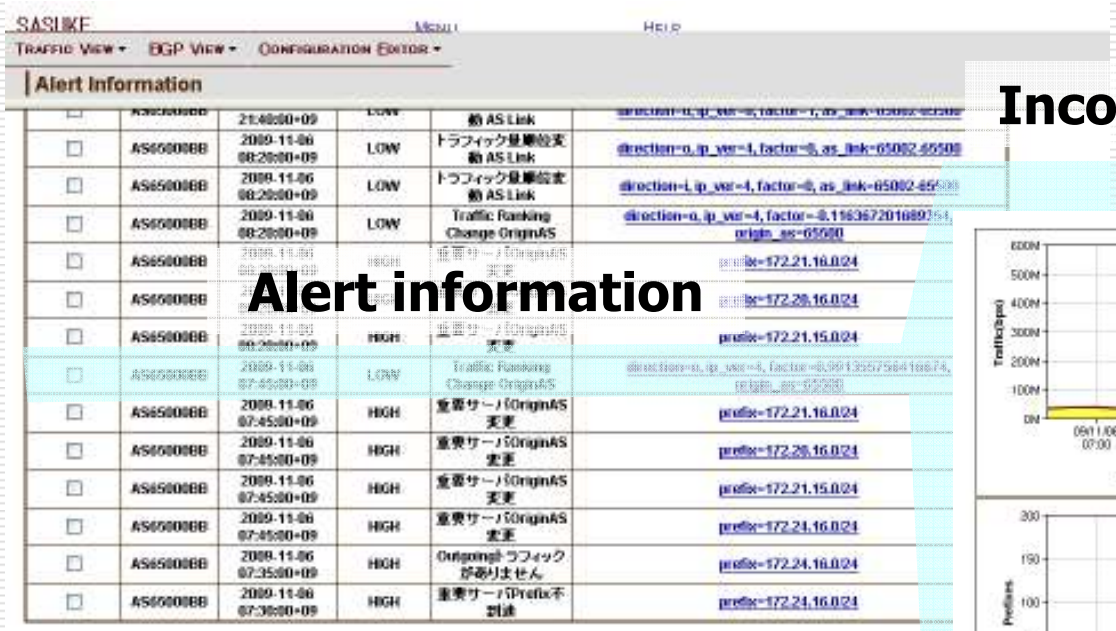

Incoming traffic

Outgoing traffic

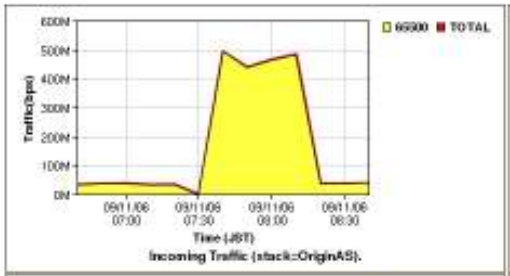

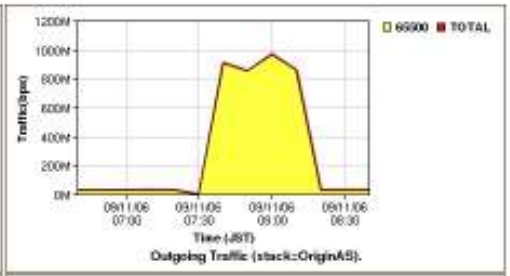

Number of prefixes

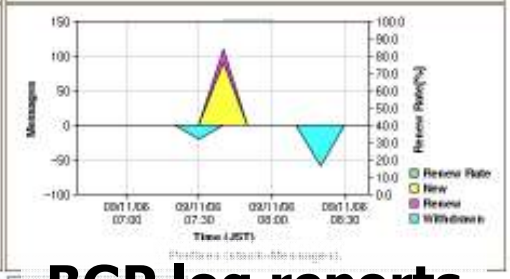

BGP log reports involved

Overall Traffic View

Incoming Traffic

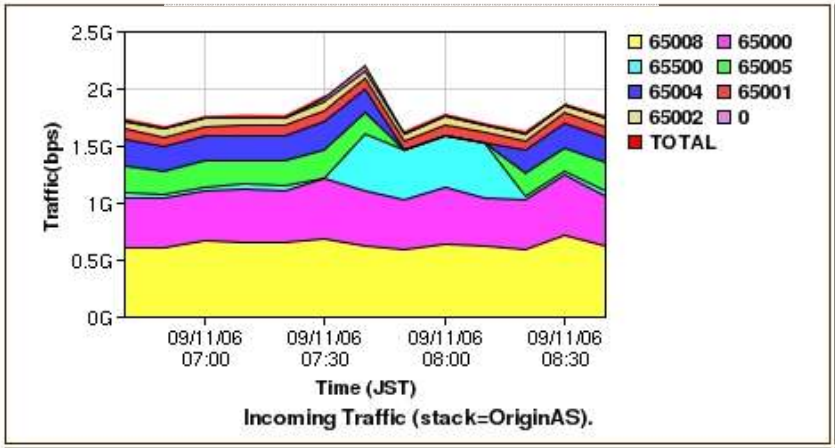

Outgoing Traffic

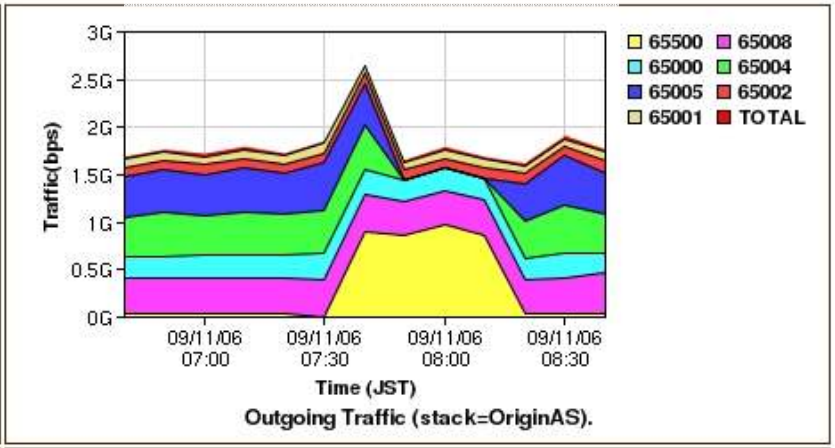

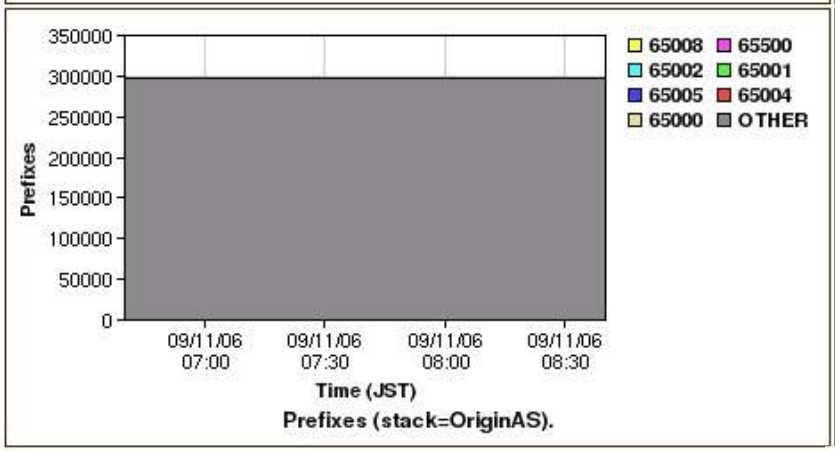

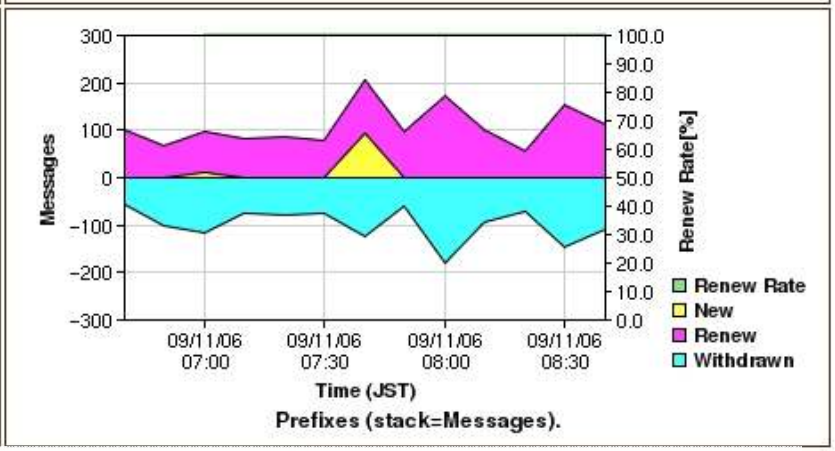

Incomming Traffic CSV出力

Outgoing Traffic CSV出力

Number of Prefix

BGP log reports

Evaluation of SASUKE Tool

- **"SASUKE" has been introduced in some commercial networks as an experimental phase.**
  - Much more performance evaluation is needed.

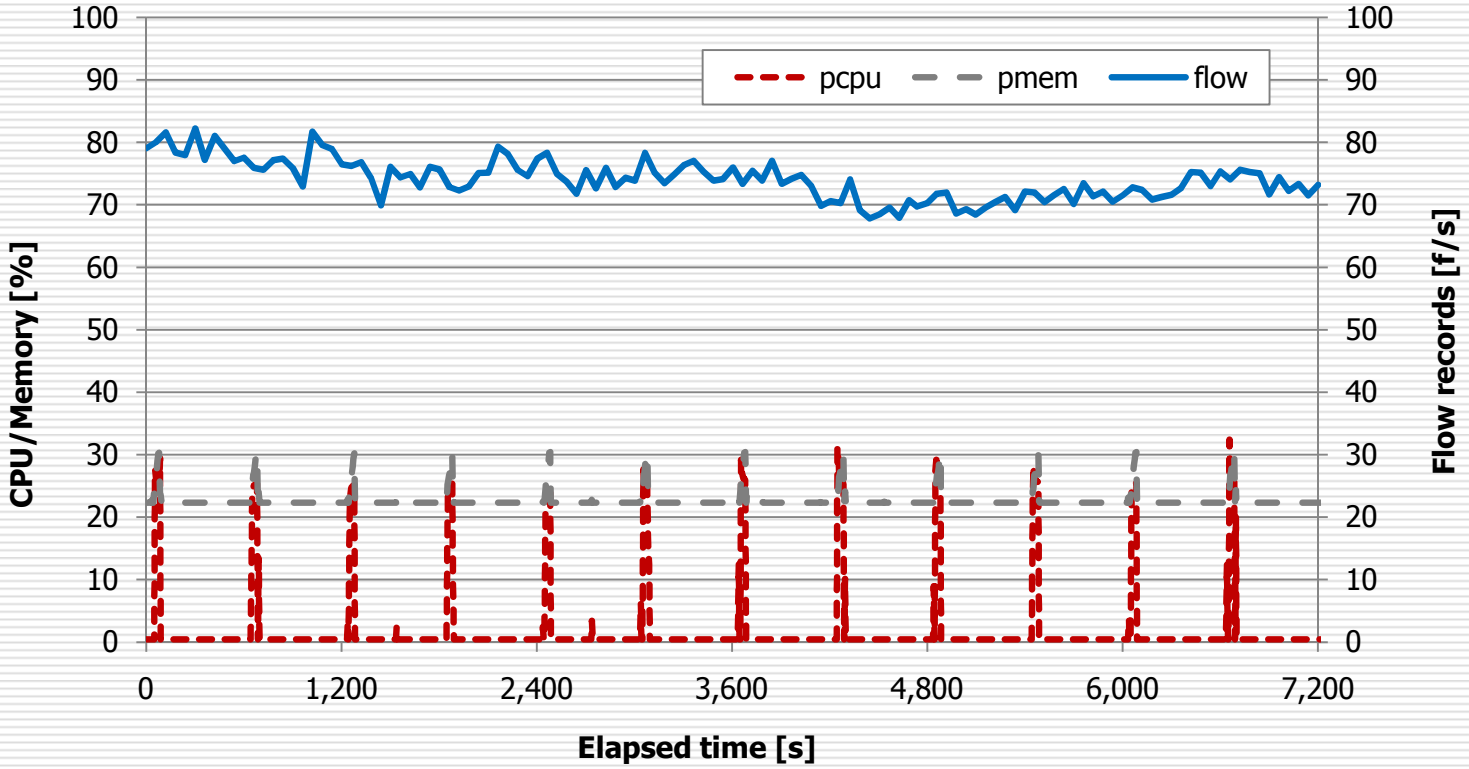

Conclusion

- **I demonstrated the traffic change detection method implemented in the "SASUKE" tool.**
  - Focuses on the similarities between time consecutive Top-N ranks in time-series.
  - Correlates between BGP log reports B and traffic change reports R.
  - Alleviates the troubleshooting load for network operators.
    - Visualizes BGP log and traffic data.
    - Links multidimensional traffic data related to BGP attributes.
- **More evaluation is needed for performance and accuracy.**

Thank you very much.

This study was supported by the Ministry of Internal Affairs and Communications of Japan.