A Flexible Policy Control Architecture for Inter-AS Routing

Osamu Akashi, Kenichi Kourai, Kensuke Fukuda,
Toshio Hirotsu, Koji Sato, Mitsuru Maruyama,
Toshiharu Sugawara
NTT Network Innovation Laboratories
Tokyo, Japan
{akashi,kourai,fukuda,hirotsu,koji,
mitsuru,sugawara}@t.onlab.ntt.co.jp

Abstract

Inter-AS policy control is important to reliably and efficiently operate the Internet, but Inter-AS routing is difficult to control since advertised routing information is modified as it spreads through ASes. Since these ASes are managed by independent administrators acting based on each own policy, cooperative distributed solutions are desirable. To cope with this problem, we propose a policy control architecture AISLE that uses distributed agents that act autonomously and adjust routing behavior based on given policy description considering environmental changes. Their cooperative actions enable each AS to control and coordinate routing behavior at the inter-AS level.

Keyword: BGP, multi-agent, policy control

Background in inter-AS routing

- Routing information is spreading over the Internet in a hop-by-hop manner using BGP-4.
  - Receive → modify → advertise
  - Each BGP entry includes AS path information that information traversed.

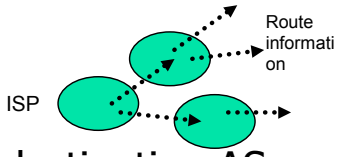

- All ASes along with source AS to destination AS should set their routing tables as source AS intends.
- Loss of connectivity, instable access, policy inconsistency

Border Gateway Protocol (BGP-4) [1] is widely used for inter-AS (Autonomous System) routing and has become the current de facto standard. The AS that receives BGP information selects only one route per destination as a best path based on the AS's policy, and passes the information to other ASes. So the routing information spreads among these ASes in a hop-by-hop manner.

Inter-AS routing controlled by BGP is not stable [2,3] and worldwide or local routing accidents have been reported [4]. Since there are more than 11000 independent ASes where each AS is controlled by a single administrative authority based on that AS's own policy, network reachability failures or unintended traffic flow are easily caused by various events such as hardware failures, routing protocol failures, routing configuration errors, and routing inconsistency among several ASes. These failures could easily cause instability or loss of access on an extensive scale. Even if reachability is retained, inconsistency with routing policy occurs more frequently.

Problems of inter-AS routing

- Difficulty in understanding the behavior
  - Routing information mutates as it spreads.
  - Independent administrative domain that has its own policy and routers are configured by hand.
  - Needs analysis by experts by hand
    - ex. Using tool such as **`Looking glass`**
- Operators cannot adapt dynamically changing environment.
  - Policy is only represented by low level primitives, namely router configuration commands.
- No inter-AS cooperative policy control scheme

Need a cooperative distributed solution

3

One of essential problems is that it is difficult for us to know or control the behavior of the spreading routing information that we are advertising. Moreover, we cannot control spreading information to fit to our management policy, because advertised information is modified and passed based on other ASes' policy independently on which routing information traverse. These means control by a single centralized system impossible.

From the architectural point, policy is represented in highly abstracted level in human and should be translated to configuration commands of a router. This configuration is statically given and cannot change its behavior reacting with environments change. So no feedback mechanism is given. The configuration only describes about one router's behavior, while the inter-AS policy control need coordination among several ASes.

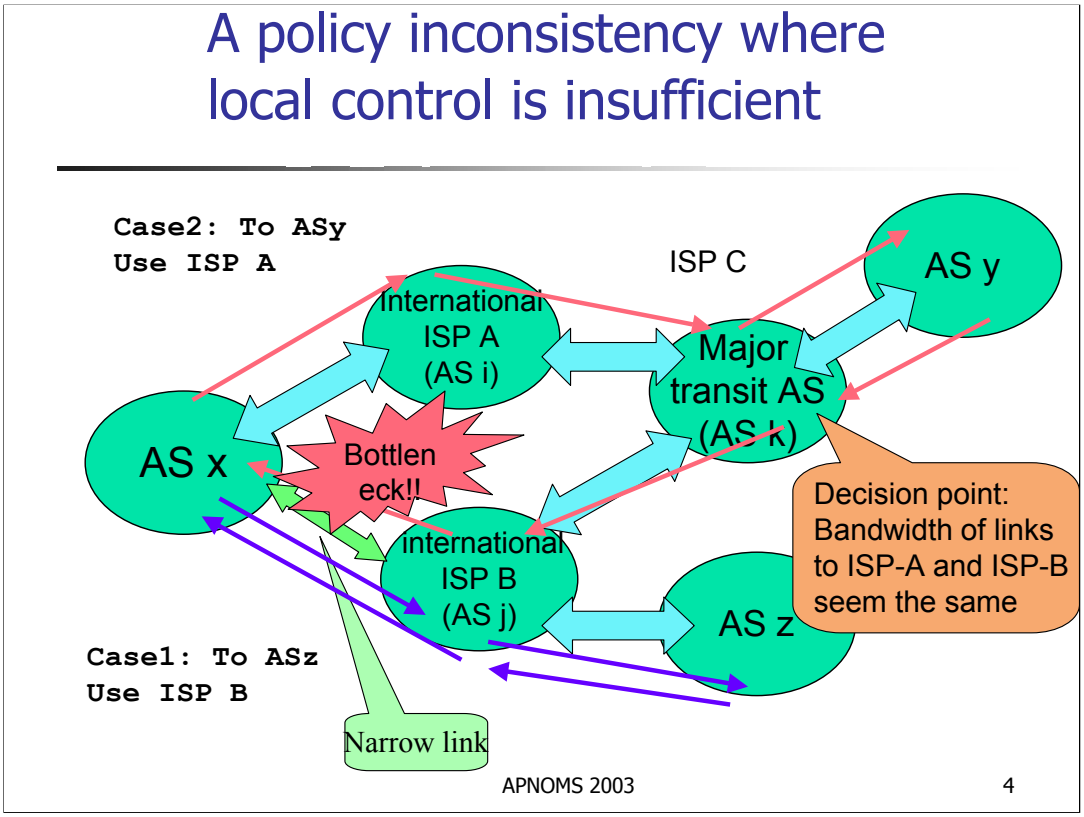

This shows a policy inconsistency example. In this example, AS x has two BGP peers, namely AS i and AS j. Advertised BGP information from AS x reaches AS k via AS i and AS j. Then AS k must select only one route as the best path. If two links between AS k and AS i and between AS k and AS j have the same bandwidth, AS k might select route via AS j as the best path for AS x. Since the bandwidth of the link between AS x and AS j is narrow, it becomes bottleneck even if route via AS i is available. Unintended route for incoming packets is used.

Our approach

- Diagnosis for inter-AS routing anomalies
  - ENCORE[4, 5]
    - Multi-agent based, cooperative analysis
- Flexible inter-AS policy control
  - AISLE

    (Autonomous and Intelligent Self-control Environment)

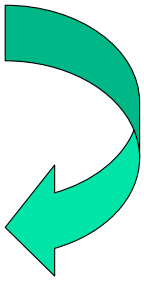

Extended

APNOMS 2003 5

These problems require that autonomous policy control process should share inter-AS routing information from outside the AS and a cooperative distributed solution should be established. Moreover, intermediate control layer between policy and router configuration is desirable.

In this paper, we propose an inter-AS policy control architecture AISLE (Autonomous and Intelligent Self-control Environment) that provides autonomous and flexible routing policy control mechanism at the inter-AS level. From the view point of diagnosis of inter-AS routing anomalies, we have proposed a multi-agent-based inter-AS diagnostic system called ENCORE [4,5]. In this system, a collection of intelligent agents located in multiple ASes perform cooperative observation and analysis. It analyzes the causes of the anomalies based on its embedded diagnostic knowledge and integrated information observed from multiple viewpoints. We use this framework to observe policy at the inter-AS level and extends these cooperative actions to achieve autonomous and flexible policy control.

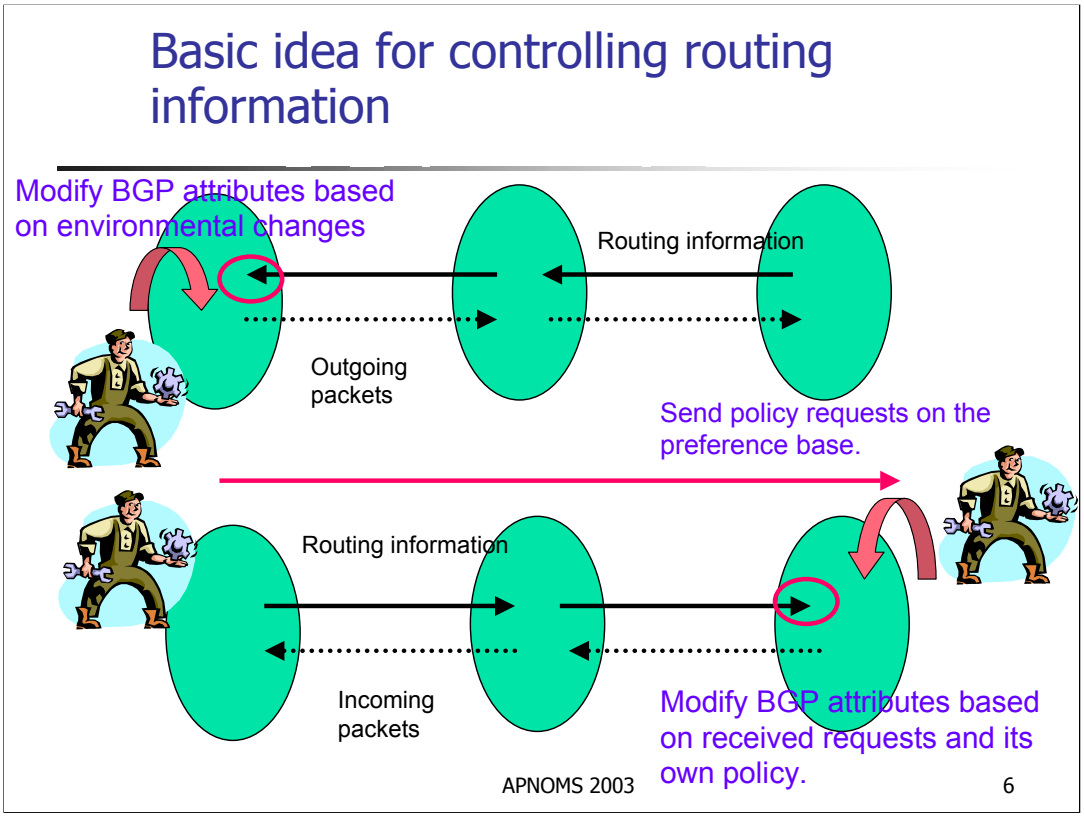

As shown in this figure, we are difficult to know the behavior of spreading inter-AS routing information or to control it to fit policy as the originator intends at the inter-AS level. BGP defines several attribute parameters for controlling policy, but policy cannot utilize these parameters effectively if the policy is controlled on a single router in an AS using only static low level configuration primitives.

In case of outgoing packets, AISLE dynamically modifies attribute values in received BGP information in the local AS to fit given policy description considering network status changes. This is feedback action based on observation results in its deployed environment. In case of incoming packets, AISLE controls attribute values in advertised BGP information in the remote AS using cooperation actions among ASes.

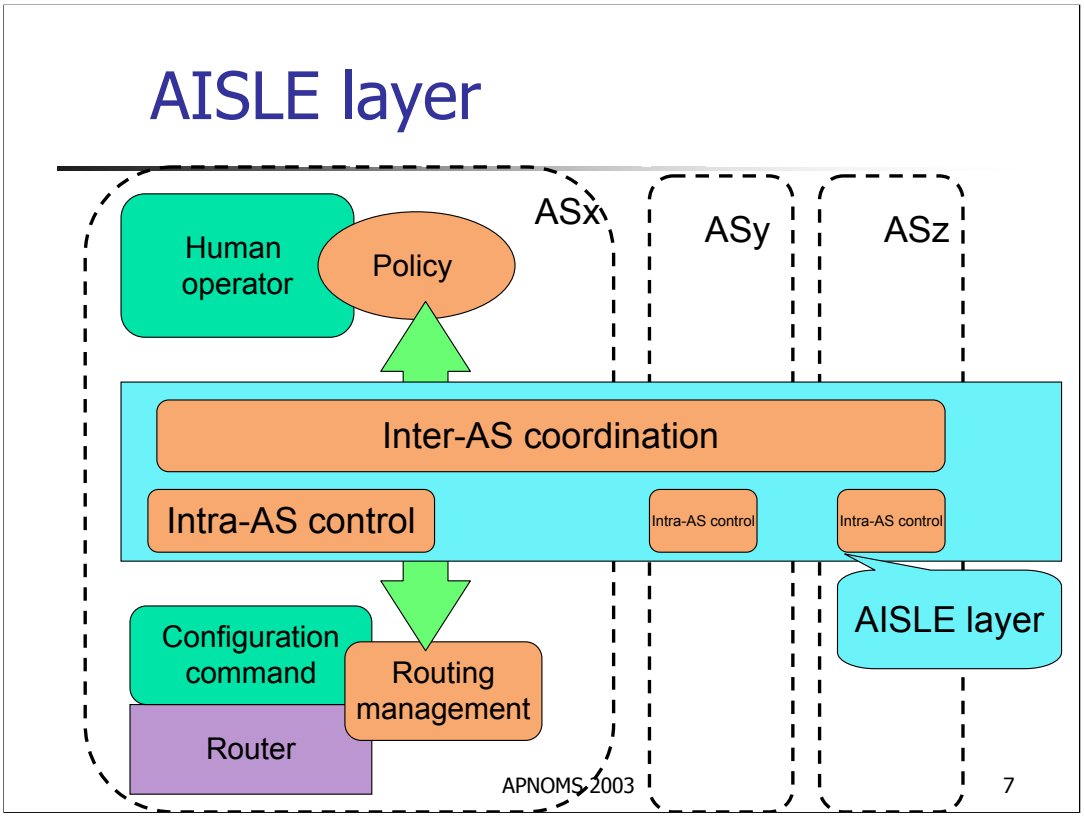

The policy control layer of AISLE is realized as it exists between policy that resides in human operators and router configuration commands. The AISLE layer consists of intra-AS control part and inter-AS control part since each AS is independently managed and a simple centralized model would be difficult to apply. We can assume that an agent can access routing information in its AS because the AS is controlled by the same administrative authority.

The intra-AS part controls policy in its AS from the local perspective. The inter-AS coordination part controls policy at the inter-AS level through the cooperative actions among multiple intra-AS parts that reside in different ASes. The cooperation model means it retains autonomy where an entity could determine actual actions based on requests from other agents. So the entity might refuse requests that are inconsistent with the entity's intention.

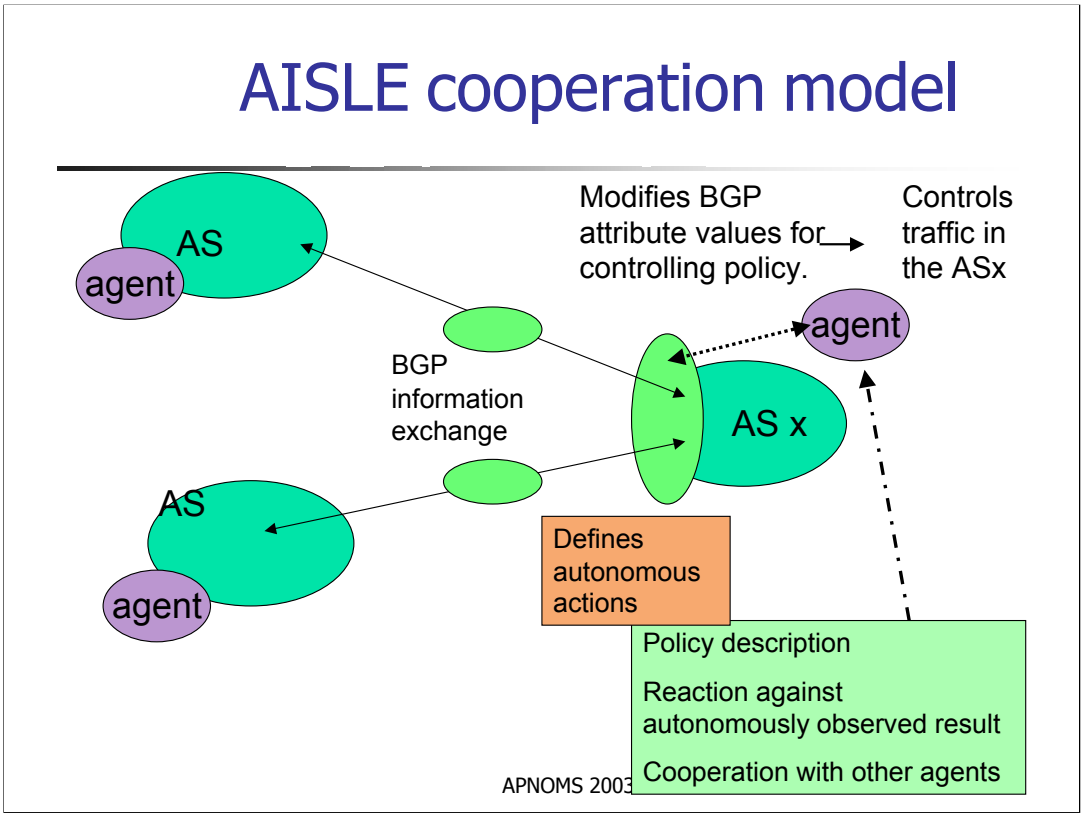

We adopt a multi-agent approach as the ENCORE model [4] where autonomous agents are distributed among multiple ASes. Intra-AS control is achieved by interaction between an agent and a border router in its AS. The agent that acts based on given policy description for the AS monitors BGP information and modifies it back to the router to reflect the policy or to adapt environmental changes. The agent can cooperate with other agents in other ASes. Through the cooperative actions, the policy at the inter-AS level is coordinated.

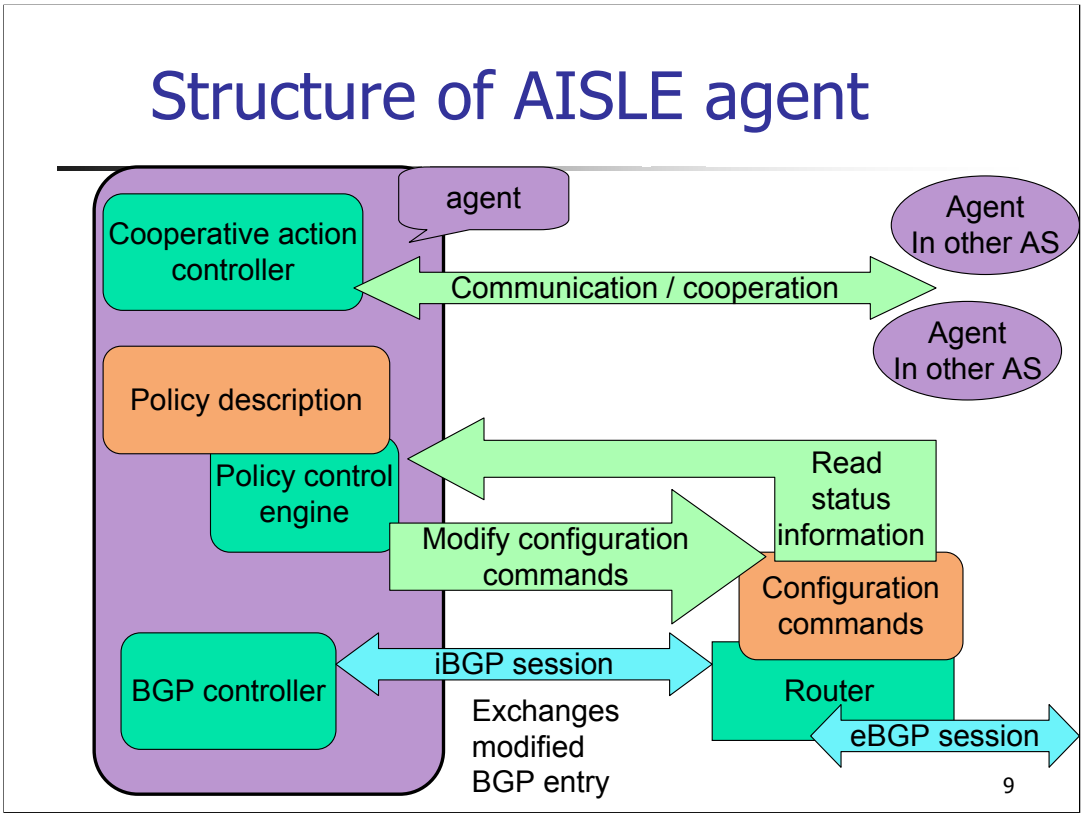

AISLE policy description provides more abstracted configuration functions than router commands. These functions are required to represent policy more intuitively and to define more complicated functions. For example, traffic control on several links to other ASes is determined after actual traffic monitoring. Because these values are not constant, we cannot give suitable configuration before observation and it cannot adapt these changes dynamically. So these parameters are described as variables in policy description and these actual values are set and modified through observation.

For controlling BGP information, the agent modifies several BGP attribute information to reflect policy and to adjust environmental changes. The agent monitors BGP entry information using internal BGP session as an internal BGP peer and inserts modified BGP entries for changing next hop AS of outbound packets. In addition to BGP information insertion using internal BGP session, the agent can send configuration commands directly using telnet session. It is intended for control at the IP level such as filtering.

Application

- Adaptation of policy for dynamic parameters
  - Auto load balancing
  - On demand advertisement of backup route
- Cooperative control
  - Preference control for incoming packets
  - Verification of routing policy
  - Defense against attacks

AISLE framework can be applied to many management situations as shown.

This policy example would be used when an AS wanted to distribute routes for outgoing packets among several peering ASes.

```
;; Distribute outbound traffic evenly
;;; in number of destination routes.
(def-strategy check-out-traffic
  (:interval (* 60 10)) ;;; 10 [min]
  (:inhibit-interval (* 60 20))   ;;; Inhibit interval against any changes
  (rule distribute-out-bound) )

(def-rule distribute-out-bound
  (acq get-balance-info)
  (eval balance-next-hop-in-number  (acq-result) ))
```

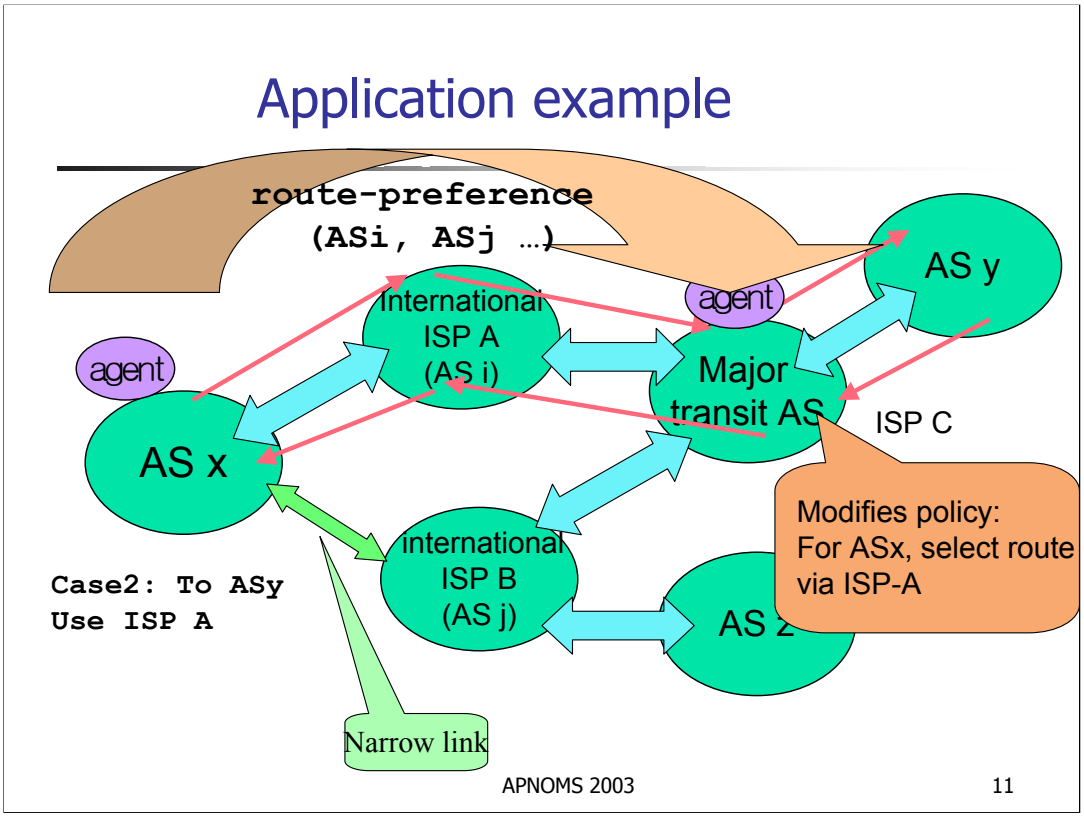

This shows an application example against previously explained policy inconsistency. In this example, unintended route for incoming packets is adjusted by sending a route preference request to a transit AS. The agent that receives preference checks its own policy. If there is no conflict, the agent could change the route destined for AS x via AS j to the route via AS j. This policy example would be used when AS x wanted to notify some major ASes to use suitable AS, namely ISP A.

```
;; Distribute inbound traffic
(def-sp-var up-stream-AS-list  (get-initial-list) )
(def-strategy check-up-agents ...)
;;;
(def-sp-var incoming-pref  '(ASi ASj))
(def-strategy modify-in-pref-list  (:interval 3600) ... )
;;;
(def-strategy check-in-traffic
  (:interval (* 3600 6)) ;;; 6[hour]
  (rule change-in-bound
    :every target-AS up-stream-AS-list ))

(def-rule change-in-bounnd
  (acq set-next-AS-preference-if-possible (incoming-pref)
        :cooperative target-AS)
  (eval report-result (acq-result)) )
```

Conclusion

- AISLE: inter-AS flexible policy control architecture
  - Multi-agent based implementation
- Needs more experiment in real internet environments
  - Verification and feedback

Intelligent routers can control outbound traffic by modifying received BGP information based on given policy description, but they do not provide cooperative actions among multiple ASes. Although some community attribute extension for policy control is proposed [7], it only defines the mechanism how to distribute additional values on BGP and does not discuss inter-AS routing adjustment or coordination functions.

For autonomous and flexible inter-AS policy control, we have proposed AISLE architecture that uses multiple cooperative agents. In this environment, multiple agents can adjust and coordinate policy at the inter-AS level and can reflect the current network status to policy control.

Reference

[1] Y.Rekhter and T.Li. ``A Border Gateway Protocol 4 (BGP-4)", 1995. RFC1771.

[2] The North American Network Operators' Group. NANOG mailing list. http://www.nanog.org.

[3] C.Labovitz, G.R. Malan, and F.Jahanian. ``Internet Routing Instability". In Proc. of ACM SIGCOMM, 1997.

[4] O.Akashi, T.Sugawara, K.Murakami, M.Maruyama, and K.Koyanagi.

Agent System for Inter-AS Routing Error Diagnosis. IEEE Internet Computing, 6(3):78 -- 82, 2002.

[5] O.Akashi, T.Sugawara, K.Murakami, M.Maruyama, and N.Takahashi.

``Multiagent-based Cooperative Inter-AS Diagnosis in ENCORE". In IEEE/IFIP Network Operations and Management Symposium, pages 521 -- 534, Apr 2000.

[6] E.Kern. http://nitrous.digex.net.

[7] draft-ietf-idr-bgp-ext-communities-06.txt, S. Sangli, D. Tappan, Y. Rekhter