# TECHNICAL RESEARCH REPORT

Interesting Examples of IBGP Configuration

*by Anuj Rawat, Mark Shayman*

**TR 2006-6**

# Interesting Examples of IBGP Configuration

Anuj Rawat and Mark A. Shayman
Department of Electrical and Computer Engineering
University of Maryland, College Park, MD 20742
Email: {anuj, shayman}@glue.umd.edu

***Abstract*— In this paper we give examples to show that if an Internal Border Gateway Protocol (IBGP) configuration using route reflections violates even one of the four conditions mentioned in the theorem given in [1], then there may be persistent oscillations or forwarding loops.**

## I. Introduction

It has been observed that scaling IBGP using route reflection may lead to problems such as routing oscillations and forwarding loops. There have been attempts to suppress these anomalies by changing the IBGP protocol and by using graph theoretic analysis. Anuj et al. [1] model the autonomous system (AS) as an IGP connectivity graph $G_I$ and an IBGP peering graph $G_L$. They then state and prove conditions on these graphs so that the IBGP configuration using route reflection is free from persistent oscillations and forwarding loops due to MED attribute and IBGP path asymmetry. In this paper we give example to show that if an IBGP configuration using route reflection violates even one of the conditions mentioned in [1] then there may be persistent oscillations or loops in the system.

Section II presents the theorem given in [1]. In section III we present the examples to show that violating any condition mentioned in the theorem can lead to problem. Finally section IV concludes the paper.

## II. Theorem

The theorem mentioned in [1] is given below.

*Theorem 2.1:* If an AS configuration with route reflection satisfies each one of the following conditions then it is free of persistent route oscillations as well as forwarding loops.

- (*i*) If nodes $u, v$ learn about paths $P, Q$ respectively, having $nextAS(P) = nextAS(Q)$ through EBGP sessions, then $u, v$ are IBGP peers.
- (*ii*) Clients of same cluster are not IBGP peers.
- (*iii*) $cost(sp(u, v)) < cost(sp(u, w))$ for all nodes $u, v, w$ such that $u, v \in$ cluster $V_i$, $w \in$ cluster $V_j$ and $i \neq j$.
- (*iv*) If $u_i \in V_i$ and $u_j \in V_j$ are client nodes and $i \neq j$, then there is a reflector $u_k \in sp(u_i, u_j)$.

## III. Examples

In this section we present examples which show that if an IBGP configuration violates even one condition of Theorem 2.1 then we may have persistent oscillations or loops.

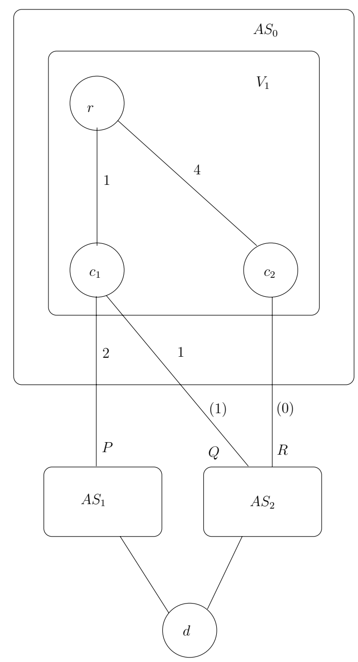


Fig. 1. Condition (*i*) violated: Persistent oscillations

### A. Example for Condition 1

In Fig. 1, $AS_0$ has one cluster $V_1$ with reflector $r$ and clients $c_1, c_2$. The only IBGP sessions are between node pairs $r, c_1$ and $r, c_2$. The IGP costs for the shortest path between the node pairs are indicated besides the lines joining them. There are three available paths to external destination $d$. Client $c_1$ learns about paths $P, Q$ through EBGP peers in $AS_1, AS_2$ respectively, and client $c_2$ learns about path $R$ through an EBGP peer in $AS_2$. The weights on EBGP links of $c_1$ represent its preference in BGP tie-breaking (lower value is more preferred). Also the MED values for paths $Q, R$ (through same neighboring AS) are given in parentheses. Now node $c_2$ will always select path $R$ and therefore $R$ is always visible at node $r$. Suppose path $R$ is invisible at node $c_1$, so out of paths $P, Q$ it selects $Q$ and advertises it to node $r$. Now, out of paths $Q, R$, node $r$ selects path $R$ (lower MED value) and advertises it to $c_1$. Now all the three paths $P, Q, R$ are visible at $c_1$ and it selects path $P$ and advertises it to $r$. Now $r$ sees paths $P, R$ and selects $P$ (based on lower IGP cost). Now path $R$ is again invisible at $c_1$ and it changes its path selection from $P$ to $Q$. So we see that the system is in a persistent oscillation and nodes $r, c_1$ keep on changing their path selections. In this example, the IBGP configuration violates condition (*i*) of Theorem 2.1,

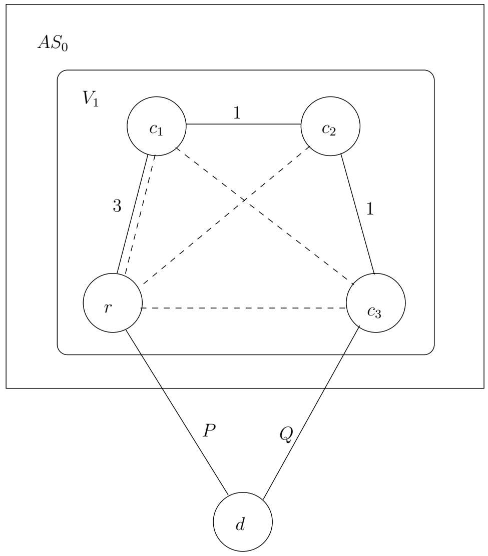


Fig. 2. Condition (*ii*) violated: Forwarding loop

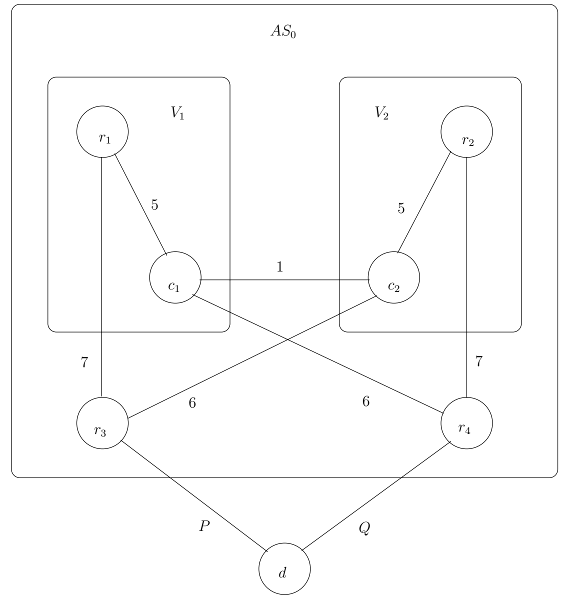


Fig. 3. Condition (*iii*) violated: Forwarding loop

since nodes $c_1, c_2$ learn about paths $Q, R$ through EBGP peers in the same neighboring AS, but are not IBGP peers.

### *B. Example for Condition 2*

In Fig. 2, $AS_0$ has one cluster $V_1$ with reflector $r$ and clients $c_1, c_2$ and $c_3$. The dotted lines represent the IBGP sessions and the solid lines represent the IGP links between the two nodes. The IGP costs are indicated besides the lines. There are two paths to external node $d$. Reflector $r$ learns about path $P$ through EBGP session and client $c_3$ learns about path $Q$ through EBGP session. Now in this example we have an extra IBGP session between clients $c_1$ and $c_3$. Now we can see that $c_1$ selects $Q$ as its best path (based on lower IGP to NEXT_HOP node) whereas $c_2$ selects path $P$ (only path known). Now it is easy to see that there is a forwarding loop between nodes $c_1$ and $c_2$. In this example, the IBGP configuration violates condition (*ii*) of Theorem 2.1, since there is an IBGP session between the client nodes $c_1, c_3$.

### *C. Example for Condition 3*

In Fig. 3, $AS_0$ has clusters $V_1$ and $V_2$, both having one reflector and one client ($r_1, c_1$ and $r_2, c_2$ respectively). There are two stand alone reflectors $r_3$ and $r_4$ (these may very well be considered to be in clusters $V_3$ and $V_4$ respectively, with both the clusters having only one node). We assume that there is a full mesh of IBGP sessions between all the reflectors and the only other IBGP sessions are between $r_1, c_1$ and $r_2, c_2$. The lines represent the IGP links between two nodes. The IGP costs are indicated besides the lines. There are two paths to external node $d$. Reflector $r_3$ learns about path $P$ through an EBGP session and reflector $r_4$ learns about path $Q$ through an EBGP session. Now $r_1$ selects $P$ as its best path (based on lower IGP to NEXT_HOP node) therefore $c_1$ also selects path $P$ (only path known). Similarly $r_2$ and $c_2$ select path $Q$. Now it is easy to observe that the shortest path from $c_1$ to $r_3$ goes through $c_2$ and the shortest path from $c_2$ to $r_4$ goes through $c_1$ so there is a forwarding loop between nodes $c_1$ and $c_2$. The problem with this configuration is that IGP cost of shortest path between $c_1$ and $c_2$ is less than the IGP cost of the shortest path between $c_1$ and $r_1$, which violates condition (*iii*) of Theorem 2.1.

### *D. Example for Condition 4*

In Fig. 4 the AS has clusters $V_1, V_2$ and $V_3$, all having one reflector and one client ($r_1, c_1$ ; $r_2, c_2$ ; $r_3, c_3$). There are three stand alone reflectors $r_4, r_5$ and $r_6$ (these may very well be considered to be in clusters $V_4, V_5$ and $V_6$ respectively, with all the clusters having only one node). We assume that there is a full mesh of IBGP sessions between the reflectors and the only other IBGP sessions are between $r_1, c_1$ ; $r_2, c_2$ ; $r_3, c_3$. The lines represent the IGP links between two nodes. The IGP costs are indicated besides the lines. There are three paths to external node $d$. Reflector $r_4$ learns about path $P_1$ through an EBGP session, reflector $r_5$ learns about path $P_2$ through an EBGP session and reflector $r_6$ learns about path $P_3$ through EBGP session. We can see that according to the IGP cost to the NEXT_HOP node $r_1$ selects path $P_2$ and therefore $c_1$ also selects $P_2$ (only path known). Similarly $c_2$ selects $P_3$ and $c_3$ selects $P_1$. Now it is easy to verify that there is a forwarding loop between node $c_1, c_2$ and $c_3$. The problem with this configuration is that there is no reflector in the shortest path between two clients $c_1$,$c_2$ (similarly for $c_2, c_3$ and $c_3, c_1$). This is in violation to condition (*iv*) of Theorem 2.1.

## IV. Conclusion

In this paper we presented examples which show that in IBGP configurations with route reflection, even if one of the conditions mentioned in Theorem 2.1 stated in [1] is violated, there may be persistent oscillations or forwarding loops in the system.

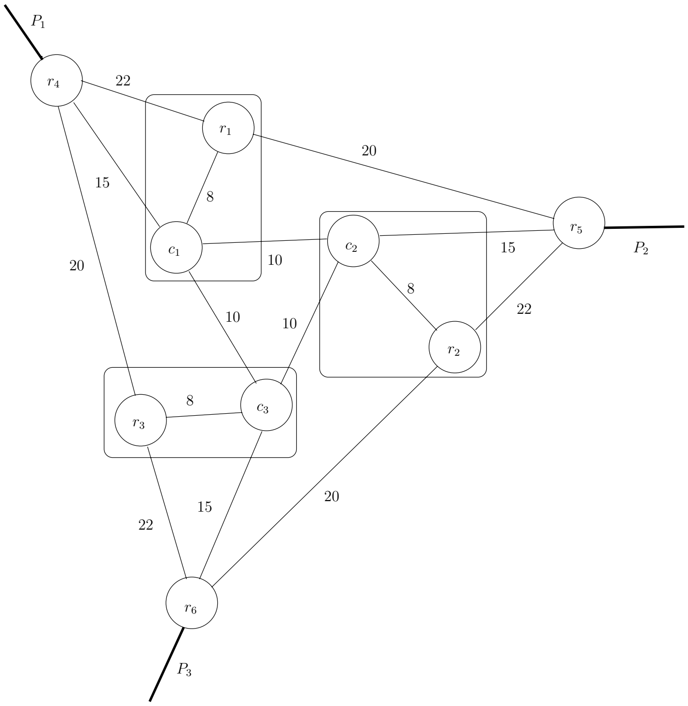


Fig. 4. Condition (*iv*) violated: Forwarding loop

# REFERENCES

[1] A. Rawat and M. Shayman, "Preventing persistent oscillations and loops in IBGP configuration with route reflection," March 2006.