




Bootstrapping evolvability for inter-domain routing with D-BGP



Raja Sambasivan

David Tran-Lam, Aditya Akella, Peter Steenkiste

This talk in one slide

Q What **evolvability features** needed in any inter-domain protocol?



A **Pass-through support**


Multi-protocol structure

D-BGP (BGP w/features): rich, evolvable Internet

The inter-domain routing infrastructure

Allows access to Internet's content (e.g., )

Today, composed of a single protocol, BGP



BGP has many well-known issues

Cannot limit ingress traffic High convergence times

No QoS

Only one best path

ASes can be spoofed



Proposed solutions

Wiser [NSDI'07]

SCION [SP'11]

NIRA [CCR'03]

HLP [SIGCOMM'05]

R-BGP [NSDI'07]

MIRO [SIGCOMM'06]

Arrow [SIGCOMM'14]



BGPsec [IETFv8]

Pathlets [SIGCOMM'09]

EQ-BGP [AINA'06]

BGP has many well-known issues

Cannot limit ingress traffic High convergence times

No QoS


Only one best path

ASes can be spoofed

⋮

BGP is rigid: requires neighbors to use it


Rigidity results in isolated islands



- AS supports new protocol
- AS supports BGP
- Island

Isolation dis-incentivizes deployment


Skirting rigidity with data-plane tunnels



- AS supports new protocol ○ AS supports BGP
- Island - ➔ Tunnel path - ➔ BGP path

Incentivizes non-deployers to fight evolution

Key contributions



The two modest evolvability features
Pass-through support **Multi-protocol structure**
Makes data-plane tunneling optional



D-BGP, which is not far from BGP
Only Required 900 lines of code
BGP already includes pass-through support



Characterization of D-BGP's benefits
Enables a rich Internet w/many protocols
Incentivizes adoption by accelerating benefits

How we identified evolvability features

Evolvable Internet

+

R-BGP [NSDI'07]

BGPsec [IETFv8]

MIRO [SIGCOMM'06]

SCION [SP'11]

Wiser [NSDI'07]

Pathlet Routing [SIGCOMM'09]

⋮



BGP → mod. BGP



BGP // Services



BGP → FIA

↓
Reqs

↓
Reqs


↓
Reqs

Global reqs

Pass-through support
(provided by BGP)

Multi-protocol structure











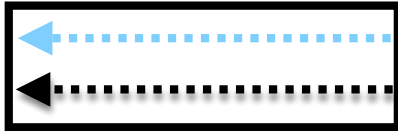
An evolvable Internet



Runs many routing protocols

All ASes support a shared baseline (B)

Taxonomy of evolvability scenarios

	BGP → mod. BGP	BGP // Services	BGP → FIA
Properties	Extra ctrl info	✱	Different ctrl info
Ex.	Wiser, R-BGP	MIRO, Arrow, ✱	SCION, HLP, Pathlets
Incentives			
Deployers	 Inc. benefits	 Profits	 Inc. benefits
Non deployers	 Joint control	 Future profits	 Joint control
Reqs	 Send across gulfs  Send in-band	 Enable discovery	 Send across gulfs  Send in-band

Evolvability scenarios (FIA)

BGP → FIA

Properties

Different ctrl info




E.g., extra paths
or link states

Ex.

SCION, HLP, Pathlets

Incentives


Deployers



Non deployers




Reqs




Send across gulfs




Send in-band




Deploying SCION, a FIA protocol




Deploying SCION, a FIA protocol



Deploying SCION, a FIA protocol



Global reqs for an evolvable Internet



Inform islands about
protocols on paths




Provide common
denominator for e-e paths


Features

Requirements


Pass-through support



Multi-protocol data structure



Disseminate across gulfs



Disseminate in-band



Enable discovery




Inform islands about protocols on paths



Provide common denominator for e-e paths

Outline



Evolvability features



D-BGP design



D-BGP eval

D-BGP overview



BGP advs
with

Multi-protocol structure

Integrated advs (IAs)




BGP processing
with

IA support & pass-through

IA processing



Island




Island

D-BGP's integrated advertisements

Dest. address: 128.2.42.52/24

Path vector


AS # Island ID - - - ➔ Abstracts within-island paths




Prevents ASes from discounting end-to-end paths that include within-island paths

An IA for a path

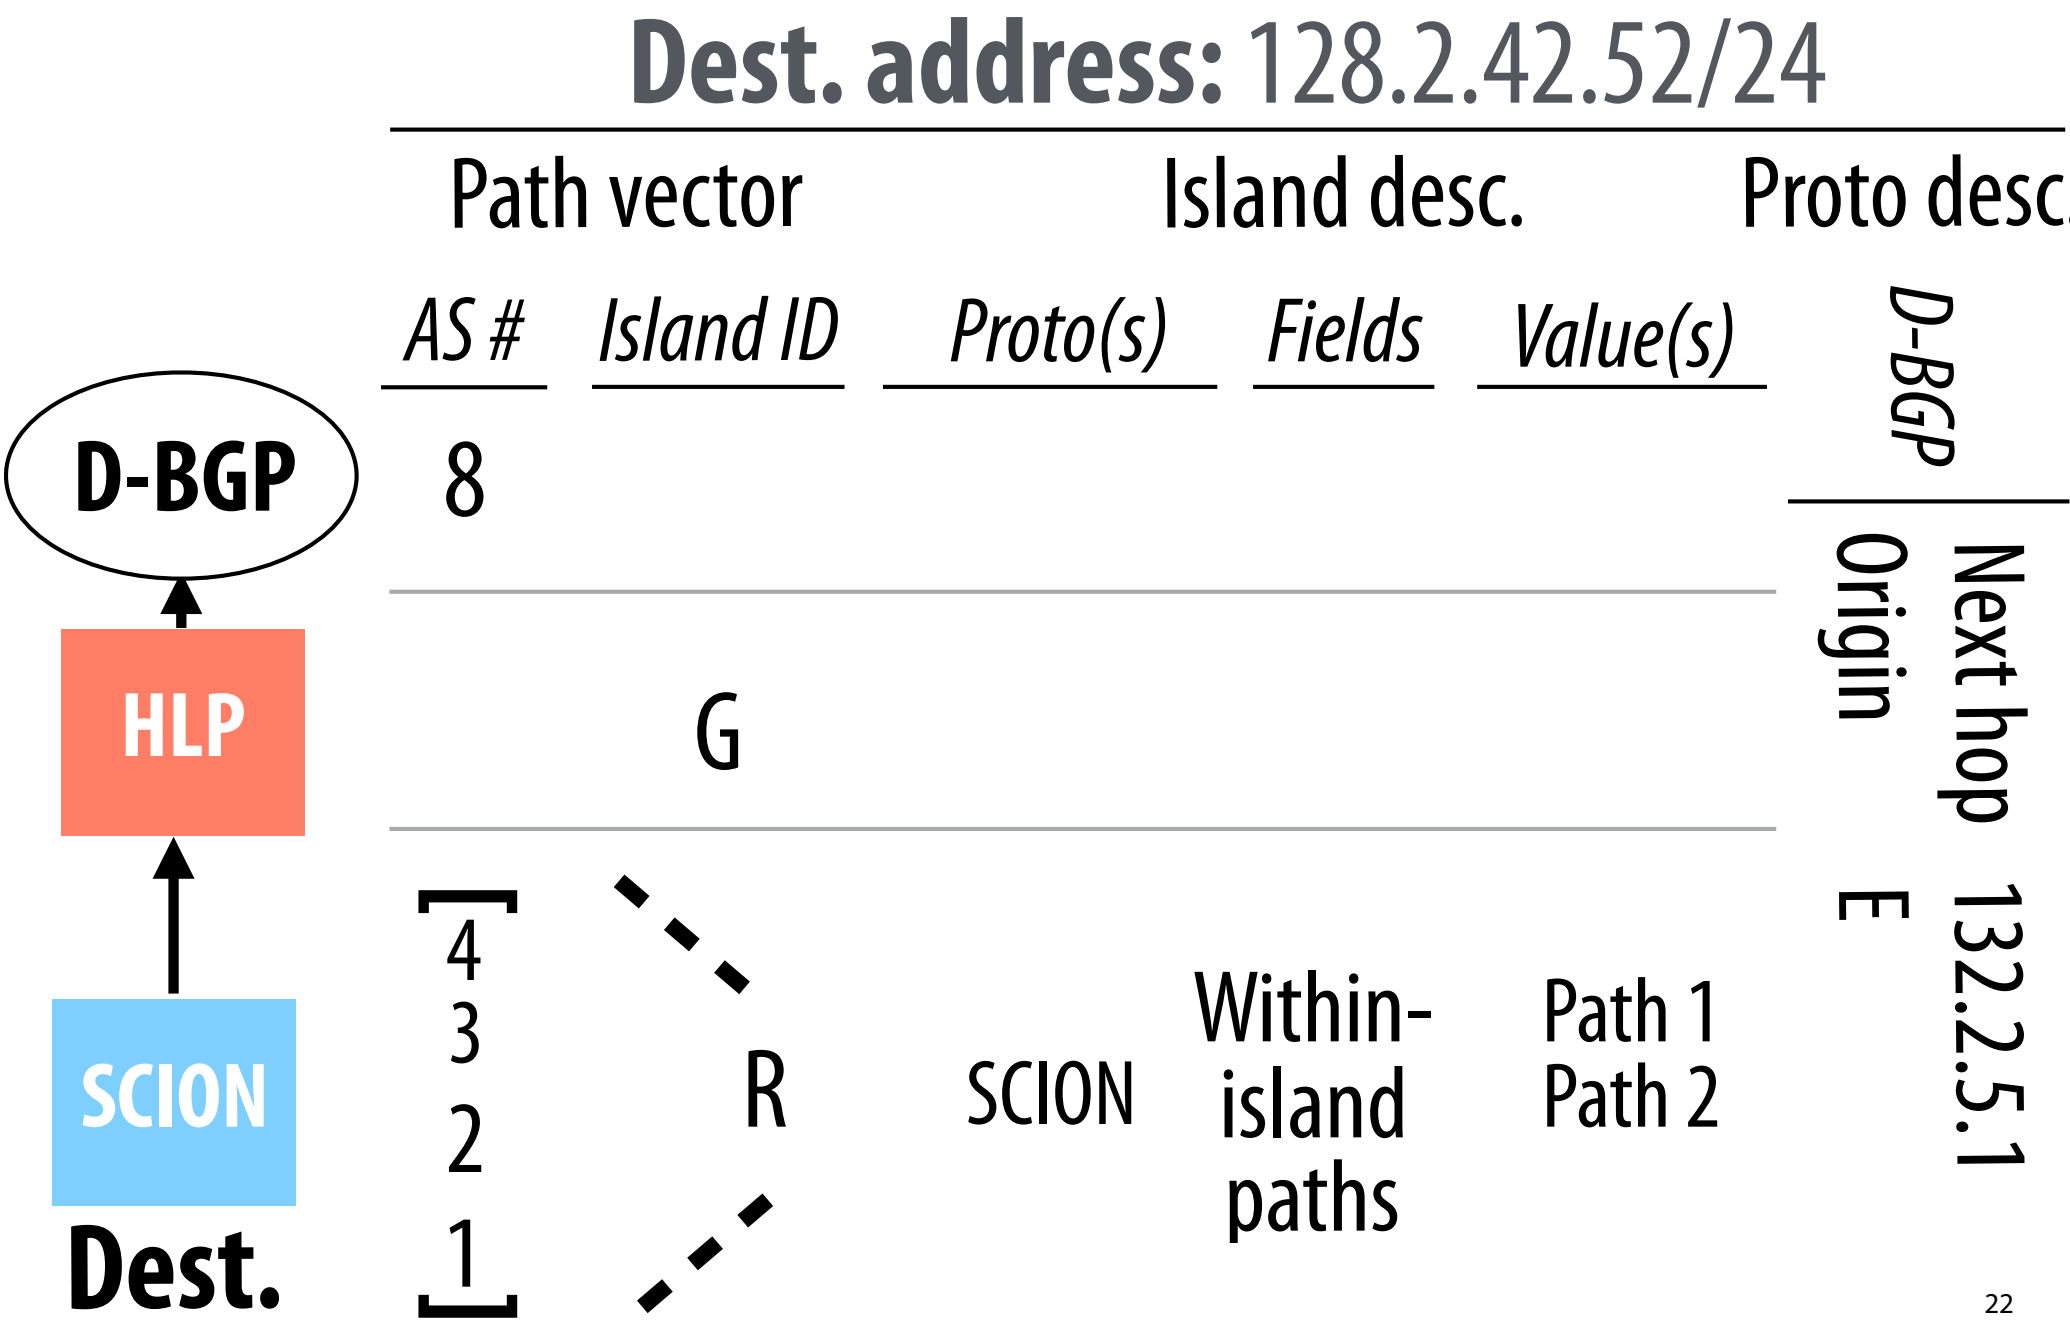
Dest. address: 128.2.42.52/24




Path vector	
<i>AS #</i>	<i>Island ID</i>
8	
<hr/>	
	G
<hr/>	
<div>4 3 2 1</div>	<div>R</div>



An IA for a path



Outline



Evolvability features



D-BGP design



D-BGP eval

Accelerating benefits

Control-plane overhead

Quagga implementation

New-protocol deployments

Accelerating benefits evaluation

Compared deployment in an Internet with:



D-BGP



BGP

Explored benefits as function of adoption


E.g., # paths to dests at upgraded edge domains

Experiments done in simulation

Used Brite [Mascots'01] to generate 1,000-node topology

Used modified routing simulator [SIGCOMM'14]

D-BGP accelerates benefits for SCION



Summary

● BGP's rigidity → Evolvability

● Two features sufficient
for evolvability

● D-BGP provides large
evolvability benefits

