Traffic Engineering: Network Geometry and Design

Alan Shieh
ashieh@cs.cornell.edu

Routing Design in Operational Networks: A Look from the Inside

David A. Maltz, Geoffrey Xie, Jibin Zhan, Hui Zhang,
Gısli Hjalmtysson, Albert Greenberg

Goals and motivation

- Routing design is important, but complicated
- Global system view is essential

High-level challenges

- Real networks are complex systems of interacting routing algorithms
- Configuration information is localized
- Wide range of implementations for the same policy
  - Backbone, enterprise, filtering, …

Methodology

- Empirical analysis of operational networks
  - Extract features from configuration files
- Construct graphs of global relationships
  - Support different goals with different abstraction levels

```
interface Ethernet0
 ip address 66.251.75.144 255.255.255.128
 ip access-group 143 in
!
interface Serial1/0.5 point-to-point
 ip address 66.253.32.85 255.255.255.252
 ip access-group 143 in
 frame-relay interface-dlci 28
!
interface Hssi2/0 point-to-point
 ip address 66.253.160.67 255.255.255.252
!
router ospf 64
 redistribute connected metric-type 1 subnets
 redistribute bgp 64780 metric 1 subnets
 network 66.251.75.128 0.0.0.127 area 0
!
router ospf 128
 redistribute connected metric-type 1 subnets
 network 66.253.32.84 0.0.0.3 area 11
 distribute-list 44 in Serial1/0.5
 distribute-list 45 out
!
router bgp  64780
 redistribute ospf 64 match route-map 8aTzlvBrbaW
 neighbor 66.253.160.68 remote-as  12762
 neighbor 66.253.160.68 distribute-list 4 in
 neighbor 66.253.160.68 distribute-list 3 out
!
access-list 143 deny 134.161.0.0 0.0.255.255
access-list 143 permit any
route-map 8aTzlvBrbaW deny 10
 match ip address 4
route-map 8aTzlvBrbaW permit 20
 match ip address 7
ip route 10.235.240.71 255.255.0.0 10.234.12.7
```

Terminology

- Router information base (RIB)
  - Protocol RIB
    - Filled by routing protocol
  - Local RIB
    - Locally generated
  - Routing RIB
    - Directs forwarding logic

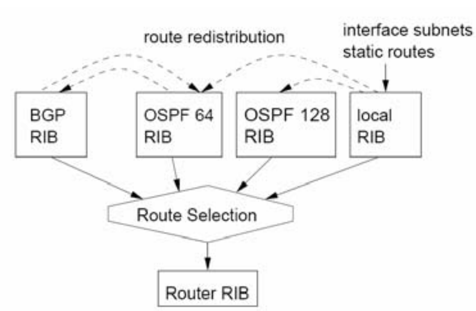

Terminology

- Interior Gateway Protocols (IGP)
  - OSPF, IS-IS, RIP, EIGRP
- BGP
  - External BGP (EBGP): Share information with peers
  - Internal BGP (IBGP): Lossless distribution of EBGP information

Routing Process Graphs

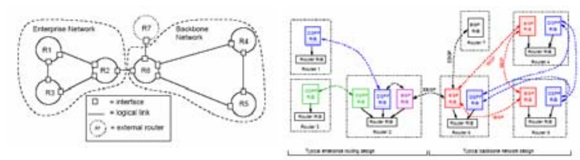

- Goal: Capture low-level interactions between routing processes on different routers
- Nodes represent RIBs
- Edges represent exchange of routes
  - Labeled with route propagation policies

Routing Instance Graphs

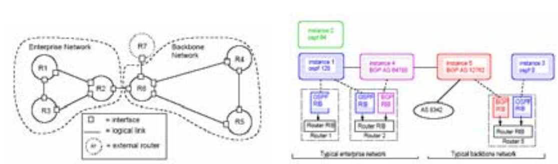

- Problem: Routing process graphs grow rapidly
  - Problematic for large networks
- Routing instance graphs digest information from routing process graphs

Routing Instance Graphs

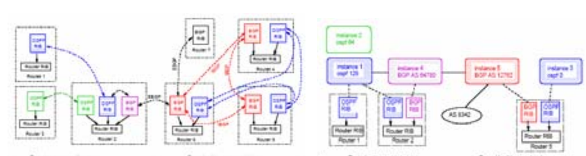

- Nodes represent instances, or collections of interacting Routing Processes
- Interaction between different routing processes running on the same routers?

Route Pathway Graphs

- Summarize interactions between instances from POV of individual router
- Search Routing Instance Graphs based on routes in Router RIB
- Ambiguity: Is this search conservative or precise?

Key implementation details

- Anonymization reduces barriers to participation
- Infer address space allocation, interior/exterior interfaces
  - Use out-of-band membership information of configuration files
  - What about out-of-band IP block assignments?

Case studies

Net5 (backbone)

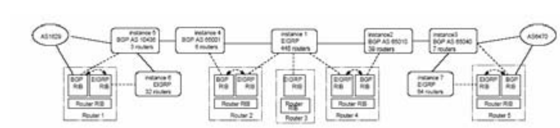

- Backbone uses IGP rather than IBGP
  - Too many routers for simple IBGP mesh
  - Use clever address allocation to increase expressiveness of IGP route map

Net15 route filter analysis

- Edge labels succinctly describe policies
- Help to determine:
  - Reachability
  - Filters on BGP route injection improve IGP scalability

Practical challenges

- Anonymization filter loses information
  - Configuration file design issue?
- Missing information from neighboring domains
  - Filtering from peers, side channels for BGP information, ...

Contributions

- Simple, elegant abstractions leading to practical, useful global analysis tools
- Analysis of operational networks using these tools
- Router configuration dataset

Discussion questions

- Comparison/validation with global configuration tools?

A First-Principles Approach to Understanding the Internet's Router-level Topology

Lun Li, David Alderson, Walter Willinger, John Doyle

Goals

- Demonstrate problems with degree-based random graph models

Technology constraints

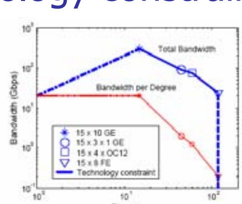

- Routers do not scale perfectly with respect to degree
  - Low degree: limited by per-interface capacity
  - High degree: extra overhead for switching fabric

Economic constraints

- Long-distance links are expensive
- Hierarchical multiplexing is cost-effective

Pitfalls of previous work

- Measurement studies: link layer interference
  - Ethernet/optical rings increase degree
  - MPLS conflates physical multihop routes into single virtual hop
- Graph generators: shortcomings of existing metrics

Novel metrics

- Performance metric
- Likelihood metric

Performance metric

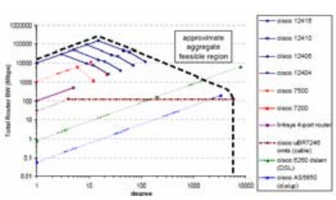

- Generate traffic demand matrix (x)
- Compute route for each pair (R)
- Compute router capacity based on degree (B)

Performance metric

$$\max_{\alpha} \quad \sum_{ij} \alpha x_i x_j$$
$$s.t \quad RX \leq B,$$

- Solve LP, compute aggregate bandwidth
- What about traffic engineering?

Likelihood metric

$$l(g) = (L(g) - L_{\min})/(L_{\max} - L_{\min}),$$

Evaluation

- Generic Random Graph, Preferential attachment
- Heuristically optimal topology (HOT)
  - Degree-preserving rewiring for explicit aggregation
- Abilene-inspired topology
- Suboptimal strawman

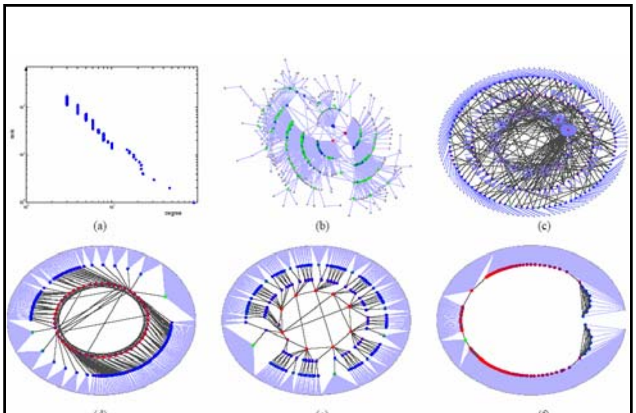

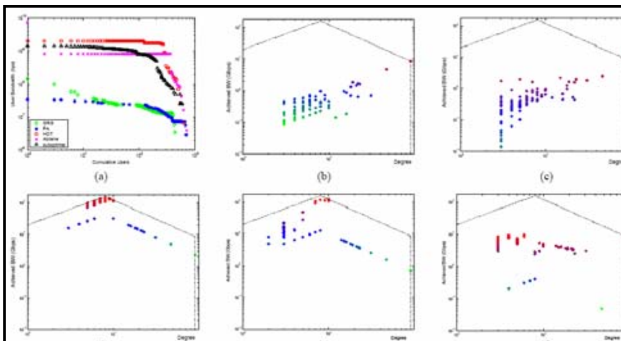

- (A) End-user bandwidth is better for HOT, Abilene, Suboptimal
- (B,C,F) Saturation at well-connected high-degree nodes limits utilization of other links

Performance vs. Likelihood

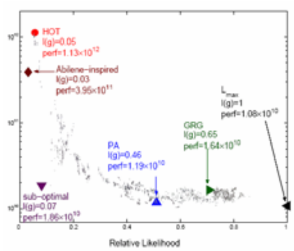

- Likely graphs have low performance
- High performance graphs are unlikely

Contributions

- Degree-based graph models do not model Internet router topology
- Domain-specific results
  - Similar class of results with specific metrics for other domains?

Discussion

- Provocative critique of large bodies of work