

Modeling the Complexity of Enterprise Routing Design

*Xin Sun (Florida International U.),
Sanjay G. Rao (Purdue U.) and
Geoffrey G. Xie (Naval Postgraduate
School)*

The costs of complexity

- “We propose that this trend [towards more complex machines] is not always cost-effective, and may do more harm than good”.
 - Patterson and Ditzel, “The Case for the RISC”, 1980.
- “Complex architectures and designs have been (and continue to be) among the most significant and challenging barriers to building cost-effective large scale IP networks”.
 - RFC 3439

Complex networks are hard to manage

```
class-map match-any QC2
match access-group 102
match access-group ACL2
class-map match-all QC3
match dscp 5 7
class-map match-any CX
```

```
policy-map QP0
class QC2
bandwidth 100
random-detect dscp-base
random-detect dscp 10 40 60
random-detect dscp 12 30 40
class QC3
bandwidth 50
random-detect dscp-base
random-detect dscp 5 20 30
random-detect dscp 7 15
policy-map PX
```

```
interface Ethernet0/1
```

```
service-policy input MarkingPolicy
```

```
interface ATM1/0.1 point-to-point
```

```
class-map match-all QC2
```

```
match dscp 10
```

```
class-map match-all QC3
```

```
match dscp 12
```

```
class-map match-all QC4
```

```
match dscp 5
```

```
class-map match-all QC5
```

```
match dscp 7
```

```
class-map match-all QC6
```

```
match dscp 10
```

```
class-map match-all QC7
```

```
match dscp 10
```

```
class-map match-all QC8
```

```
match dscp 10
```

```
class-map match-all QC9
```

```
match dscp 10
```

```
class-map match-all QC10
```

```
match dscp 10
```

```
class-map match-all QC11
```

```
match dscp 10
```

```
class-map match-all QC12
```

```
match dscp 10
```

Over 80% of IT budget in enterprises devoted to maintaining status quo yet configuration errors account for 62% of network down time, and .. enable 65% of cyber-attacks

(Yankee Group, USITS 2003)

Could we quantify “complexity” ?

“ When deciding between two approaches in networking, complexity is usual an important factor. However, the term ‘complexity’ is rarely well defined, and decisions on complexity are mostly made on subjective terms.”

– IRTF Network Complexity Research Group
Charter, 2011

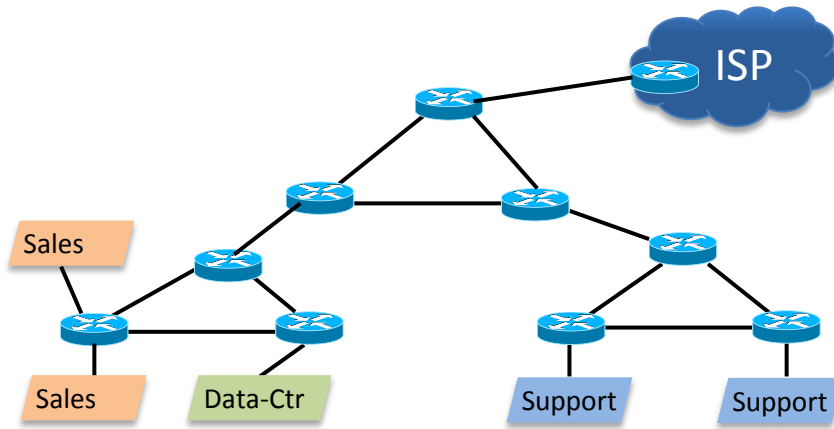
What this paper is about...

- A first framework for quantifying complexity of enterprise routing designs
- Models that relate design to difficulty managing configurations
 - Facilitate design comparisons, what-if analysis
- Focus on Enterprise Routing Design
 - Critical, widely prevalent, time-consuming

Rest of the talk...

- Enterprise Routing Design
- Modeling design complexity
- Modeling details
- Validation
 - Longitudinal snapshots of Purdue's configurations

Routing Design Objectives



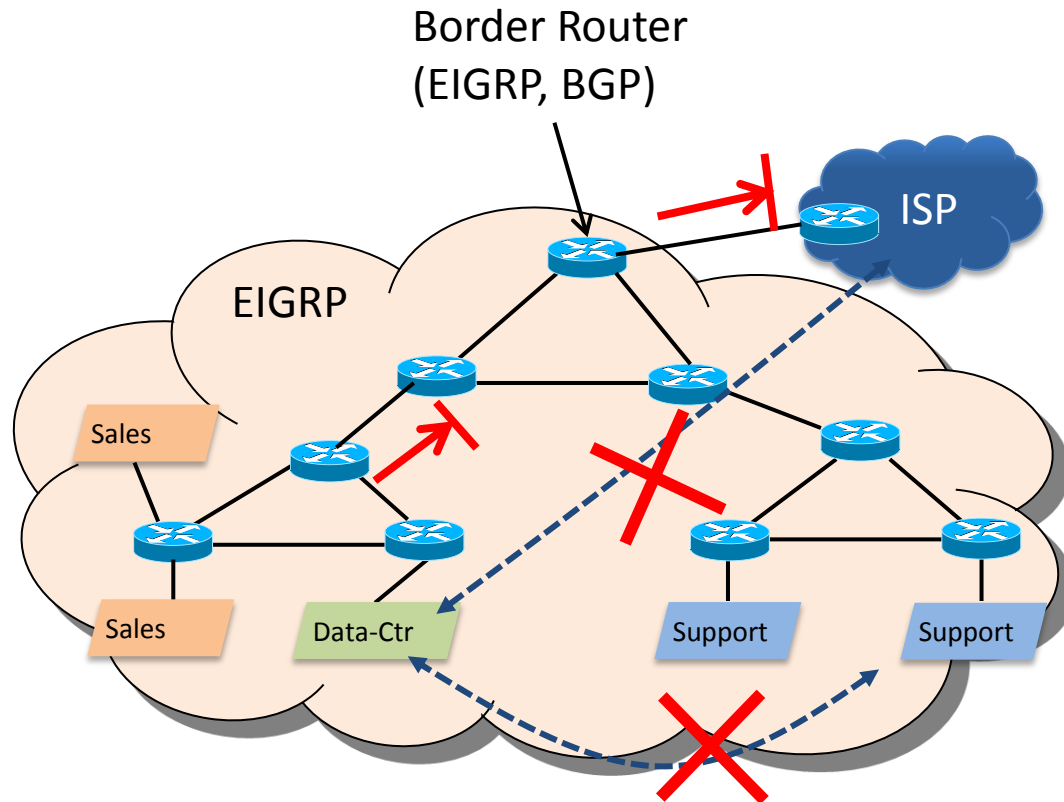
Policy Groups: Subnets with similar reachability policies [variant of IMC09]

	Data-Ctr
Sales	Y
Support	N
INT	N

Reachability Matrix

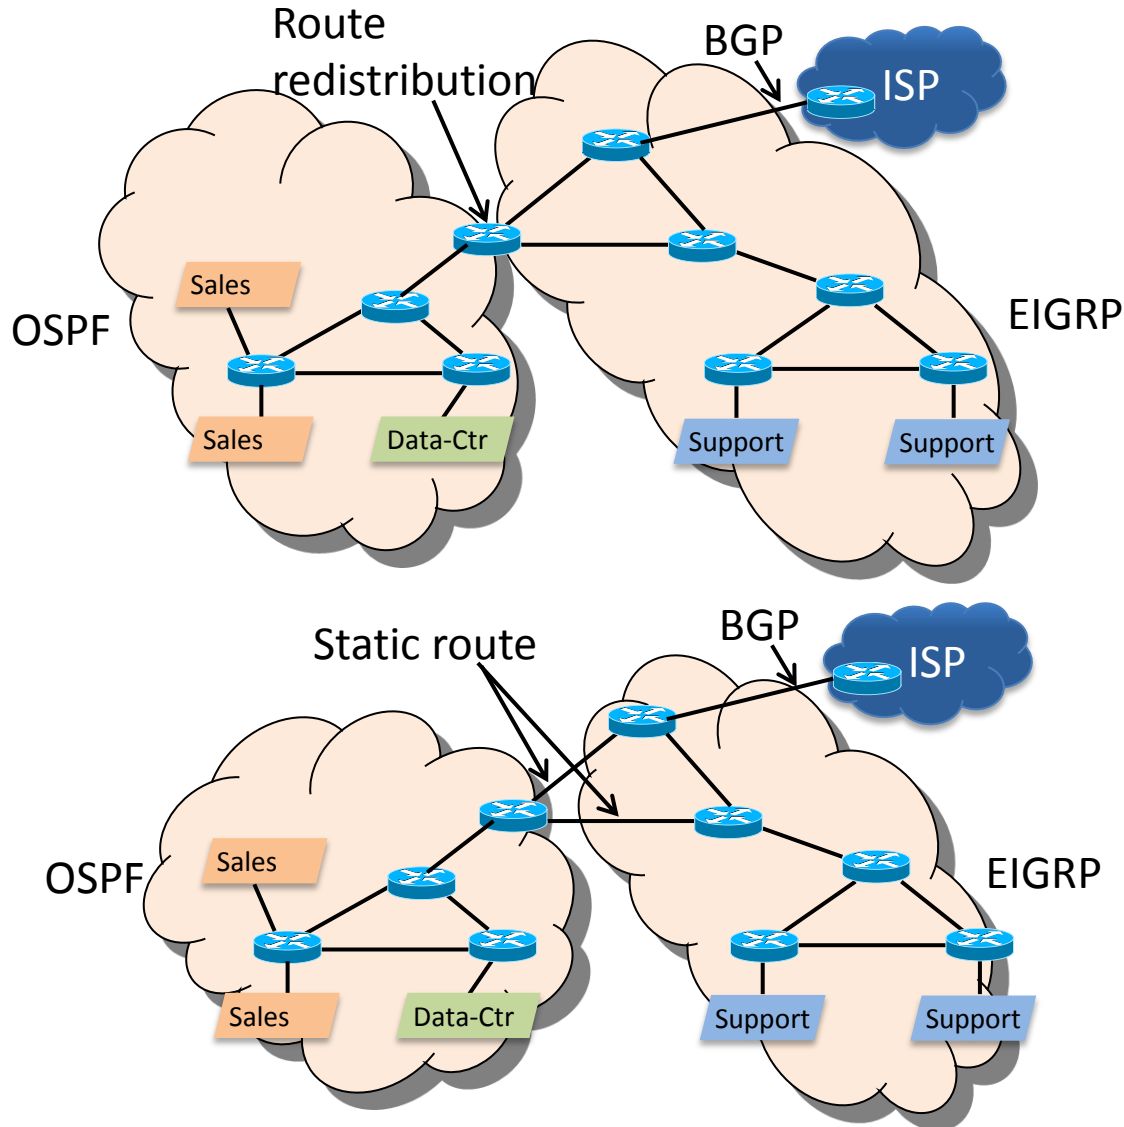
Other objectives: resiliency, traffic engineering etc.

Routing Design Primitives



- [Routing Instance](#) [Maltz et al, Sigcomm 2004]
- Route Filters

Connecting Primitives



Choosing a Routing Design

- Many acceptable choices for operators:
 - Number of instances, mapping routers to instances, connecting primitives etc.
- Design complexity can provide guidance
 - Complexity: important, neglected, subjective
 - Complement performance metrics (e.g., # of hops)

Rest of the talk...

- Enterprise Routing Design
- Modeling design complexity
- Modeling details
- Validation

Prior efforts at quantifying complexity

- Protocol complexity [Chun et al, NSDI 08]
 - Based on state of distributed protocols
 - Dependencies leading to given state
 - E.g. Distance Vector Vs. Link State
- Configuration complexity [Benson et al, NSDI 09]
 - Family of metrics to capture complexity of network configurations
 - Correlation with difficulty managing networks established through operator interviews

Measuring Configuration Complexity

- Key metric: *# of configuration dependencies (referential Links)*

```
class-map match-any QC2
match access-group 102
match access-group ACL2
class-map match-all QC3
match dscp 5 7
class-map match-any CX
.....
!
```

```
policy-map QP0
class QC2
bandwidth 100
random-detect dscp-based
random-detect dscp 10 40 60 10
random-detect dscp 12 30 40 10
class QC3
bandwidth 50
random-detect dscp-based
random-detect dscp 5 20 30 5
random-detect dscp 7 15 20 5
policy-map PX
.....
!
```

```
interface Ethernet0/1
```

```
service-policy input MarkingPolicy
```

```
!
```

```
interface ATM1/0.1 point-to-point
```

```
rate-limit output access-group 102 15 20 20 \
```

```
conform-action set-dscp-transmit 10 \
```

```
exceed-action set-dscp-transmit 12
```

```
rate-limit output access-group 103 2 4 4 \
```

```
conform-action set-dscp-transmit 5 \
```

```
exceed-action set-dscp-transmit 7
```

```
service-policy output QP0
```

```
!
```

```
access-list 102 permit ip any any dscp 10
```

```
access-list 102 permit tcp any any eq www
```

```
access-list 103 permit ip any any
```

```
ip access-list extended ACL2
```

```
permit ip any any dscp 12
```

```
!
```

```
router bgp 1
```

```
no synchronization
```

```
neighbor 10.10.10.101 remote-as 1
```

```
neighbor 10.10.10.101 update-source Loopback0
```

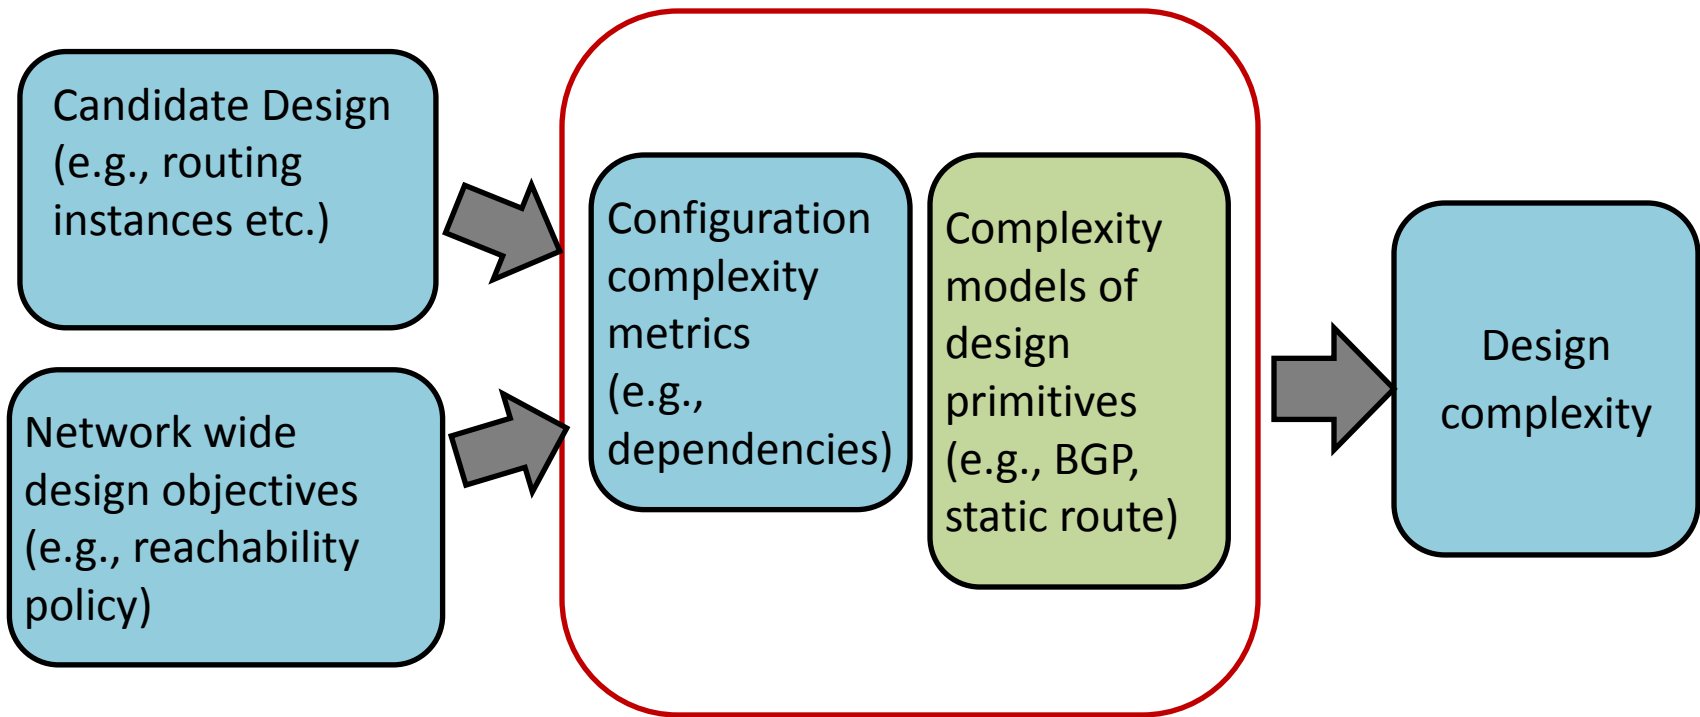
```
no auto-summary
```

```
!
```

Challenge: Network Design Complexity

- Reason about “higher-level” network designs
 - Not just “lower-level” configurations
- Understand sources of complexity
 - E.g., misalignment of routing instances and reachability policies
- What-if Analysis
 - E.g., different set of routing instances ?
 - E.g., replacing static routes with BGP?
- Greenfield network design
 - No access to configuration files

Modeling design complexity



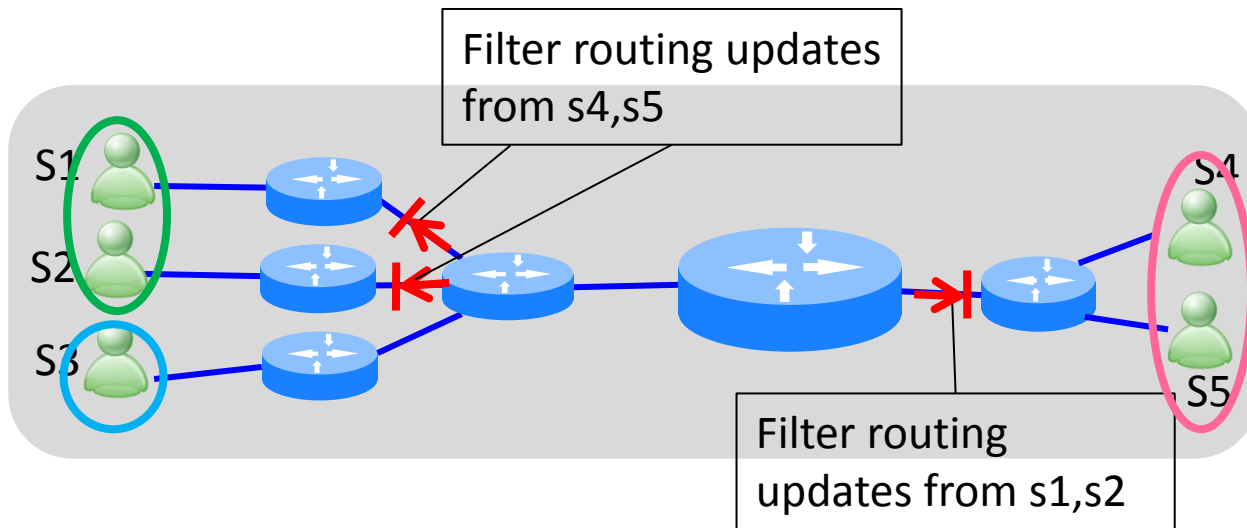
Facilitates green-field design, what-if analysis etc.

Rest of the talk...

- Enterprise Routing Design
- Modeling design complexity
- Modeling details
 - Intra-Instance complexity
 - Inter-Instance complexity
- Validation

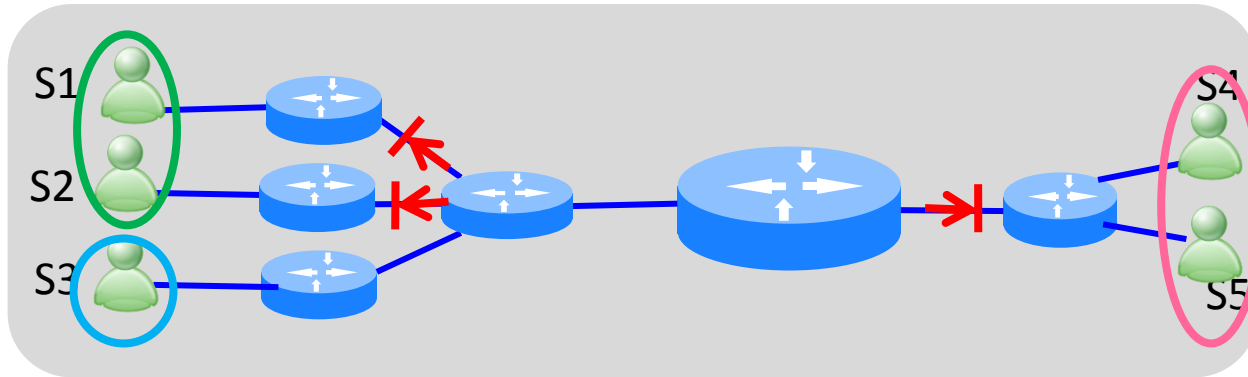
Modeling Single Instance Complexity

- Key cause of complexity:
 - Multiple **policy groups** within an instance



	s1	s2	s3	s4	s5
s1	-	Y	Y	N	N
s2	Y	-	Y	N	N
s3	Y	Y	-	Y	Y
s4	N	N	Y	-	Y
s5	N	N	Y	Y	-

Modeling Single Instance Complexity



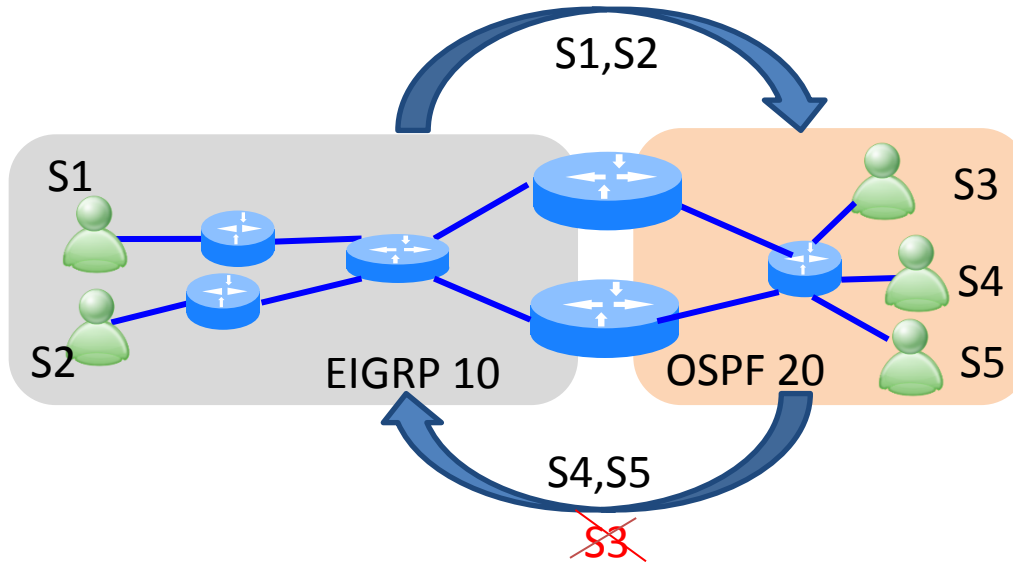
	s1	s2	s3	s4	s5
s1	-	Y	Y	N	N
s2	Y	-	Y	N	N
s3	Y	Y	-	Y	Y
s4	N	N	Y	-	Y
s5	N	N	Y	Y	-

- Complexity depends on:
 - Number of policy groups
 - Topology (# of paths between policy groups, edge-cut sets)
 - # of subnets that must be filtered between policy group pairs
- Estimation details described in paper.

of filters

Filter configuration complexity

Modeling Inter Instance Complexity



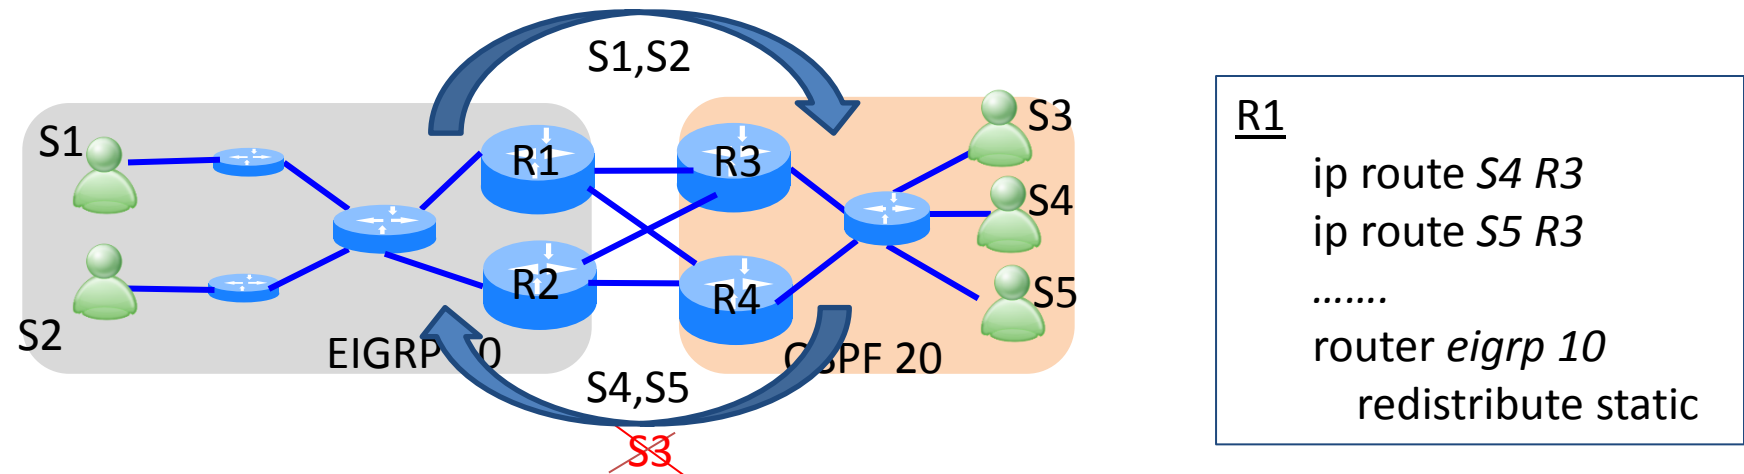
Sources of Complexity: Propagation of routes across instances while meeting

- Reachability requirement
- Resiliency requirement

Different **connecting primitives** may lead to different complexity

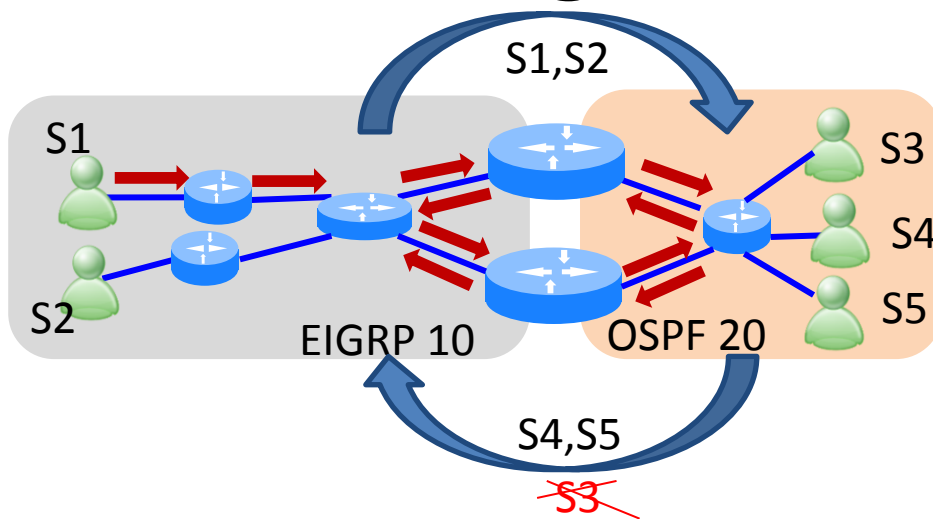
- Route Redistribution
- Static Routes
- BGP

Modeling Static Routes



- Key issue: Failure handling.
 - Configuration for automatic re-routing on failures
- Complexity depends on
 - # of border routers, # of arcs across instances
 - # of propagated routes
 - Basic Propagation, Failure handling

Modeling Route Redistribution



- Key Issue: Preventing Route Feedback
 - Route filters, tags
- Complexity depends on
 - # of border routers
 - # of propagated routes
 - Basic propagation, feedback prevention
 - Fraction of routes propagated

Which primitive lowers complexity?

- Depends on several factors
 - # of border routers
 - # of propagated routes
 - Fraction of routes propagated
- Static Route:
 - Single Border Router, small # of routes
- Route Redistribution
 - Single Border Router, lots of routes, most propagated.
- BGP
 - Multiple Border Routers, most routes propagated

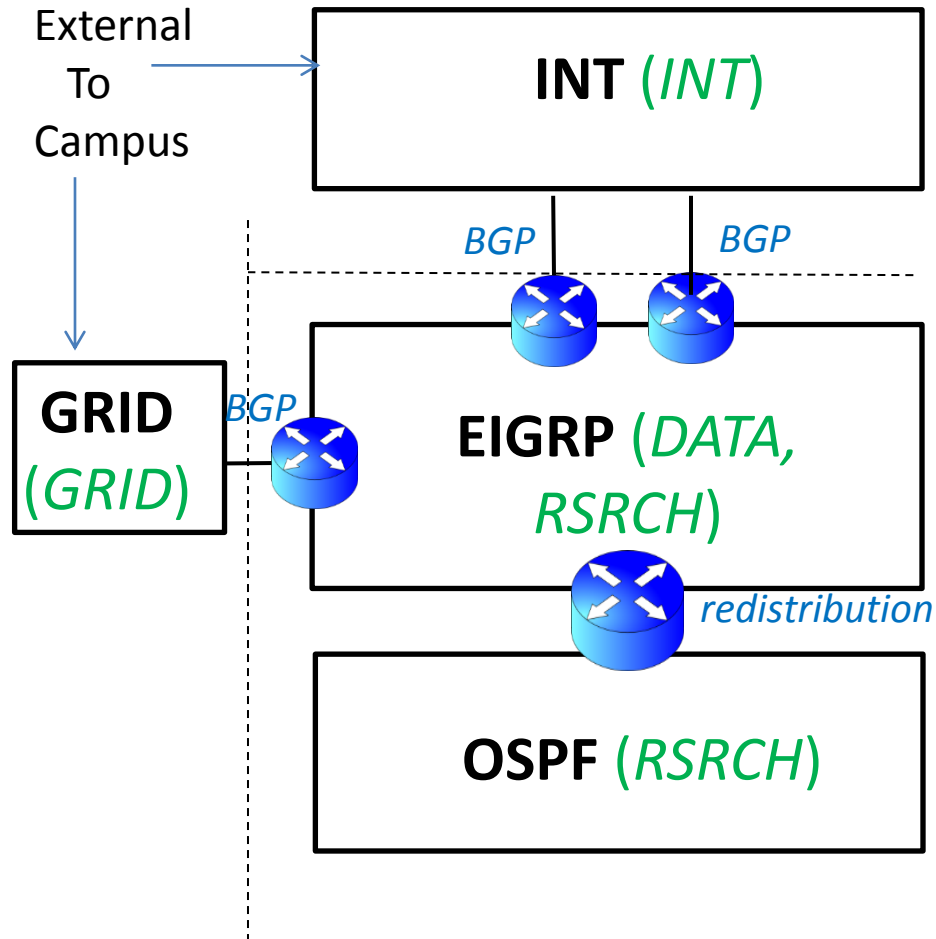
Rest of the talk...

- Enterprise Routing Design
- Modeling design complexity
- Modeling details
- Validation

Evaluation Study Overview

- Data-set
 - Longitudinal configuration snapshots of Purdue
 - 2009 – 2011
 - Major redesign in 2010
 - Physical topology data from CDP
 - ~100 routers, 1000 switches, 700 subnets
- Key Questions
 - Do our models match configuration-based metrics?
 - Yes, see paper
 - Feasible to lower complexity of operational designs?

Purdue Campus Design (2009)

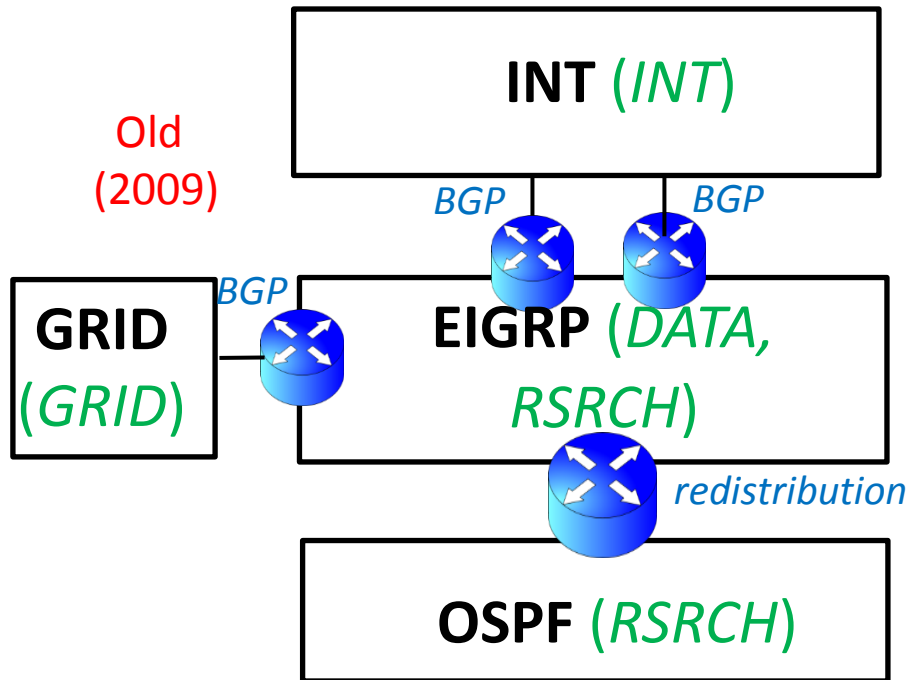


Reachability matrix

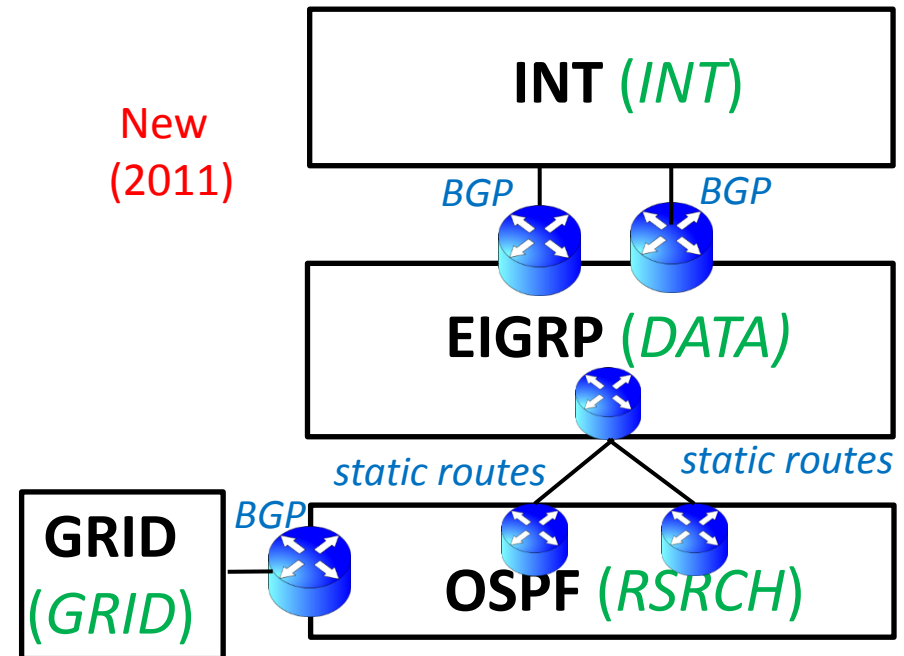
	DATA	RSRCH	GRID	INT
DATA	-	Partial	×	all
RSRCH	all	-	all	all
GRID	×	Partial	-	×
INT	Partial	Partial	×	-

Case Study of a Redesign

Old
(2009)



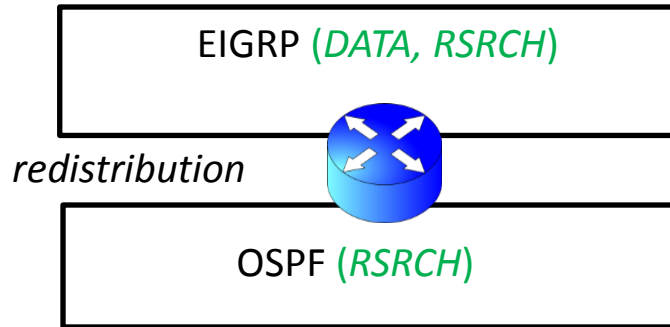
New
(2011)



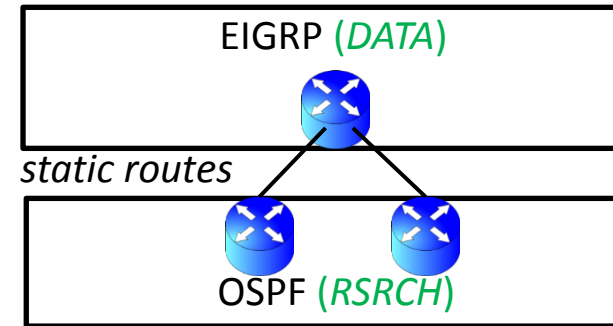
	EIGRP	OSPF	GRID	INT
EIGRP	$\Delta=-7$	$\Delta=29$	$\Delta=-1$	$\Delta=0$
OSPF	$\Delta=1$	$\Delta=0$	$\Delta=1$	-
GRID	$\Delta=-6$	$\Delta=6$	-	-
INT	$\Delta=0$	-	-	-

Δ : new - old

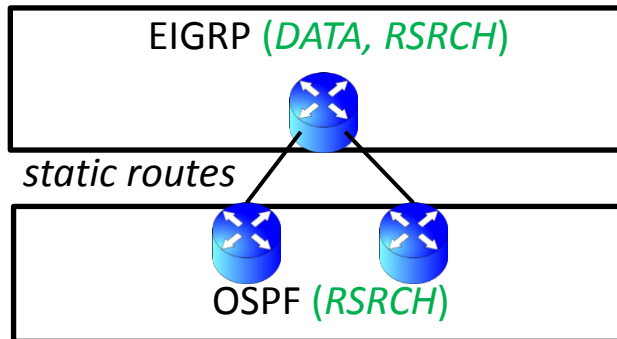
Are There Better Alternatives?



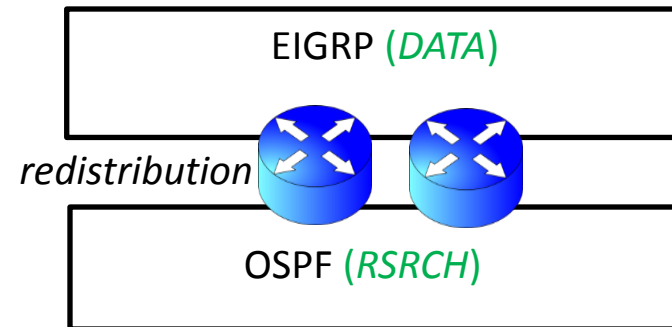
Old



New

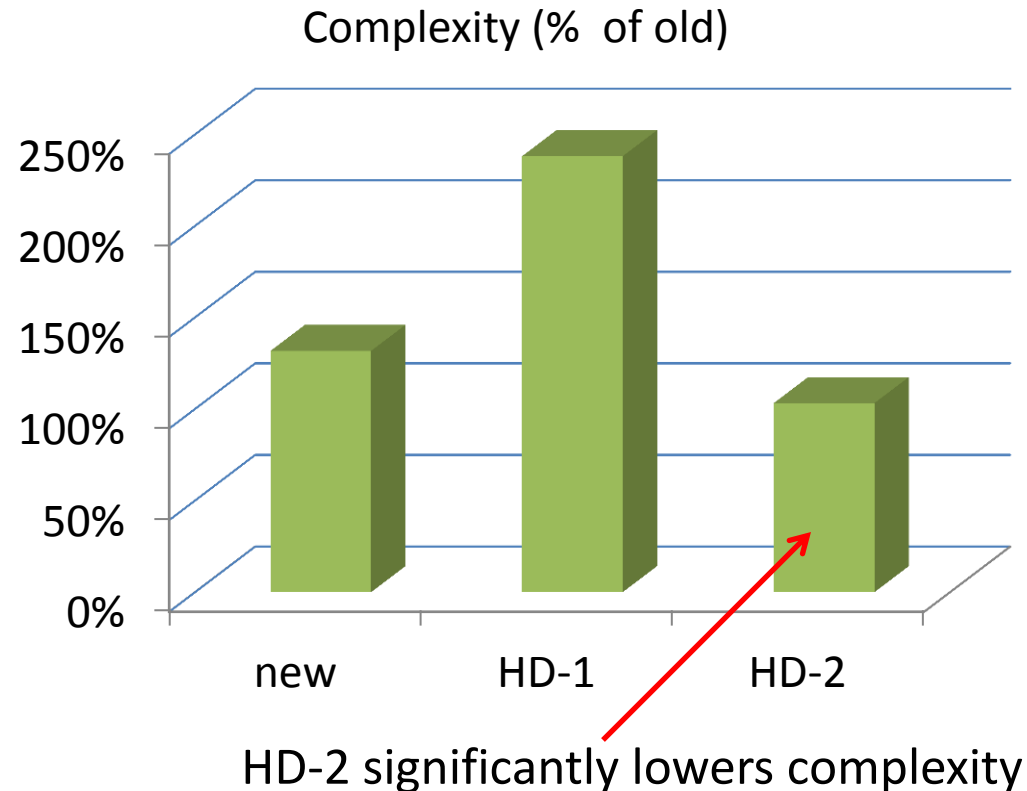
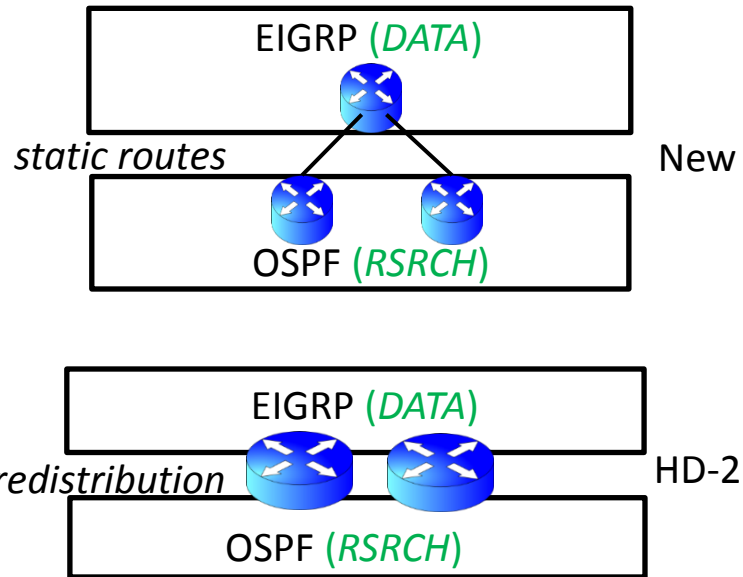


Alternate Design HD-1



Alternate Design HD-2

Are There Better Alternatives?



- Operators confirmed HD-1 would have been the ideal choice
 - However, operator group with diverse skill sets
 - Preferred static routes since less “knowledge” required for students

Conclusions

- Show it is feasible to
 - Quantify complexity of enterprise routing designs
- Distinguishing Aspect:
 - Design Complexity [Vs. Protocol/Configuration]
 - Enables what-if analysis, green-field designs etc.
- Substantial opportunity to lower complexity in an operational network
- Future work: Other design tasks, more complexity metrics, larger-scale validations