

# Minimum Disclosure Routing for Network Virtualization

Masaki Fukushima, Teruyuki Hasegawa, Toru Hasegawa  
KDDI R&D Laboratories Inc.

Akihiro Nakao  
The University of Tokyo

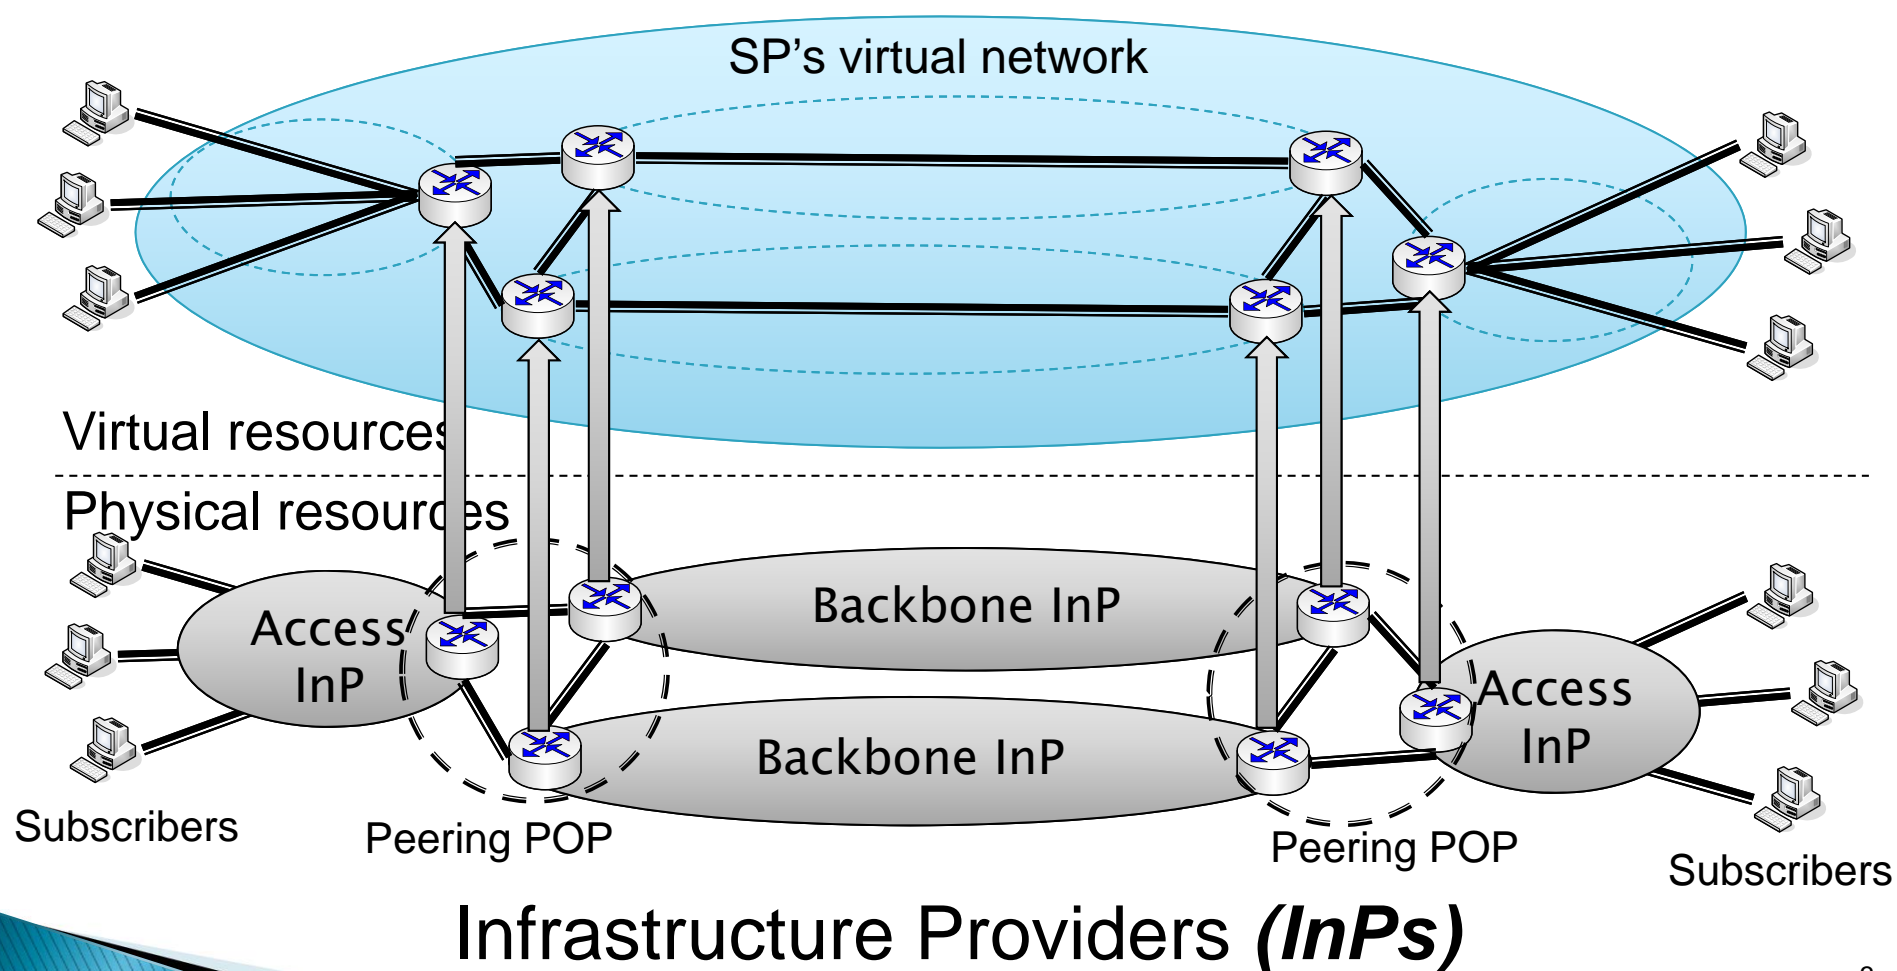
# Outline

- ▶ Background & Motivation
  - Network Virtualization (NV)
- ▶ Problem
  - Minimum Disclosure Routing (MDR)
- ▶ Related Problem
  - Secure Multiparty Computation (SMC)
- ▶ Our solution for MDR
- ▶ Feasibility of our solution
- ▶ Conclusions & Future work

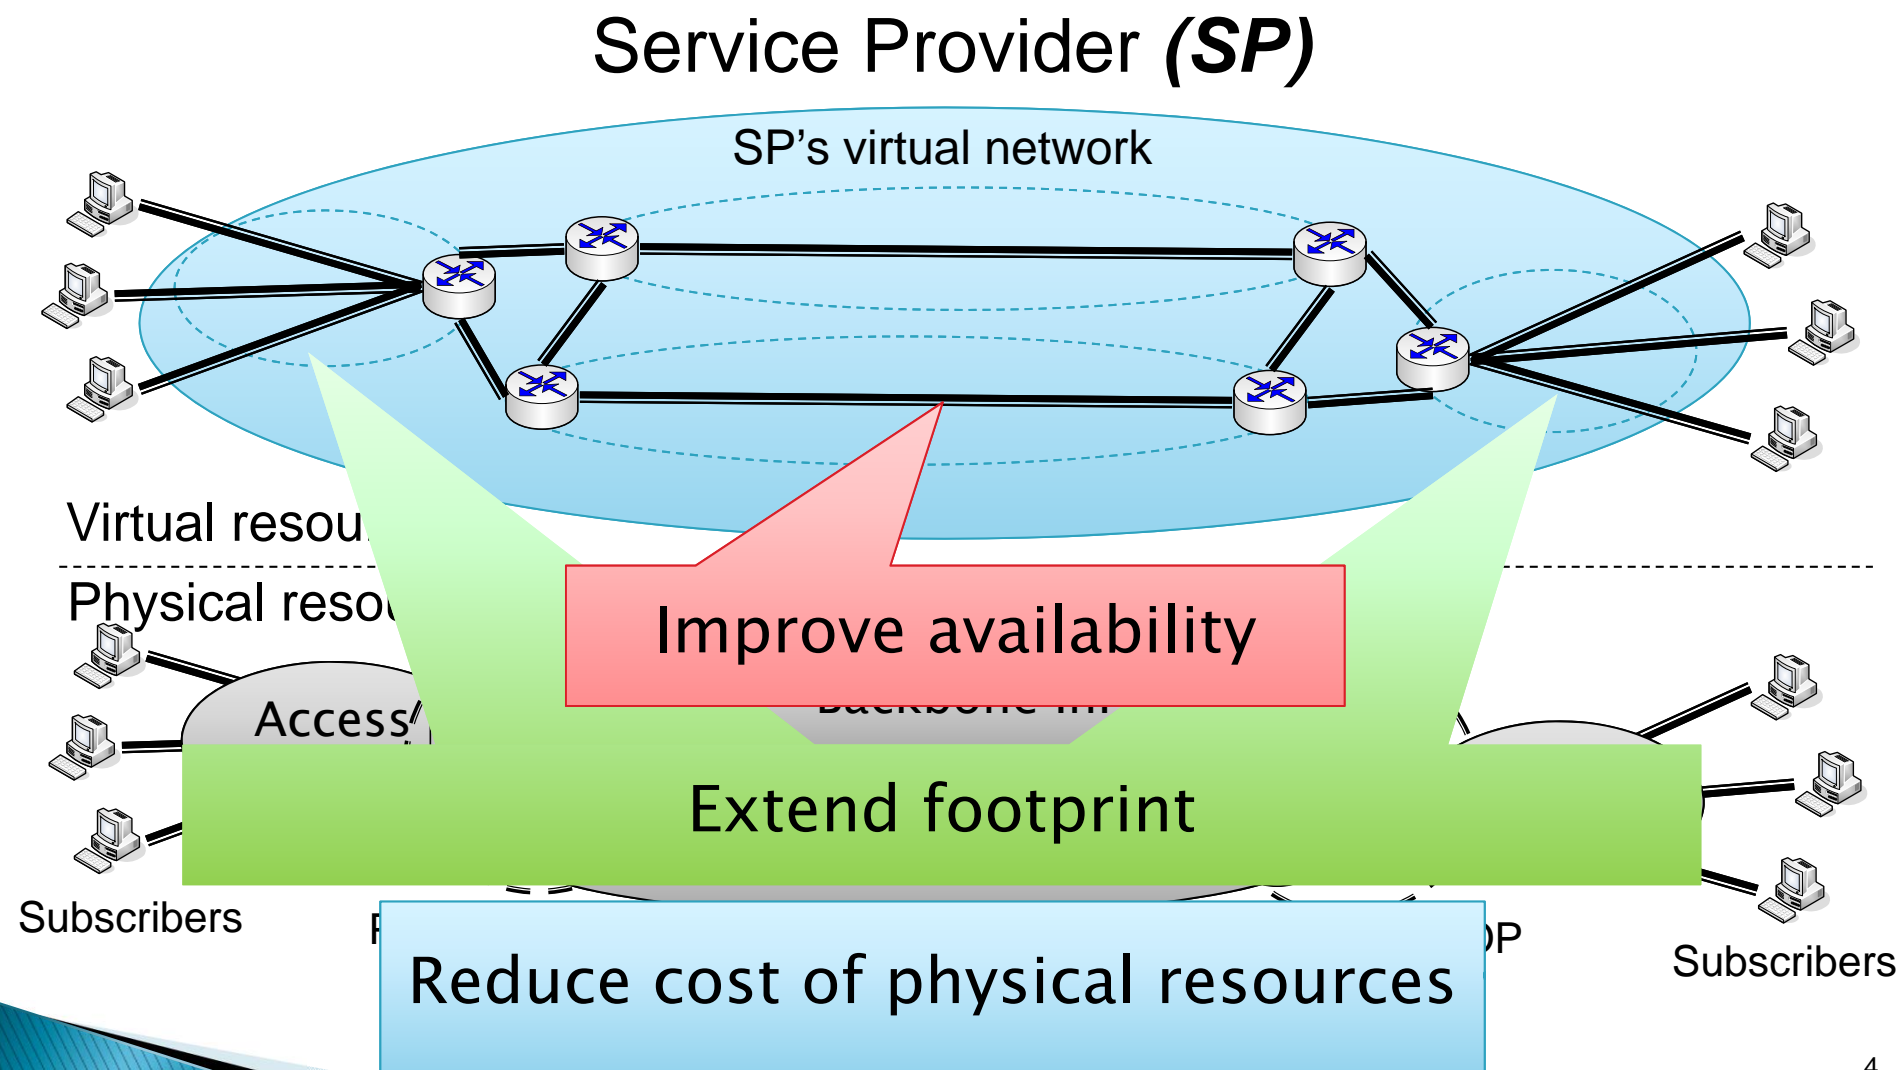


# Background: Network Virtualization (NV) Environment

## Service Provider (**SP**)

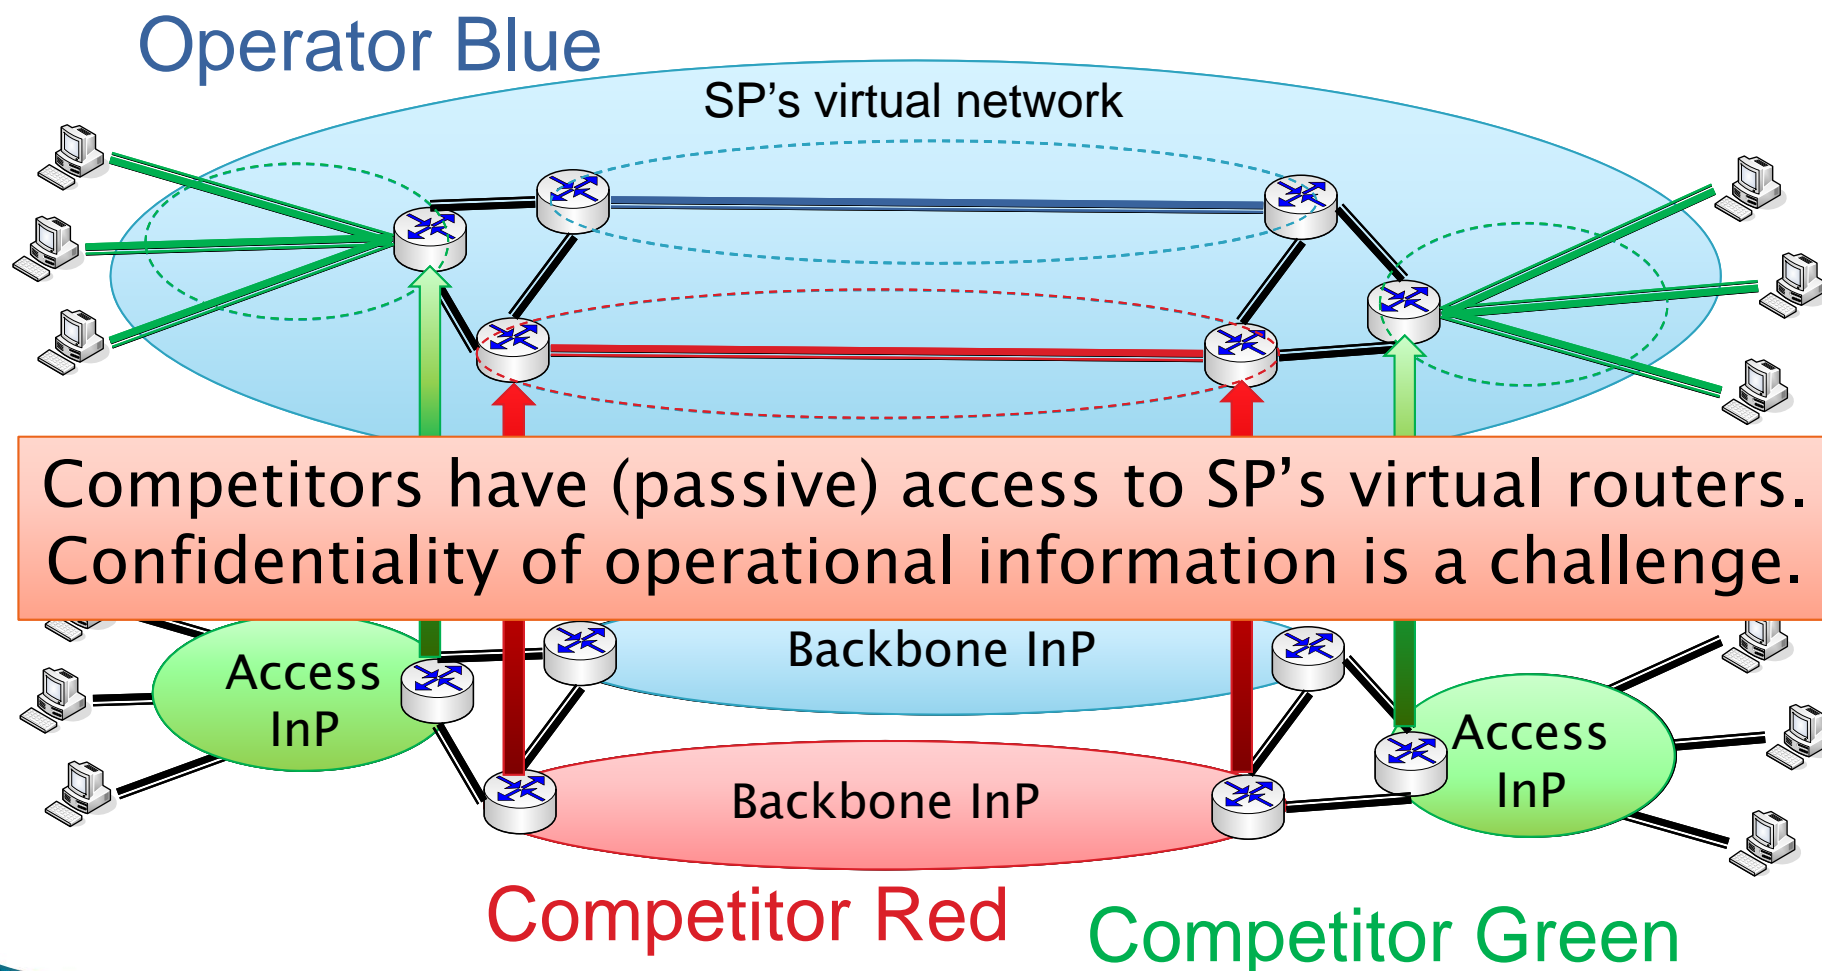


# Background: NV brings merits to SP




# Motivation:

## NV endangers confidentiality of SP



# Motivation:

How **Blue** preserves confidentiality of intra-domain routing?




This router (hosted by **Red**) needs to compute its next-hop for  $dst$ , without knowing  $d_1$  and  $d_2$

This router (hosted by **Green**) needs to compute its next-hop for  $dst$ , without knowing  $d_2$  and  $d_3$



# Motivation:

## Existing IGPs (OSPF, RIP) do not preserve confidentiality




- ▶ Routers exchange routing information
  - including link costs, path costs or even whole topology
- ▶ Underlying InPs can observe routing information
- ▶ Encrypting IGP messages does not help
  - The InPs also have access to the keys on routers

# Problem:

## Minimum Disclosure Routing (MDR)

- ▶ SP's intra-domain routing, where each router
  - Provides local topology information as input
  - Learns next-hop information as output
  - *Learns nothing else*






# Related Problem:

## *Secure Multiparty Computation (SMC)*

- ▶ An example of Secure 3-party Computation

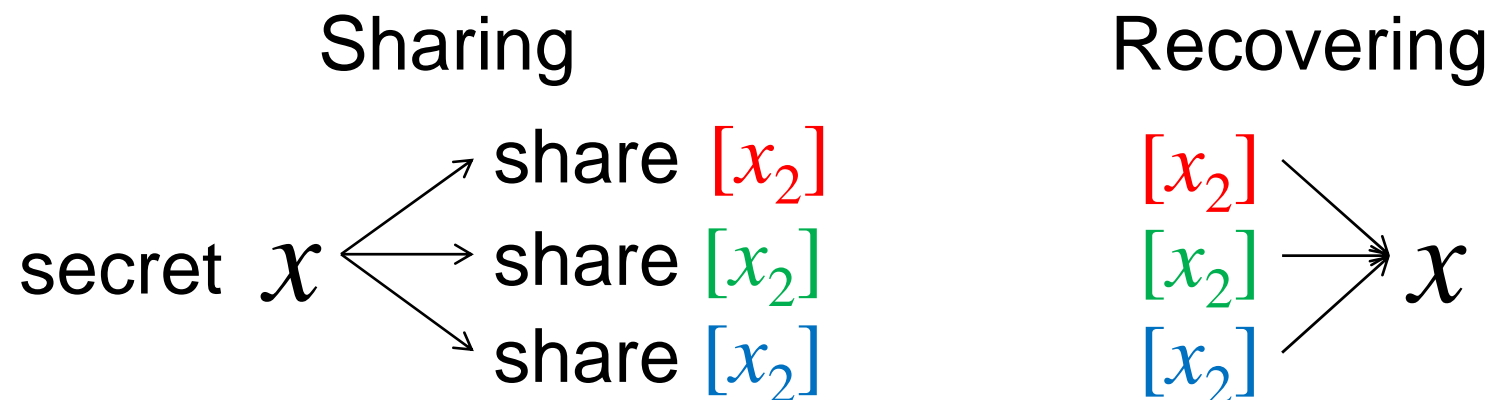


- ▶ MDR is a kind of SMC
- ▶ Some generic SMC protocols are known
  - Applicable to any function  $f$

# Generic SMC protocol uses secret sharing scheme

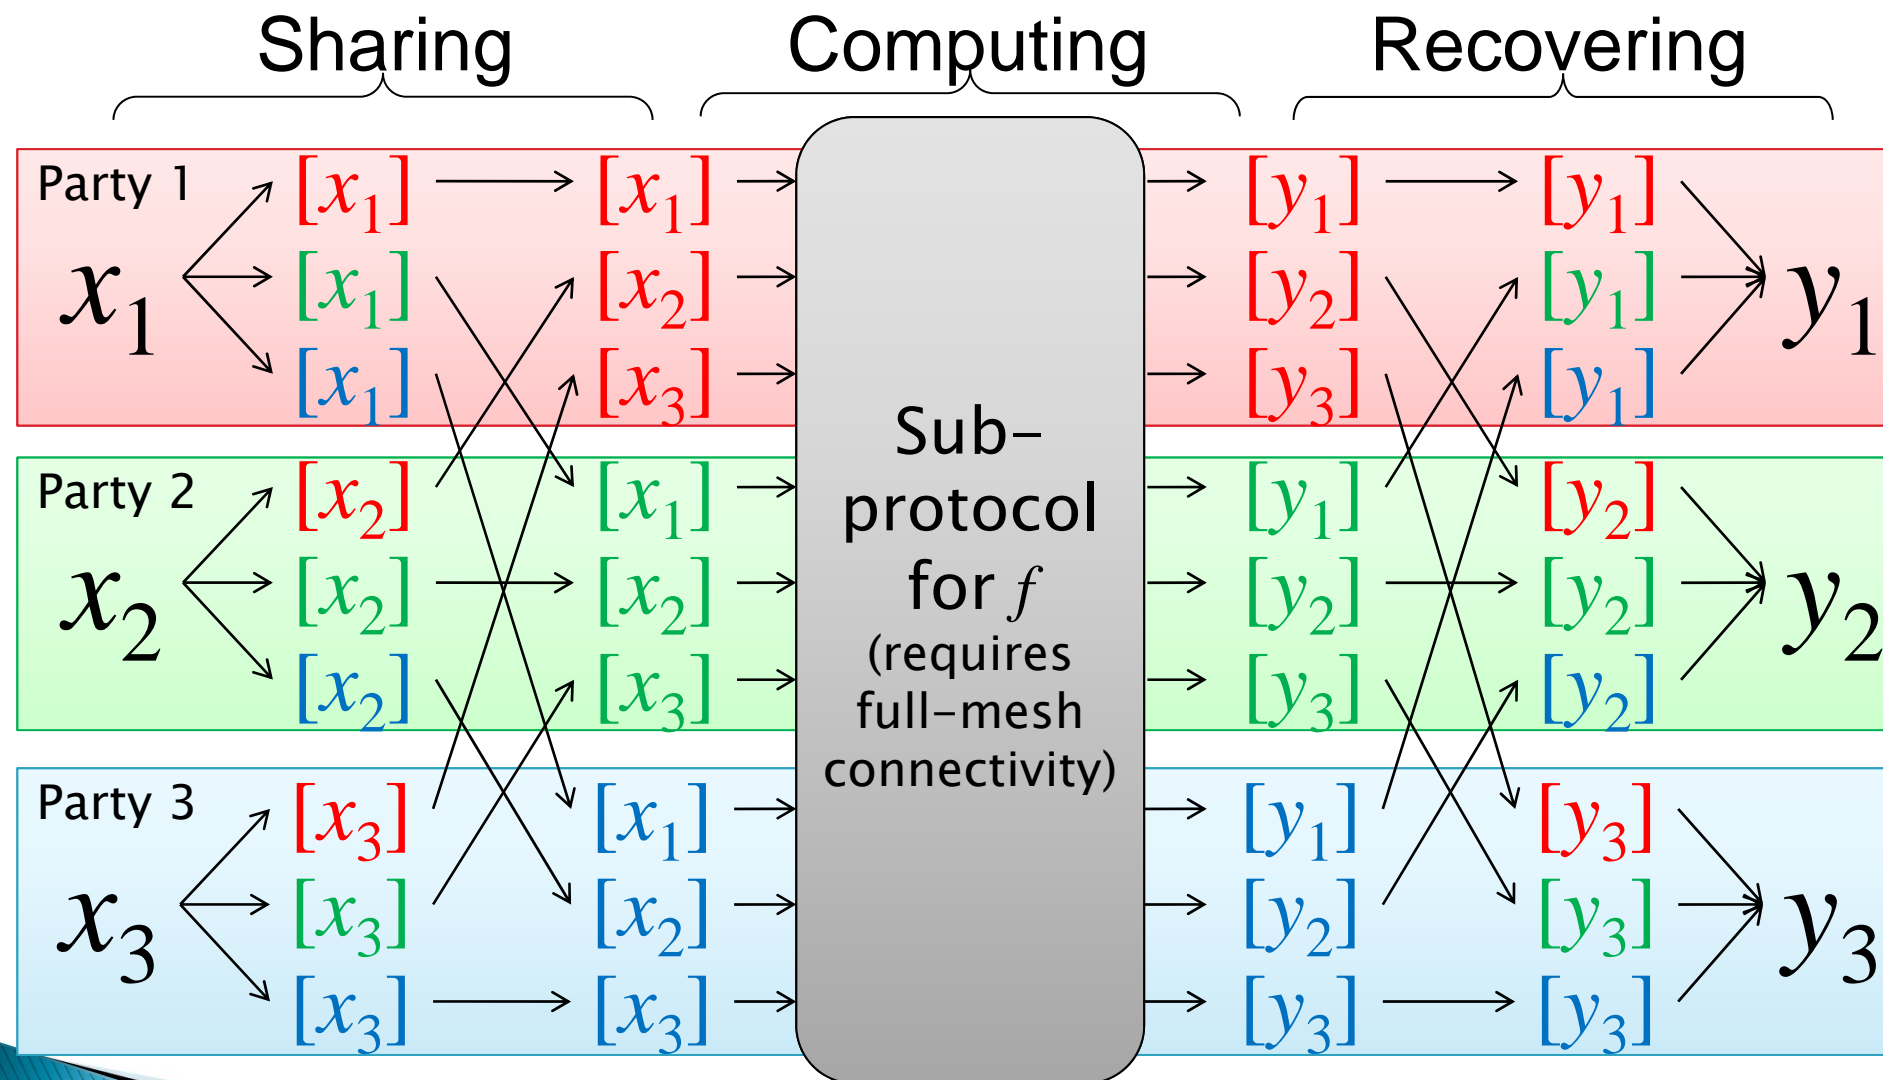
## ▶ *Secret sharing scheme*

- Encode a secret information into multiple fragments called *shares*
- Any single share cannot recover the secret
- All shares can be combined to recover the secret





# Generic SMC protocol for

$$(y_1, y_2, y_3) = f(x_1, x_2, x_3)$$



# Generic SMC protocols cannot be applied to routing

- Generic SMC protocols are applicable only if all parties are **fully connected**




- MDR is a problem for **partially connected** routers to establish such a full-mesh IP connectivity

# Our solution:

## Overview


- ▶ A *clique* virtually works as a big router
  - Fully-connected routers collocated at a peering POP
- ▶ Cliques run a distance-vector routing algorithm
  - In each clique
    - Routing information is encoded into shares
    - Computations are performed by a generic SMC protocol
  - Shares are transferred between neighboring cliques



# Our solution:

## A walk-through scenario (1 / 6)


- ▶ How each router obtains its next-hop for  $dst$
- ▶  $[d_1] = \text{SHARE}(d_1)$ ,  $[d_2] = \text{SHARE}(d_2)$ 
  - Encode link cost into shares
  - Distribute these shares in a clique



# Our solution:

## A walk-through scenario (2/6)

- ▶ **TRANSFER( $[d_1]$ )**
  - Transfer shares of link cost from the left clique to the right clique






# Our solution:

## A walk-through scenario (3 / 6)


- ▶  $[d_1 + d_2] = \text{COMPUTE}([d_1] + [d_2])$ 
  - The right clique run a generic SMC protocol for addition



# Our solution:

## A walk-through scenario (4/6)


- ▶ Similarly, shares of the distance of the other path  $[d_1+d_3]$  is obtained



# Our solution:

## A walk-through scenario (5/6)


- ▶ Suppose  $d_2 < d_3$  (the upper path is shortest)
- ▶  $[upper] = \text{COMPUTE}([d_1+d_2] < [d_1+d_3] ? [upper] : [lower])$ 
  - The right clique run a generic SMC protocol



# Our solution:

## A walk-through scenario (6/6)

- ▶  $upper = \text{RECOVER}([upper])$ 
  - Recover the route in the right clique
- ▶ Each router in the clique directs its route towards the upper path



# Feasibility of our solution:

## Assumptions of evaluation


- ▶ Generalized version of our solution
  - Arbitrary numbers of cliques and destinations
  - Scalability to the number of destinations
- ▶ Metric: latency of SMC protocol to update distances to destinations
  - Most time-consuming part of our solution
- ▶ Analysis model and parameters
  - Computation latency (in each router)
    - GPGPU (1.35 GHz, 240 cores)
  - Communication latency (between routers in a clique)
    - One-way delay is 1 msec
    - Bandwidth is 1Gbps or 10 Gbps



# Feasibility of our solution:

## Latency of SMC protocol

- $K$ : number of neighboring cliques
- $B$ : bandwidth of the links within a clique



- ▶ An invocation of SMC protocol requires less than 100 msec
- ▶ Total Convergence requires less than 1 second
  - Number of invocations are upper-bounded by network diameter
  - Diameter  $< 10$  even in a large Tier-1 network.

# Conclusions

- ▶ NV poses a new problem, MDR
  - confidentiality of operational information
- ▶ None of existing protocols solves MDR
  - Existing routing protocols do not preserve confidentiality
  - Generic SMC protocols cannot be applied to routing
- ▶ We proposed a solution for MDR
  - Extend SMC protocol to routing problem
  - Feasible in a large network if it is run on state-of-the-art hardware





# Future work

- ▶ Implementation
  - Currently implementing our solution by extending Quagga on Linux
- ▶ Evaluation with implementation
  - Preliminary experiment results show that our analysis results are reasonable

