Diamond: Nesting the Data Center Network with Wireless Rings in 3D Space

Yong Cui[1], **Shihan Xiao[1]**, Xin Wang[2], Zhenjie Yang[1], Chao Zhu[1],
Xiangyang Li[1,3], Liu Yang[4], and Ning Ge[1]
[1]Tsinghua University
[2]Stony Brook University
[3]University of Science and Technology of China
[4]Beijing University of Posts and Telecommunications

Data center networking

- Existing DCNs
  - Hierarchical layers
  - Oversubscription
  - Static & symmetrical topology
- Challenges
  - Large-scale: complex cabling & maintenance
  - Dynamic traffic (e.g., random hotspots):

One static & symmetrical topology does not fit All the traffic patterns

Dynamic topology?

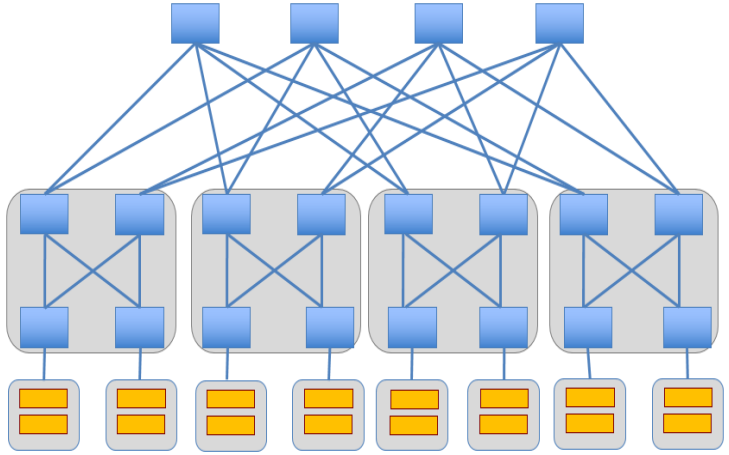

Fat-tree, Mohammad Al-Fares et al. Sigcomm08

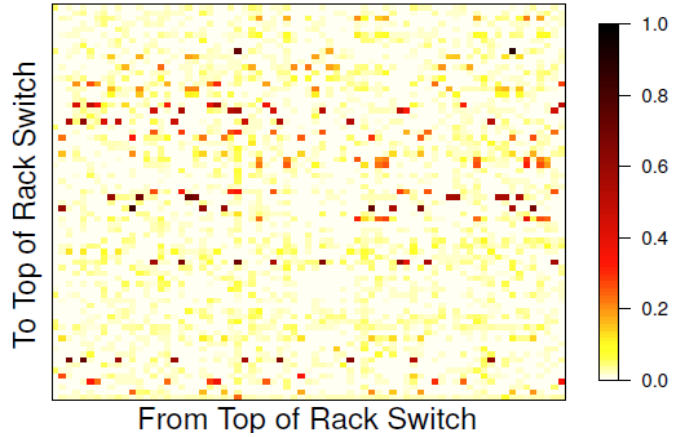

Figure source: Daniel Halperin et al. Sigcomm 2011

Hybrid data center networking

- *Wireless* hybrid networking: *Flyway, 3D-Beamforming, Firefly…*
  - Deploy *directional wireless* radios (60GHz or Free-Space-Optic (FSO)) at ToR
  - Direct rack-to-rack wireless links: built on demand to remove dynamic hotspots

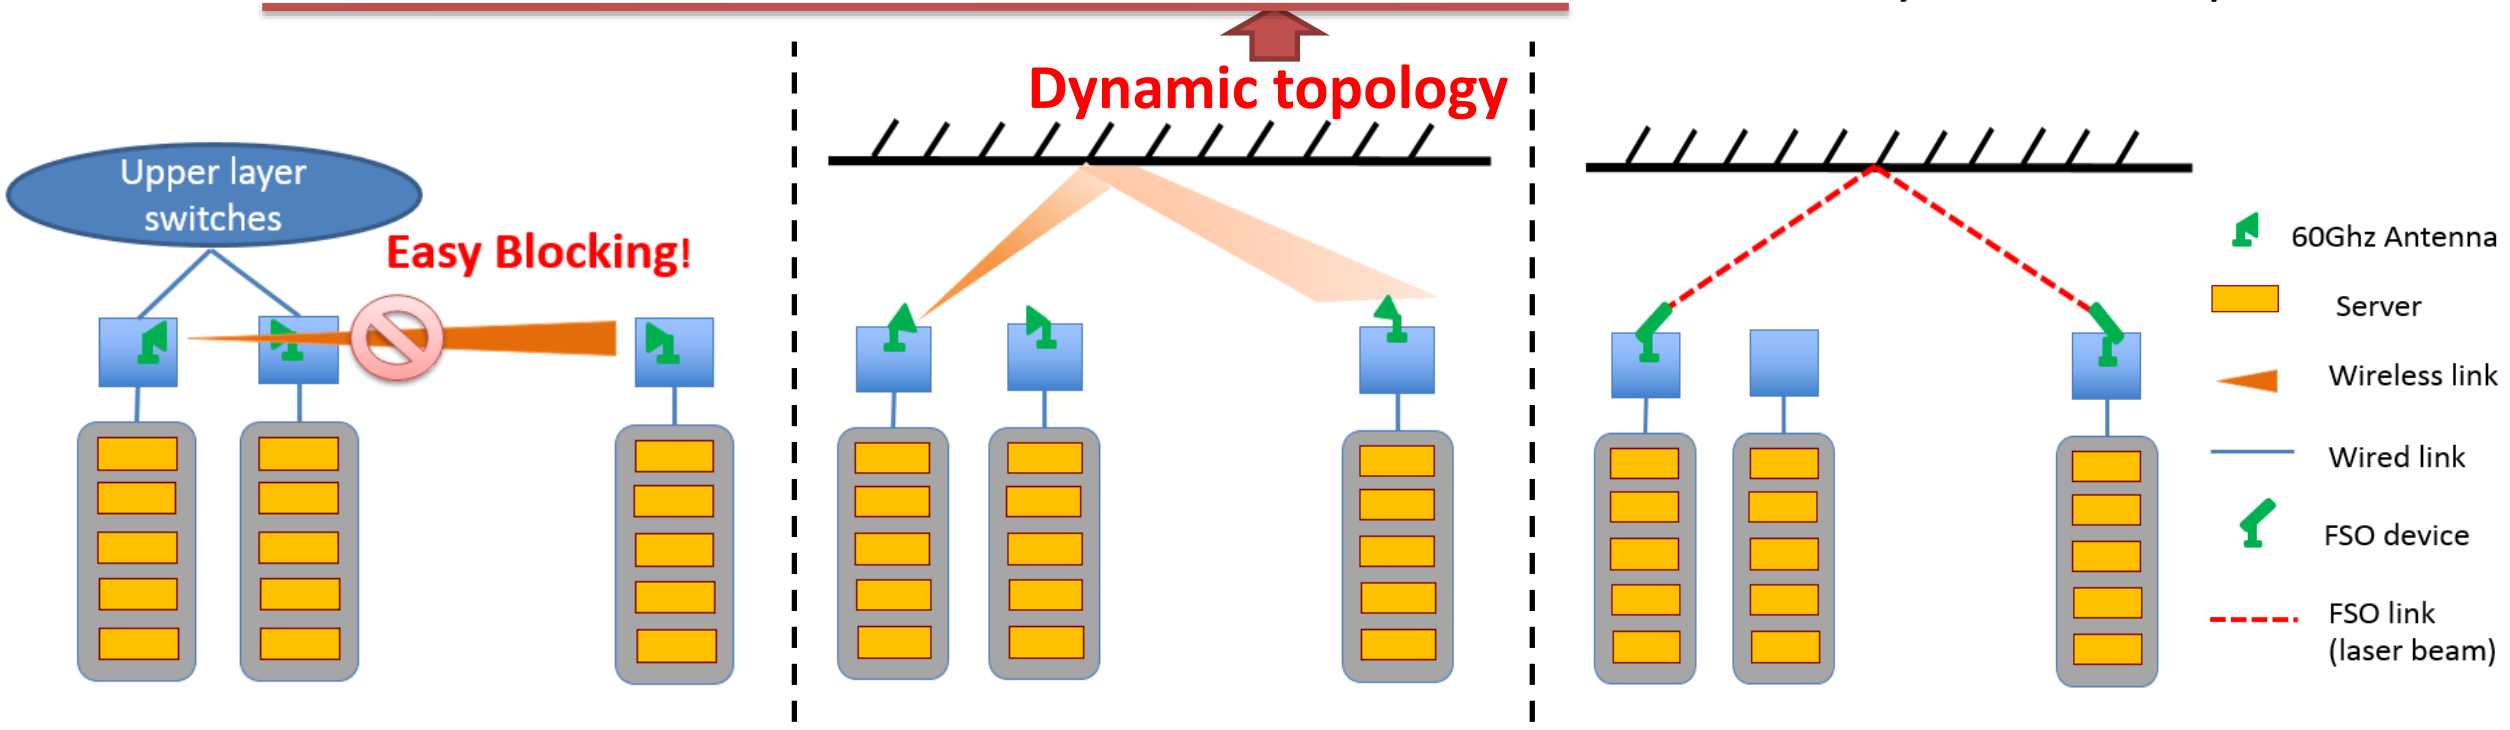

Flyway [Halperin et al, Sigcomm 2011] *3D-Beamforming [Zhou et al, Sigcomm 2012]* *Firefly [Hamedazimi et al, Sigcomm 2014]*

Hybrid data center networking

- Existing wireless hybrid DCNs
  - Wireless radios on top of rack
  - Wireless network on top of existing wired network
  - Rack-level reconfigurable topology to fit dynamic traffic
- Challenges
  - Limited wireless links: small rack size & dense interference
  - Easy blocking: ceiling mirror is unavailable in modern data centers
  - Difficult cooperation: the wired part is kept unchanged, hence hard to cooperate with newly added wireless part

Not hybrid enough!

Challenge—Limited wireless radios & links!

- *Wireless on Top of Rack?*
  - The top of each rack can hold *at most 8* wireless radios
  - Small rack size: more radios on top of rack lead to denser interference
- *Ceiling mirror?*
  - Unavailable mirror: requires a restricted-height (3 meters) clear space above rack
  - Modern data centers: complex steel structures & air conditioner plan above racks

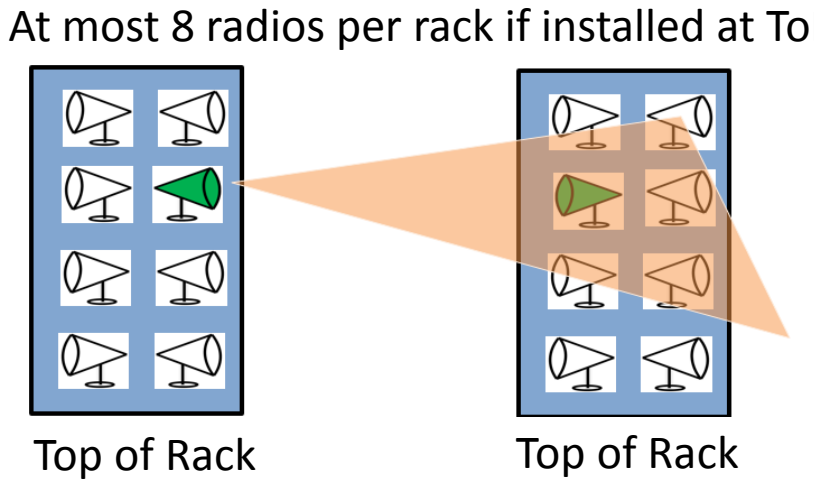

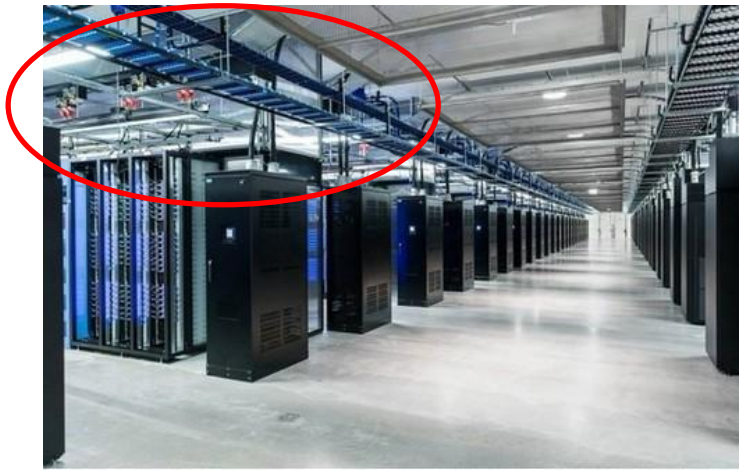

Solution 1—Multi-reflection Ring

- Motivating example

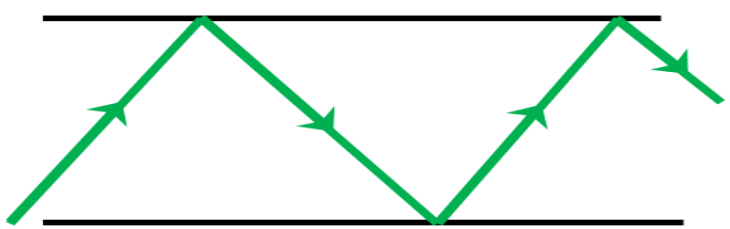

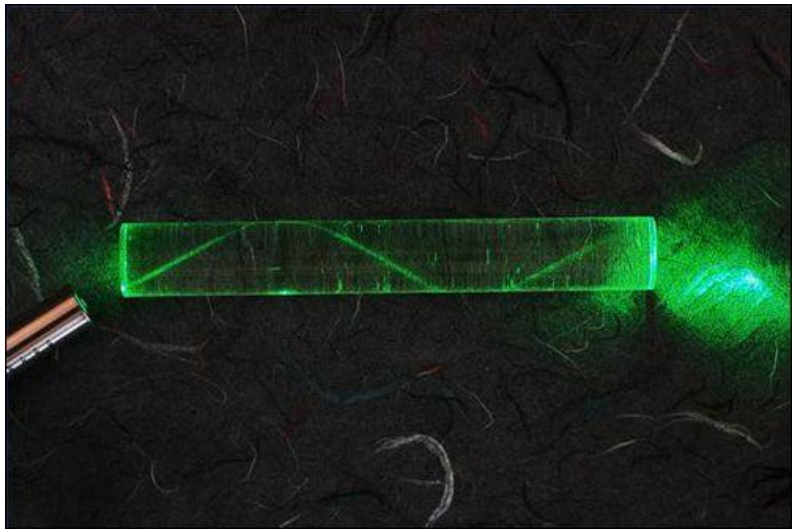

Wireless Ring: any two racks (e.g., A & B) on the ring can communicate with multi-reflections

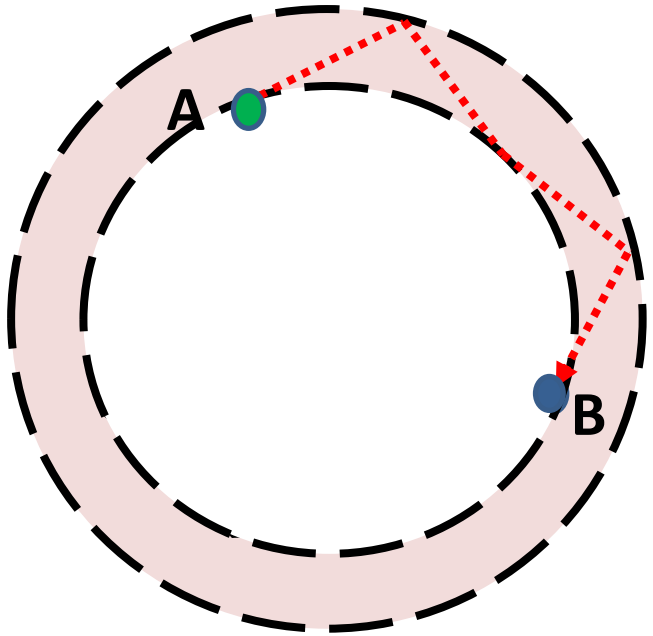

Solution 1—Multi-reflection Ring

- Scaling: add more *wireless rings*!
- But *circular reflector* board? *Hard&costly* to produce in industry…

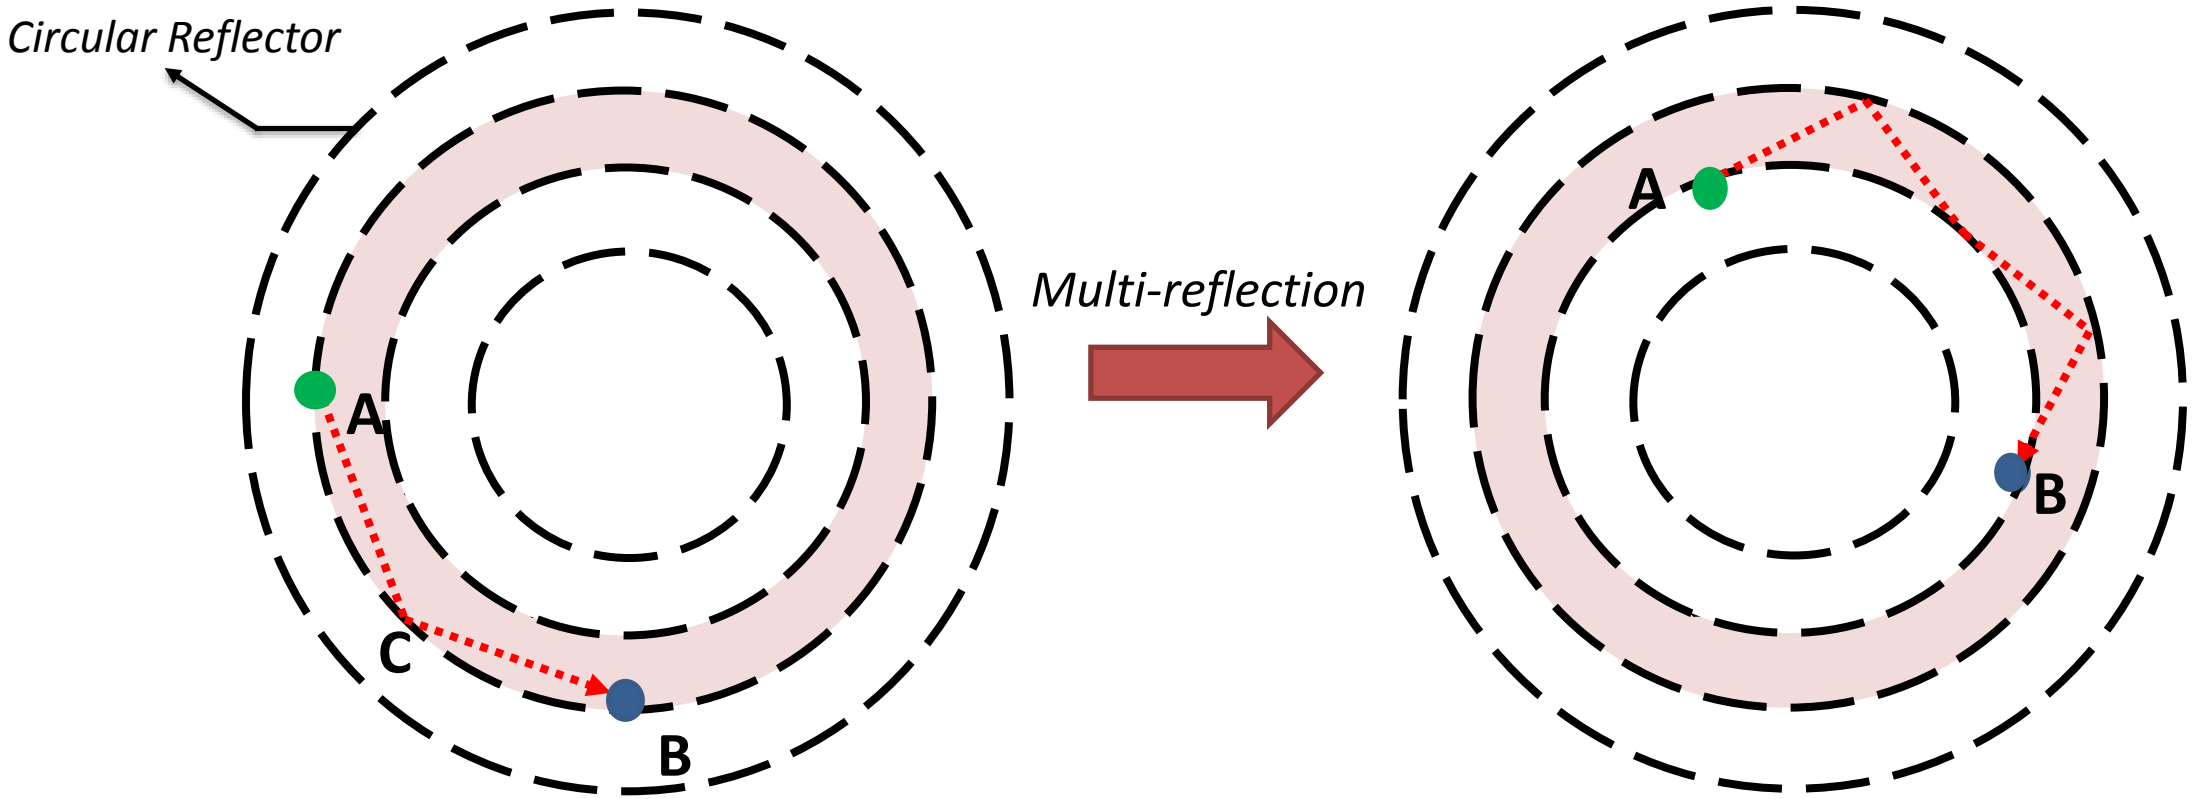

Solution 1—Multi-reflection Ring

- Using *equal-length flat reflection board* instead: easy&cheap for production
  - *Racks* are placed at the *vertex points* of regular polygon
  - *Reflection boards* are placed at the *edges* of regular polygon

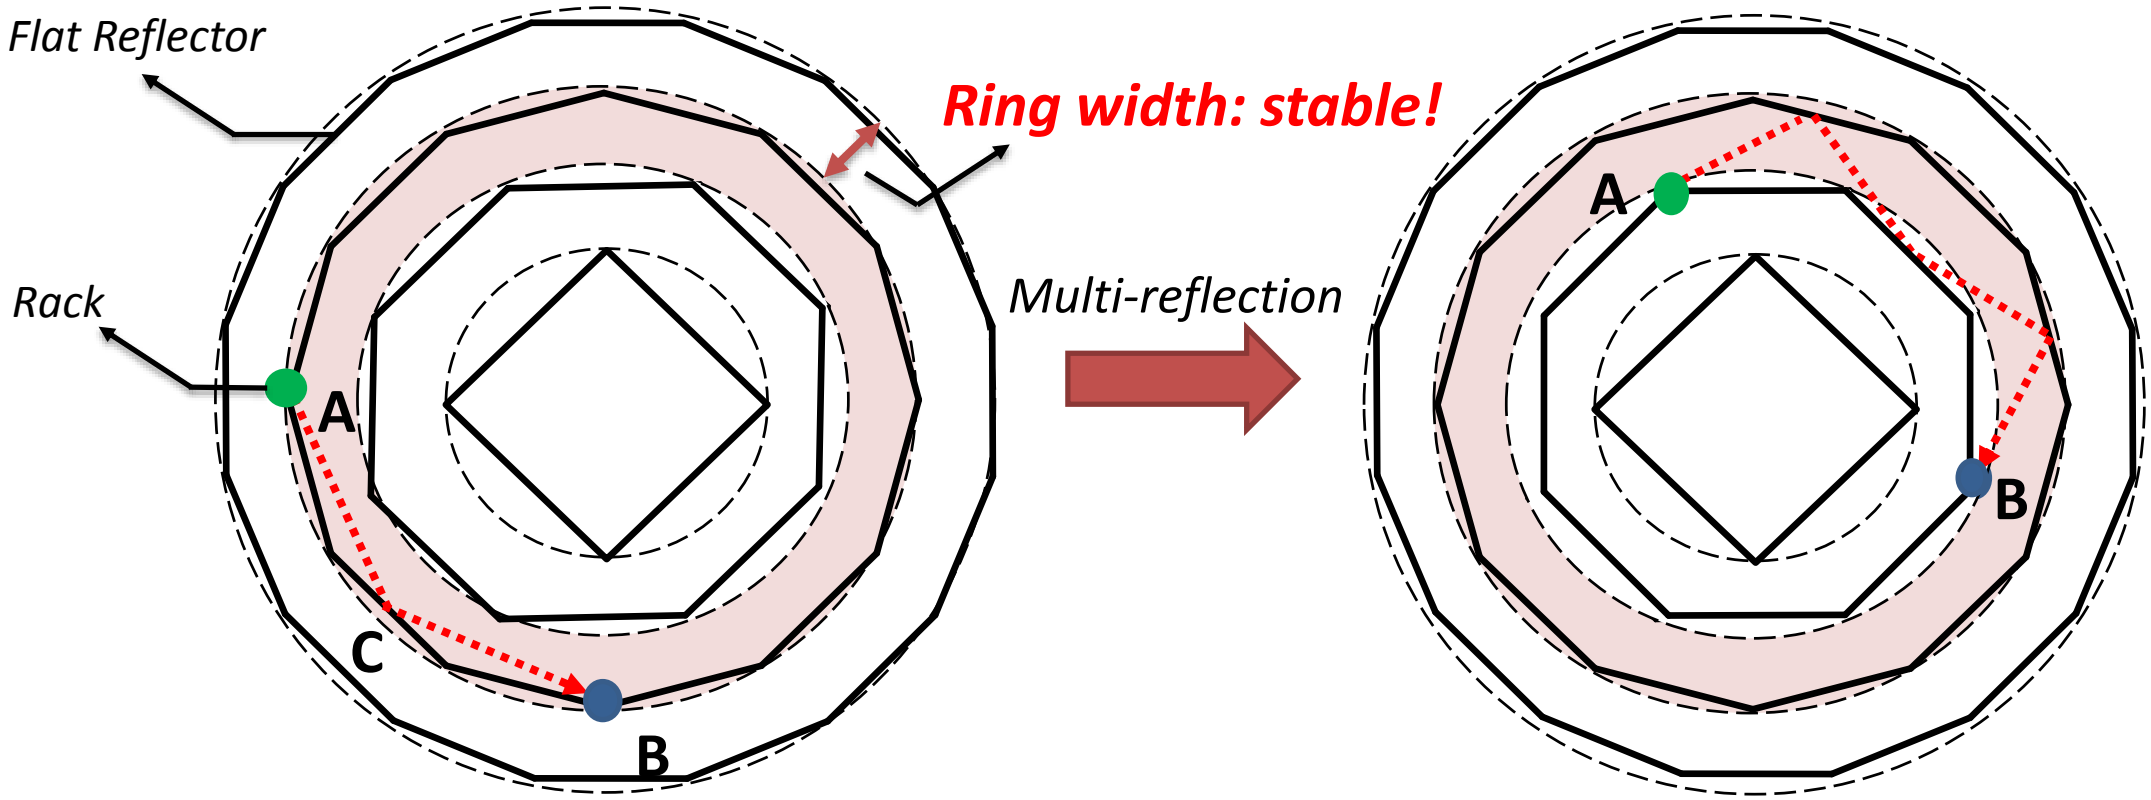

Solution 1—Multi-reflection Ring

- 3D Reflection in ring space: offering much higher flexibility

Deploy wireless radios on servers:
Enable a large number of direct server-to-server wireless links

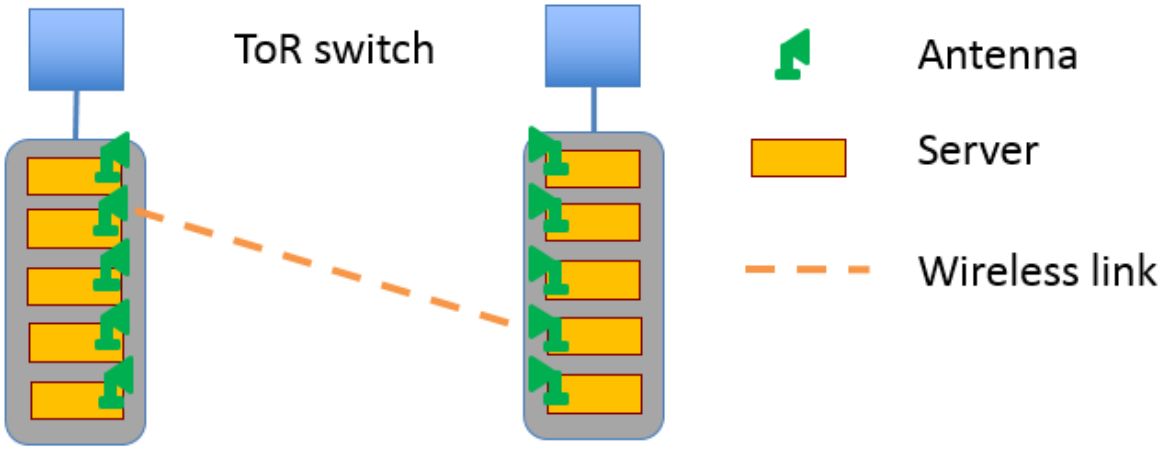

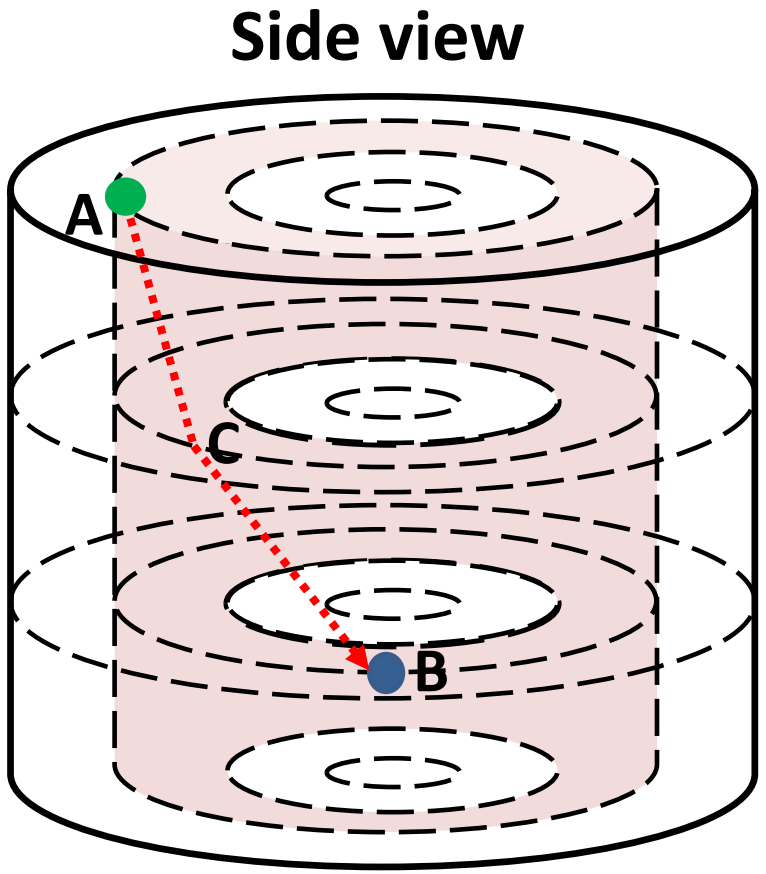

Challenge—Interference

- Directional wireless link (60GHz) is not “ideal thin line”: it has certain *beam width* and small *side-lobes* to create interference

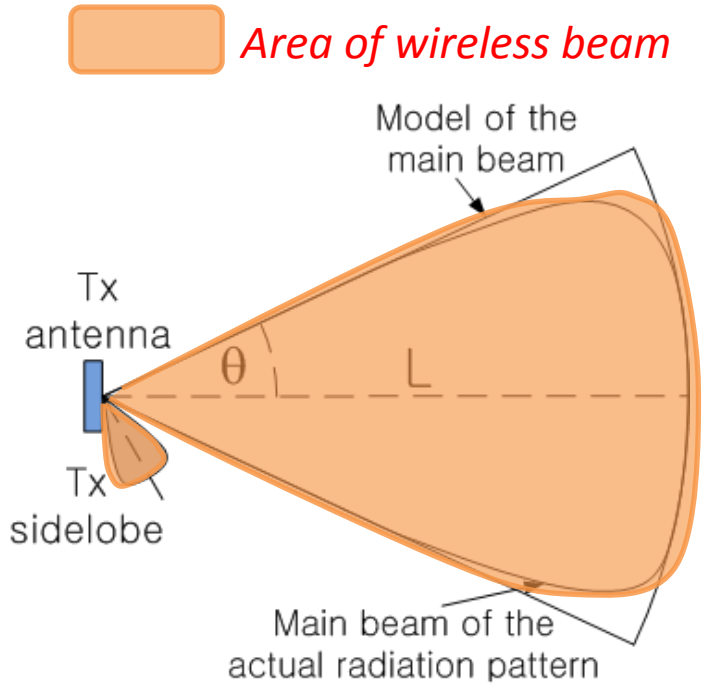

Figure source: Ji-Yong Shin, et al. ANCS’12

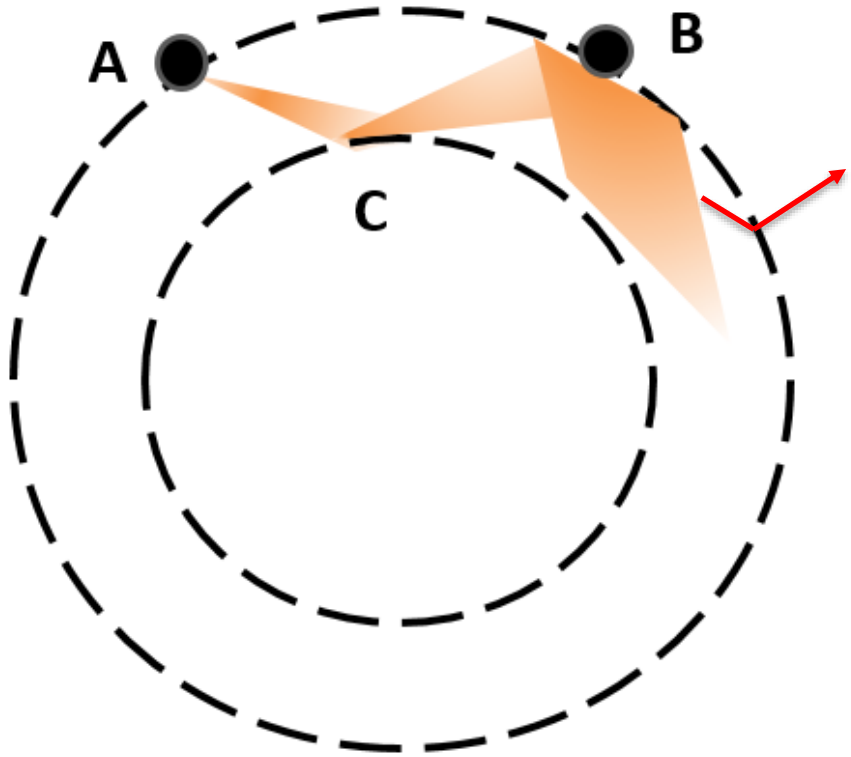

Undesired interference in the ring!

Solution 2—Precise reflection

- Filling the reflection board with absorbing paper, while only leaving special small holes for intended reflection points

Any 60Ghz wireless signal will be completely absorbed if it hits the “absorbing paper”

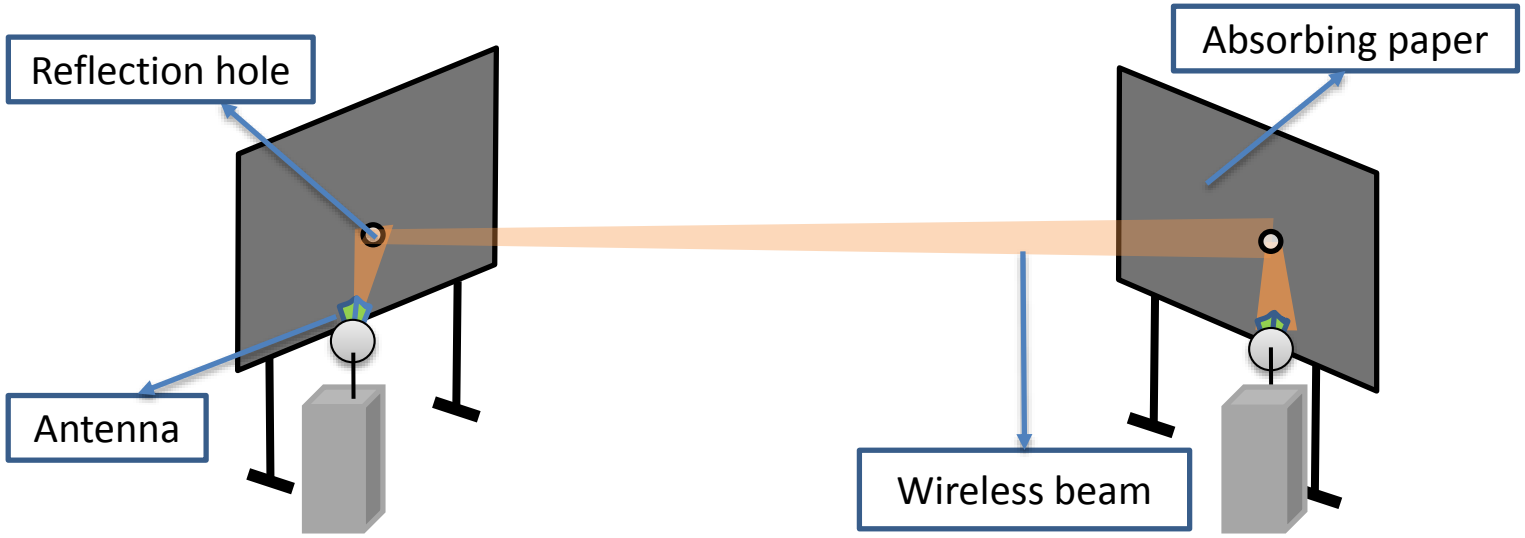

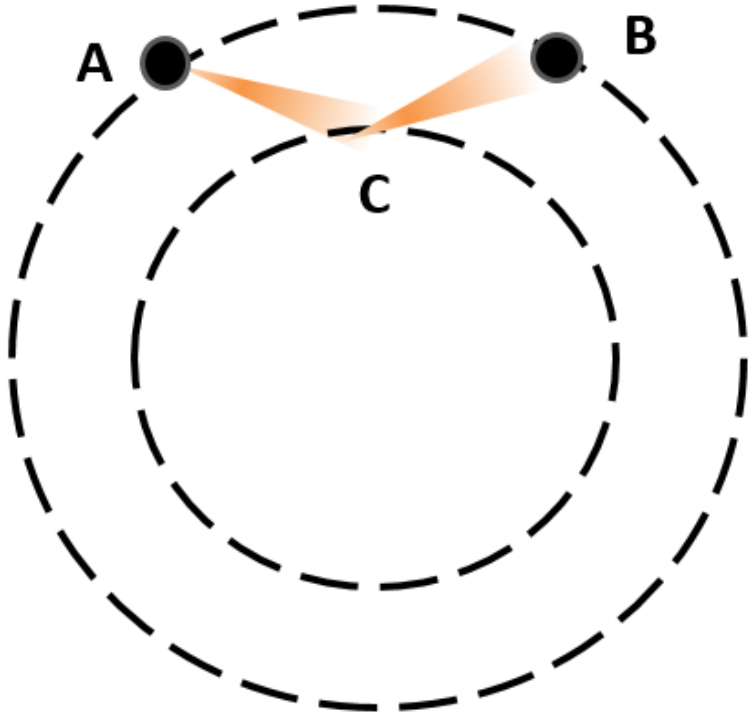

Solution 3—Cooperation with wires (Diamond)

- Function of wireless part: handling in-ring transmissions
- Function of wired part: handling cross-ring transmissions

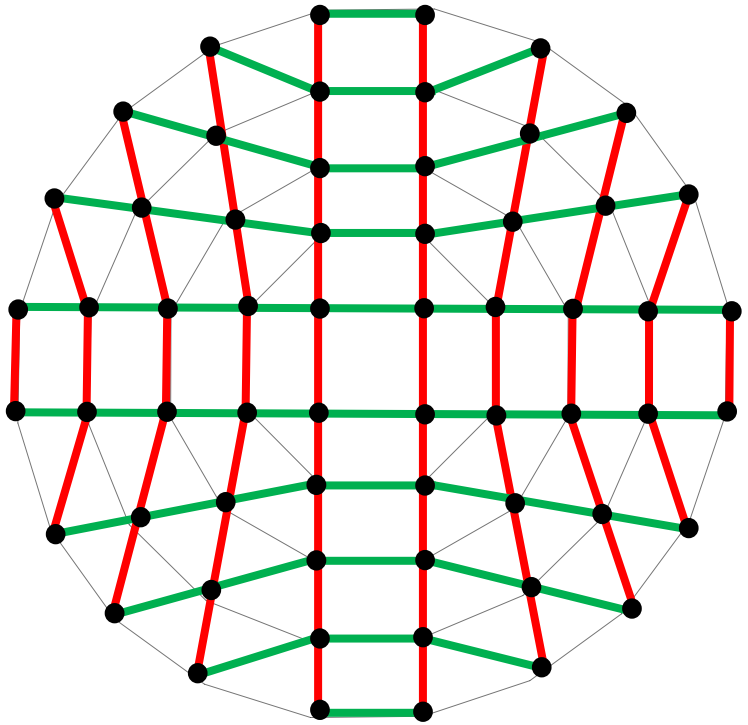

Overview of our Diamond architecture

A real diamond…

Solution 3—Cooperation with wires (Diamond)

- Function of wireless part: handling in-ring transmissions
- Function of wired part: handling cross-ring transmissions

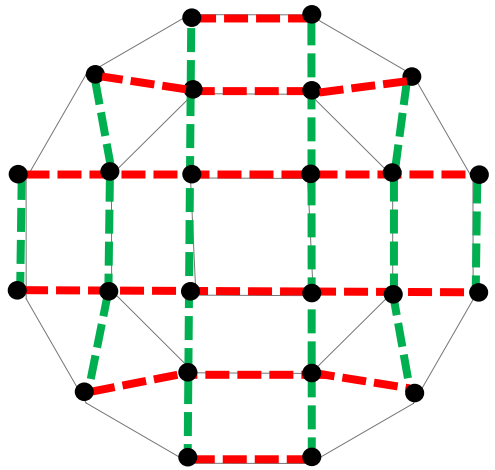

(a) Logical view

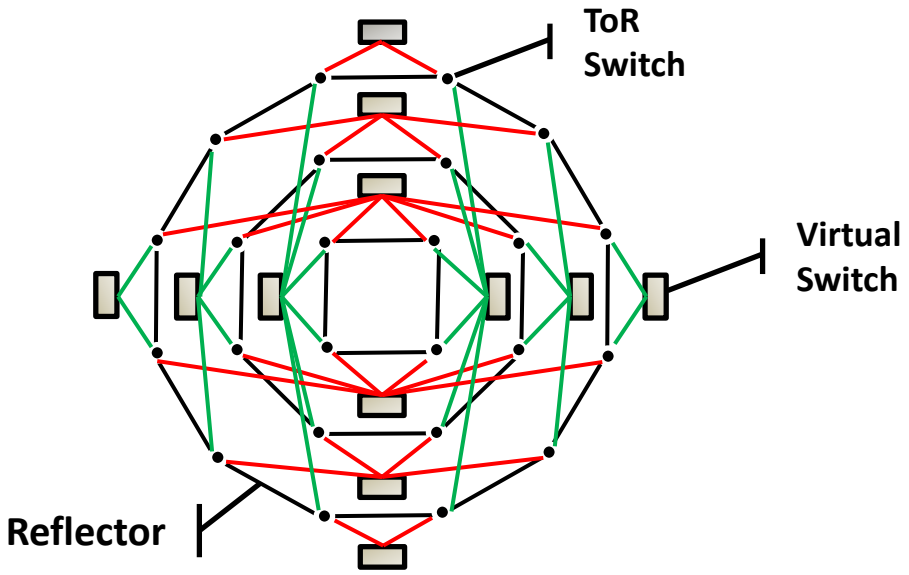

(b) Physical view

ToR Switch
Virtual Switch
Row
Virtual Link (Row)
Column
Virtual Link (Column)
Reflector

Solution 3—Cooperation with wires (Diamond)

- Design of virtual switch: De-Bruijn graph
  - Without additional switches
  - Well-defined recursive routing structures
  - Logarithmic network diameter
- Design of routing
  - Hotspot traffic: designated centralized routing (centralized scheduled by controller)
  - Non-hotspot traffic: real-time hybrid routing (distributed scheduled by server)

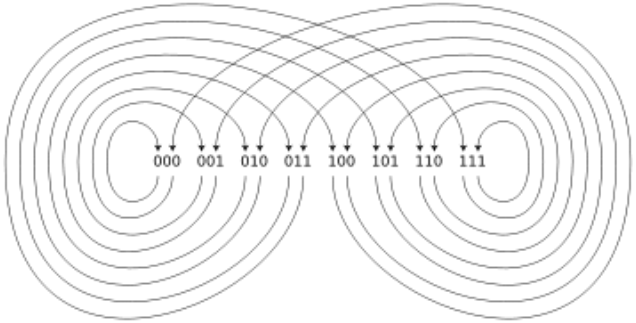

Figure source: Wikipedia "De-Bruijn graph"

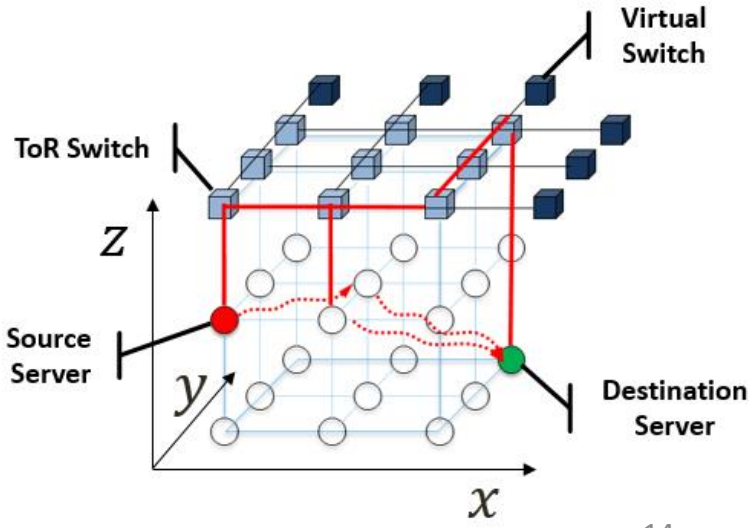

Testbed

- Single & Double reflection tests

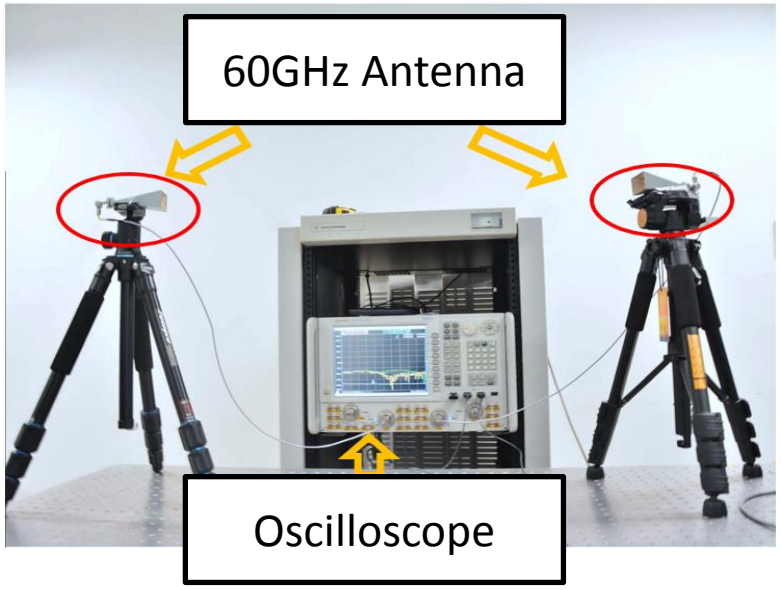

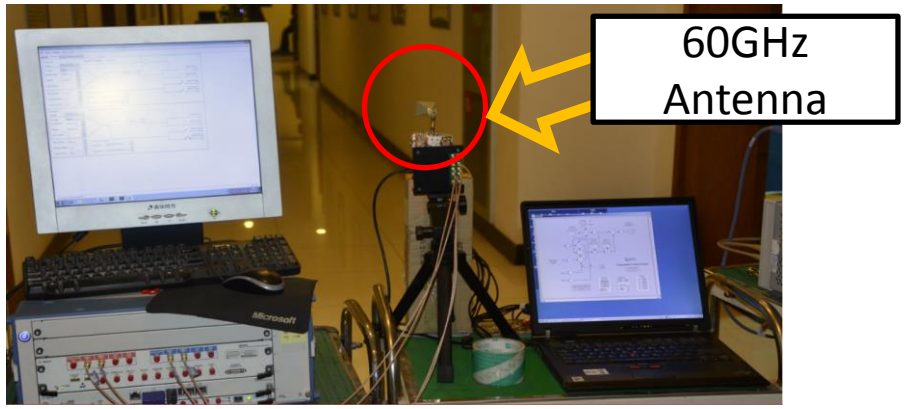

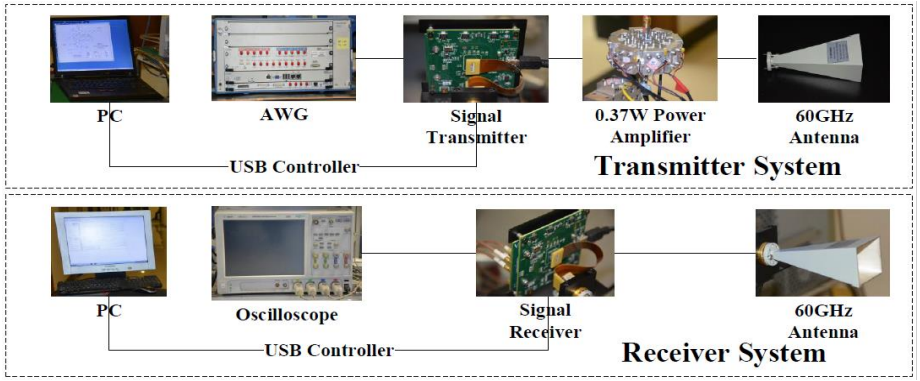

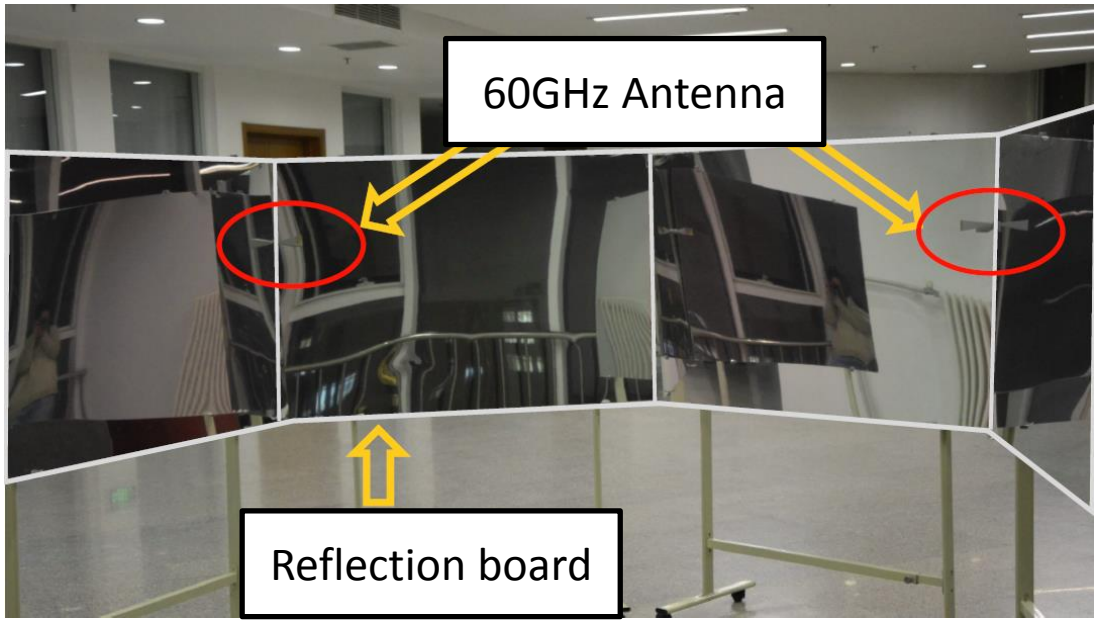

Experiment result

- Misalignment
  - Potential beam width is about 20° : a certain degree of *fault tolerance* on antenna misalignment
- Reflection hole
  - *Proper hole size* (diameter): 10cm
  - *Hole reusing*: above 50% reflection holes can be reused for different wireless links (*symmetrical structure*)
- Multi-reflection
  - *Little energy loss* when using flat metal board
  - *Little energy loss* when using *10cm reflection holes* on the flat metal board *filling with absorbing paper*

Hole size

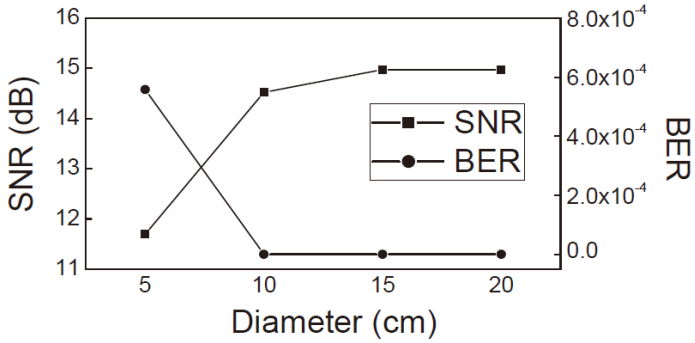

Reuse ratio

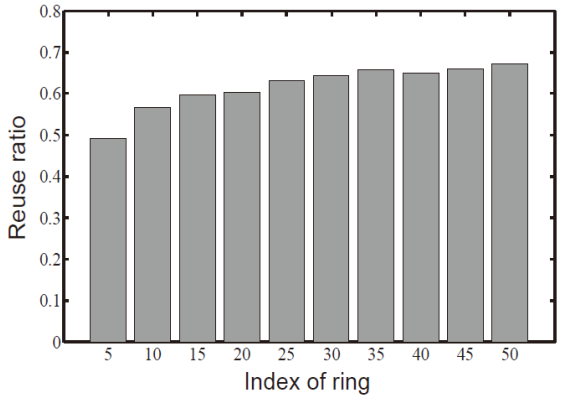

Simulation result

- Cover range
  - Cover 90% of ring within 3 reflections when ring number <10
  - Roughly, 1000 servers have potential *0.1 million* wireless links *within 2 reflections*
- Different traffic patterns
  - Average *5 times higher* throughout than others
  - Average *70% less* flow completion time than others
- Scheduling delay
  - Greedy runs each schedule *within 100ms*, while Optimal runs with exponential time of the problem scale
- Architecture cost
  - Diamond's cost is highest (*comparable to Firefly*), while it *trades off a larger number of wireless links* than others

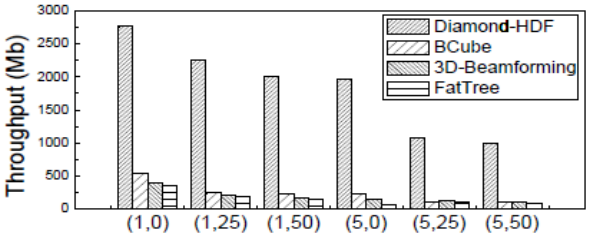

(a) Flow Throughput

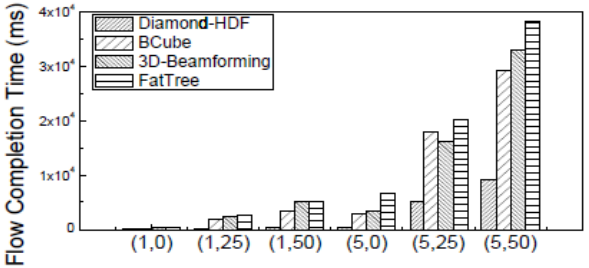

(b) Flow Completion Time

Topology #	Cost (k$)					Power (kw)
	NIC	Switch	Radio	Wire	Total	
FatTree	80	2080	-	80	2240	3486
3DB	80	2080	192	80	2432	3486
FireFly	80	416	2400	16	2912	4281
Diamond	240	832	1920	32	3024	3428

Conclusion

- Diamond can bring significant performance benefits for topology-reconfigurable DCNs
  - No need of the restricted-height clear ceiling space/ceiling mirror
  - Enable a large number of highly-flexible server-level wireless links
  - Better cooperation between wireless and wired transmission components
- Future vision: running FSO (Free-Space-Optics) in Diamond
  - Potential Tbps bandwidth
  - Nearly zero beam width: little interference
- Try it out for fun:

@ http://www.4over6.edu.cn/cuiyong/app/diamond.apk

Thank You !