Service Differentiation in the Internet for All

Sergey Gorinsky

IMDEA Networks Institute, Spain

Space Internetworking, Xanthi, Greece

10 June 2014

Delay

- Causes of delay
  - Propagation time in space
  - Limited available communication capacity
  - Storage in intermediate nodes (also a means to reduce delay)
- Application-specific requirements for delay
  - 100 ms of round-trip delay for a human conversation
  - 1 night to backup a dataset
- Internet handling of delay
  - Single best-effort service
  - Alternatives with service differentiation

Past visions of service differentiation

- IntServ (Integrated Services)
  - Performance guarantees for end-to-end flows
    - Delay and throughput
  - High complexity
- DiffServ (Differentiated Services)
  - Reduced complexity
  - Performance guarantees for classes of flow

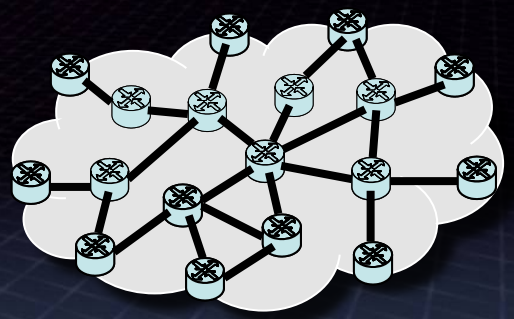

No end-to-end deployment

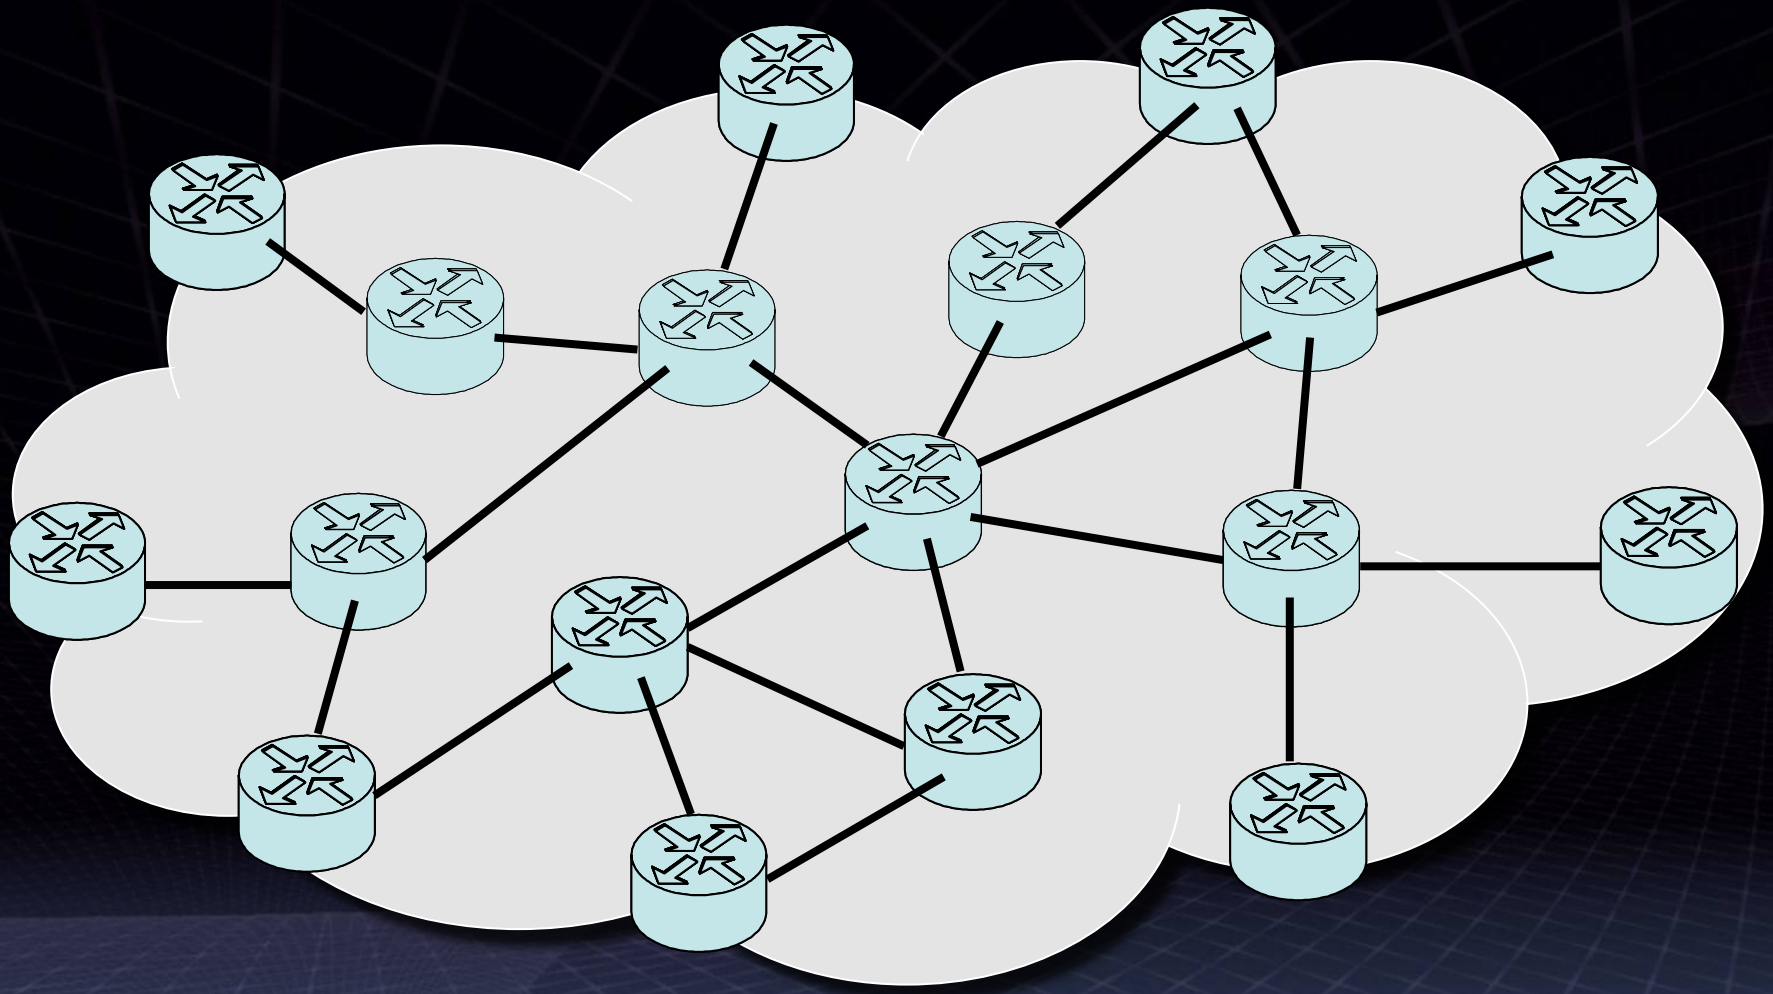

This is not an Internet. (© Magritte)

Internet of multiple stakeholders

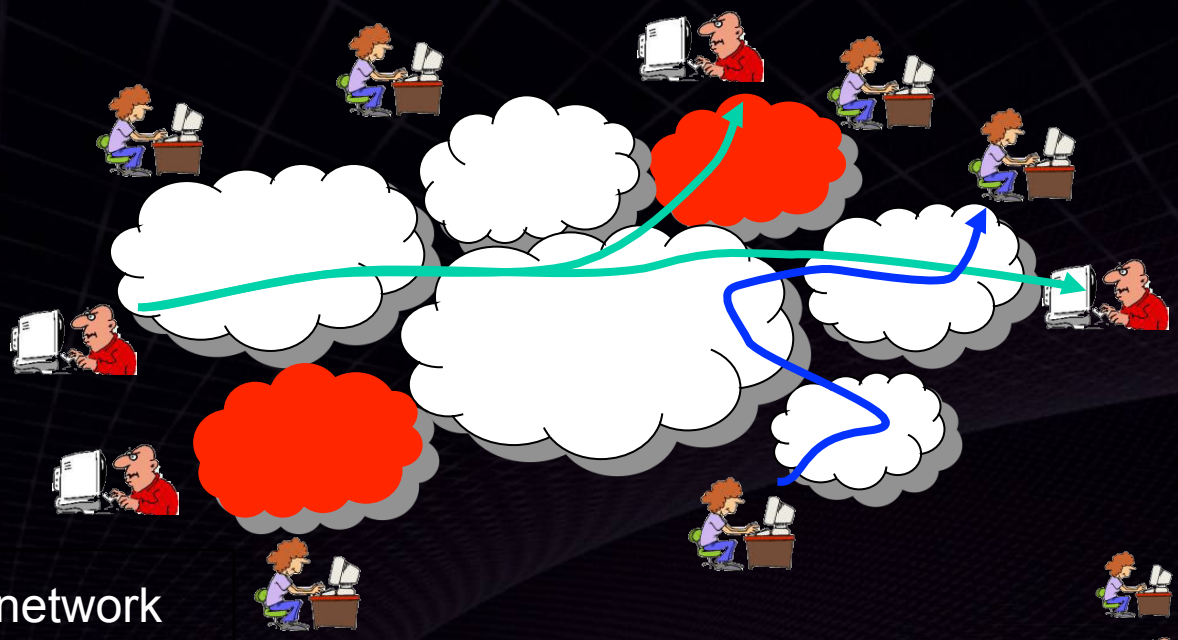

legacy network
adopting network

legacy user
adopting user

- Networks with different economic interests
- Ineffective service differentiation under partial deployment
- Differentiated charging for non-differentiated delay?

RD (Rate-Delay) network services

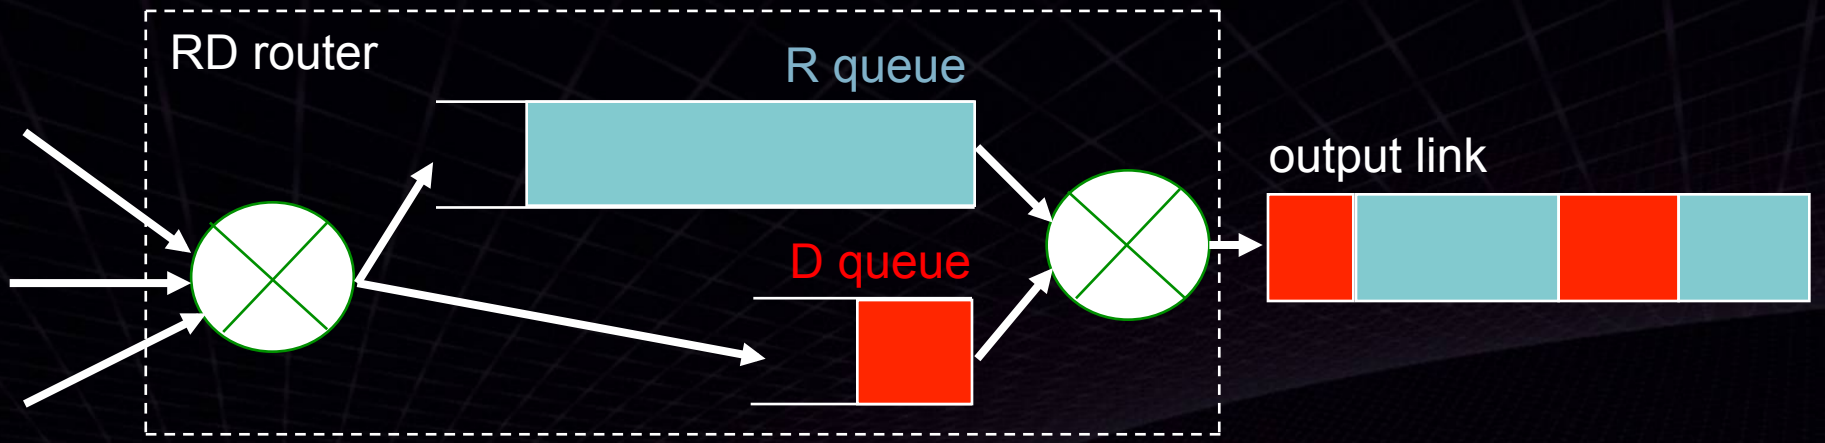

- Differentiation with free choice between two services
- R (Rate) service of higher throughput
- D (Delay) service of lower delay
- Router implementation via link scheduling and buffer sizing

Effective service differentiation without different prices

Current economy of Internet connectivity

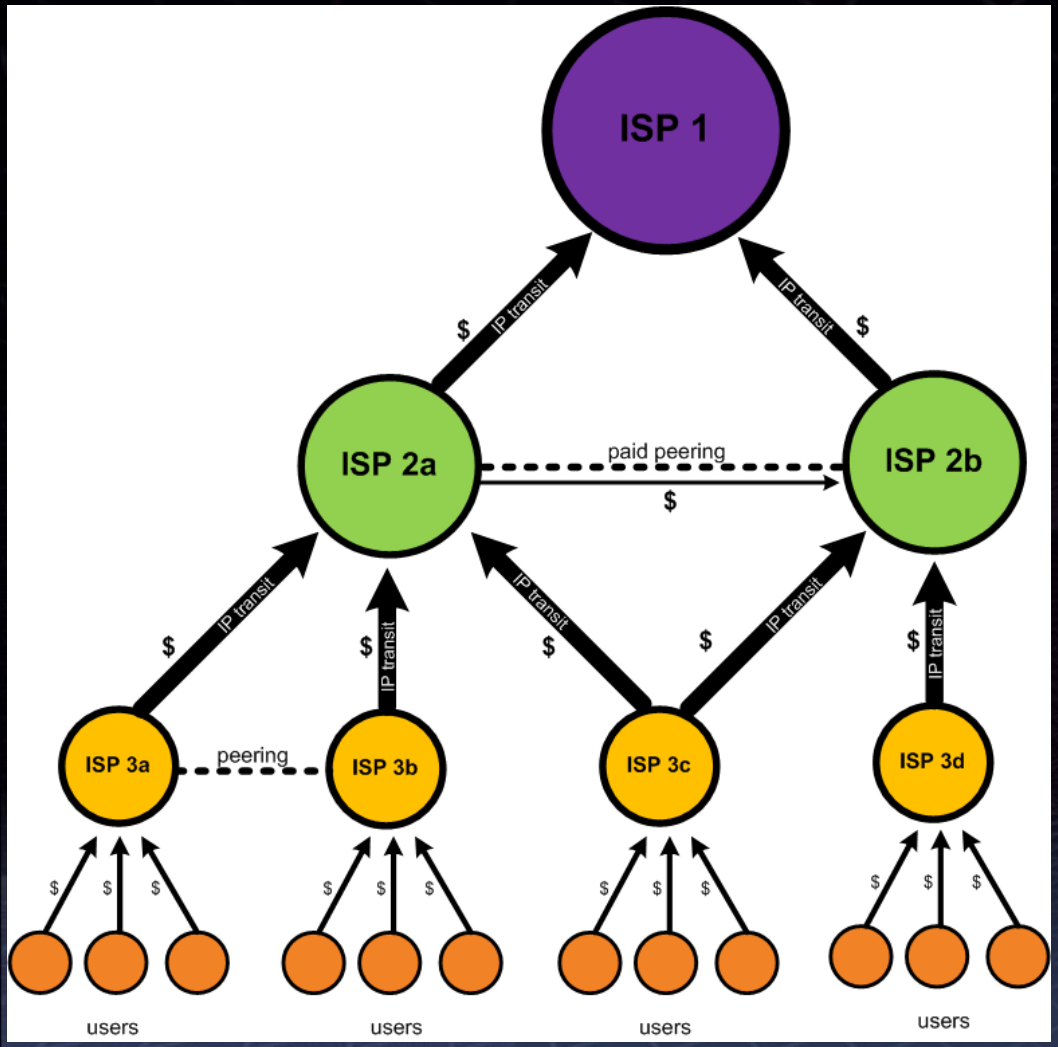

- Single best-effort service
- Delay reduction
  - By higher capacity
  - Due to flatter structure
- Diversification
  - Specialized ISPs (Internet Service Providers): access, content, transit
  - Interconnections: transit, peering, paid peering

Transit

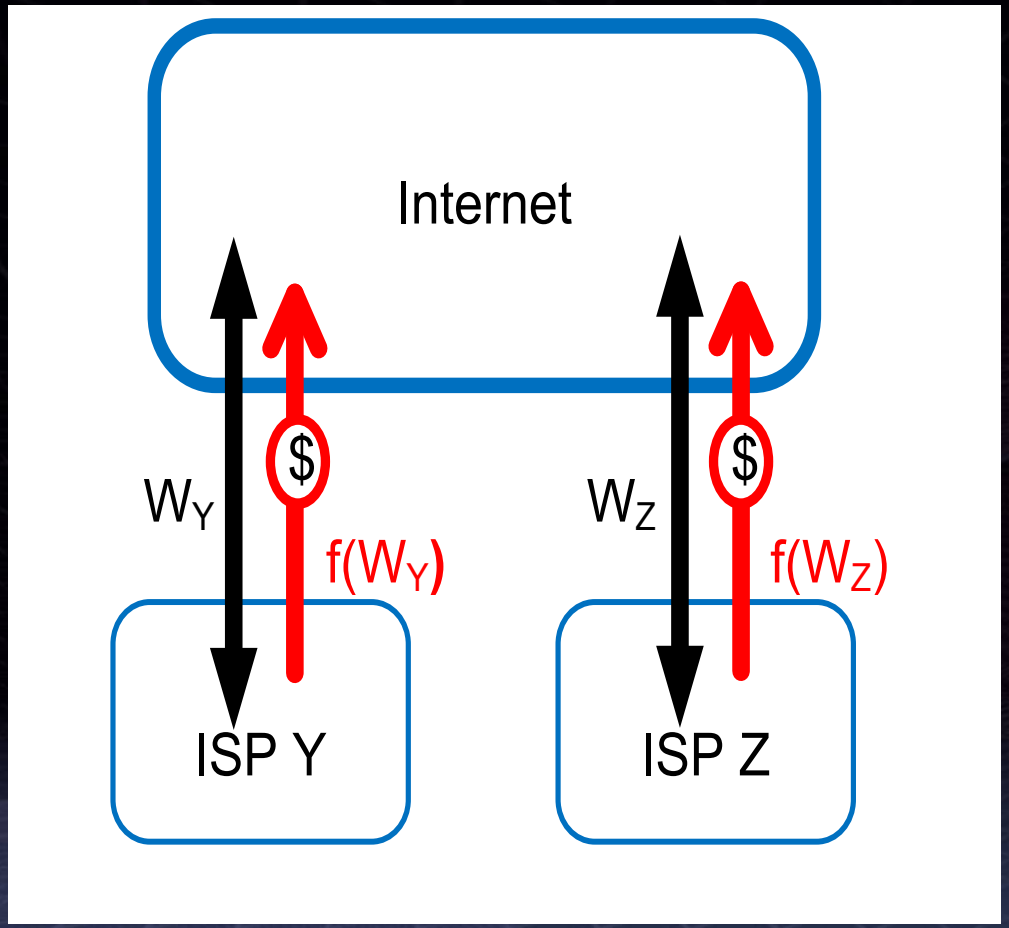

- The customer pays the provider for reaching the global Internet
- Billing depends on bidirectional traffic
- f() is a pricing function

Settlement-free peering

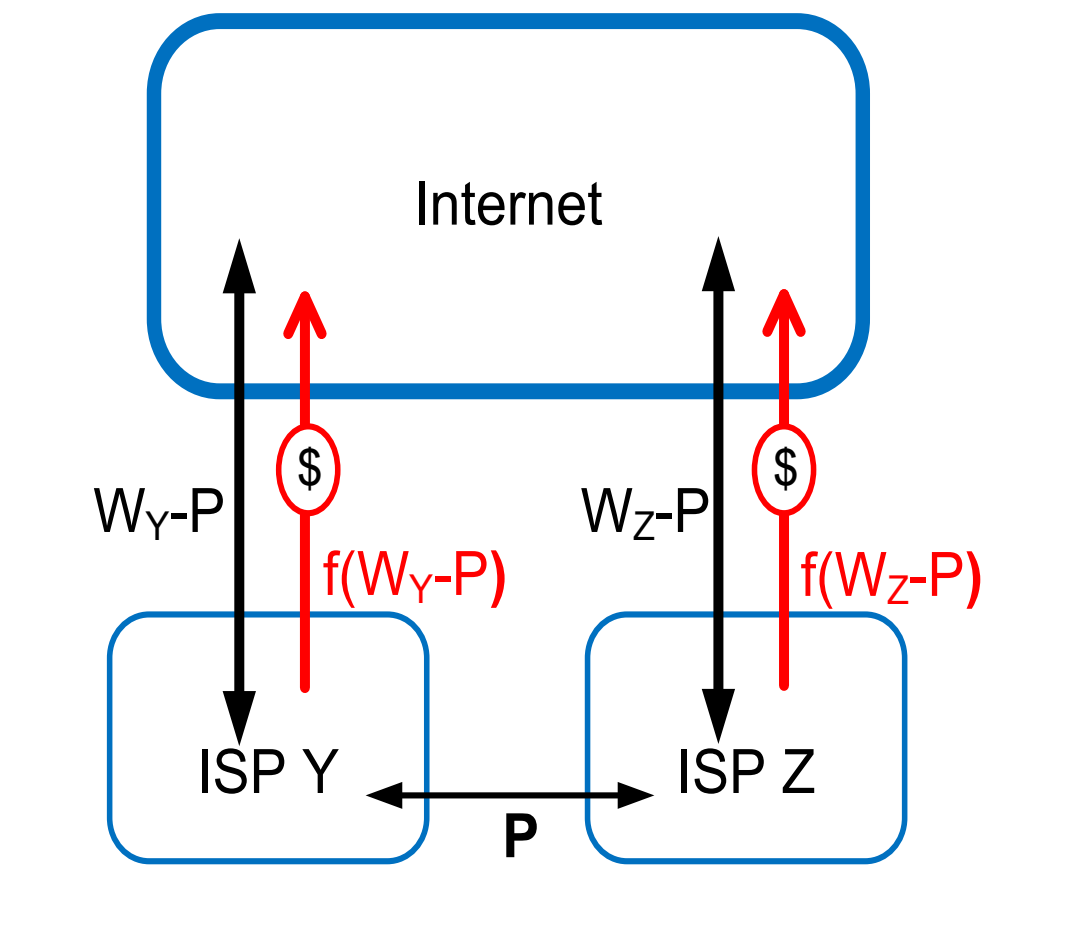

- Bilateral exchanges of own customer traffic P
- Reciprocity as the only compensation for the service
- Smaller transit traffic for ISPs Y and Z

Paid peering

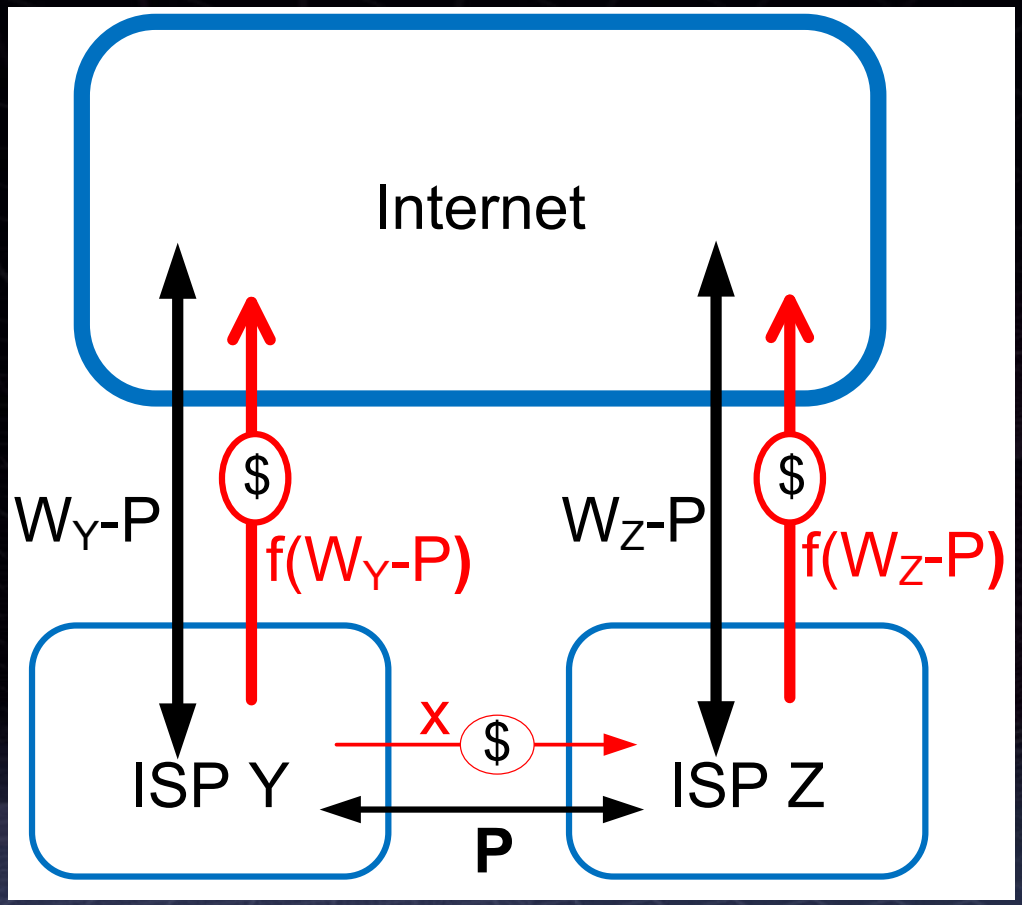

- Bilateral exchanges of own customer traffic P
- Compensation by reciprocity **and monetary payments**
- Imbalance of costs and market power (e.g., access vs. content ISPs)
- Pricing mostly unknown

T4P (Transit for Peering)

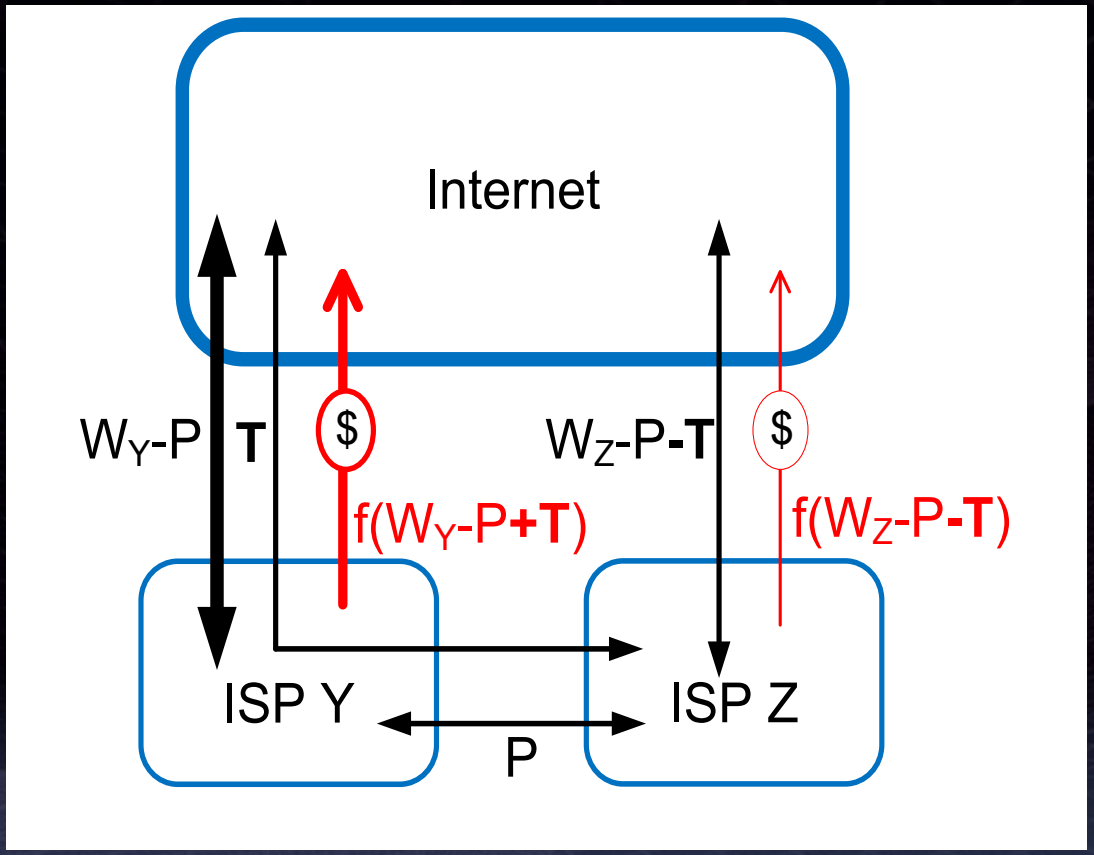

- Bilateral exchanges of own customer traffic P
- Compensation by reciprocity **and provision of transit services** for some traffic T
- Reduction in the combined transit costs of ISPs Y and Z

Interconnection innovations reduce costs

Concluding thoughts

- Delay has value and cost
  - Protocol costs of service differentiation are low
- Delay as a basis for economic and political tussles
  - Network neutrality vs. paid prioritization
- Uncertainty about the future: Internet vs. Internets
  - Economies of scale, network effect, …
  - "National" Internets, application-specific Internets, …
  - Economic efficiency vs. security vs. human rights