

Evolution of the Internet Ecosystem

Dmitri Krioukov

Marina Fomenkov, Ryan Koga, kc claffy

CAIDA/UCSD

Srinivas Shakkottai

Texas A&M University

Motivation

- # Many Internet evolution models exist. Why another one?
 - # There is none which would be simultaneously
 - realistic
 - parsimonious
 - having all its parameters measurable
 - analytically tractable
 - “closing the loop”
 - # Only a model satisfying all these requirements can shed some light on how the Internet really evolves
-

Multiclass preferential attachment: PA + Internet-specific modification

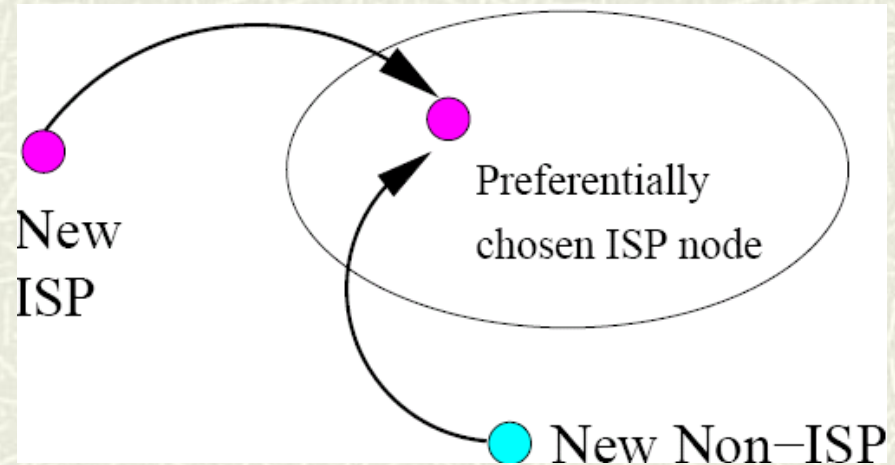
- # All ASs can roughly be split into two classes: ISPs and non-ISPs
 - # New ASs can preferentially attach to ISPs, but they cannot connect to non-ISPs at all, as those do not provide Internet connectivity services
 - # A majority of ASs (~70%) are non-ISP
-

The two key observations

- # This simple modification of PA captures a bulk of the Internet topology properties
 - # All other improvements and modifications (such as peering, bankruptcy, multihoming, geography, etc.) lead to much finer corrections
-

ISPs vs. non-ISPs

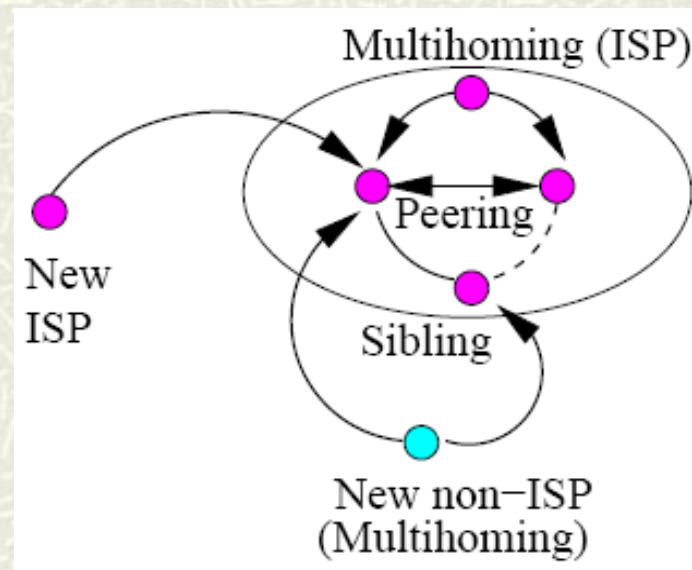
- Time unit: 1 new ISP
- Non-ISPs per time unit: ρ . The measured value of ρ is $\rho = 7/3$
- Analytic solution for the degree distribution yields $P(k) \sim k^{-2.3}$
- In the real Internet $P(k) \sim k^{-2.1}$



$$P(k) \sim k^{-\left(2 + \frac{1}{1+\rho}\right)}$$

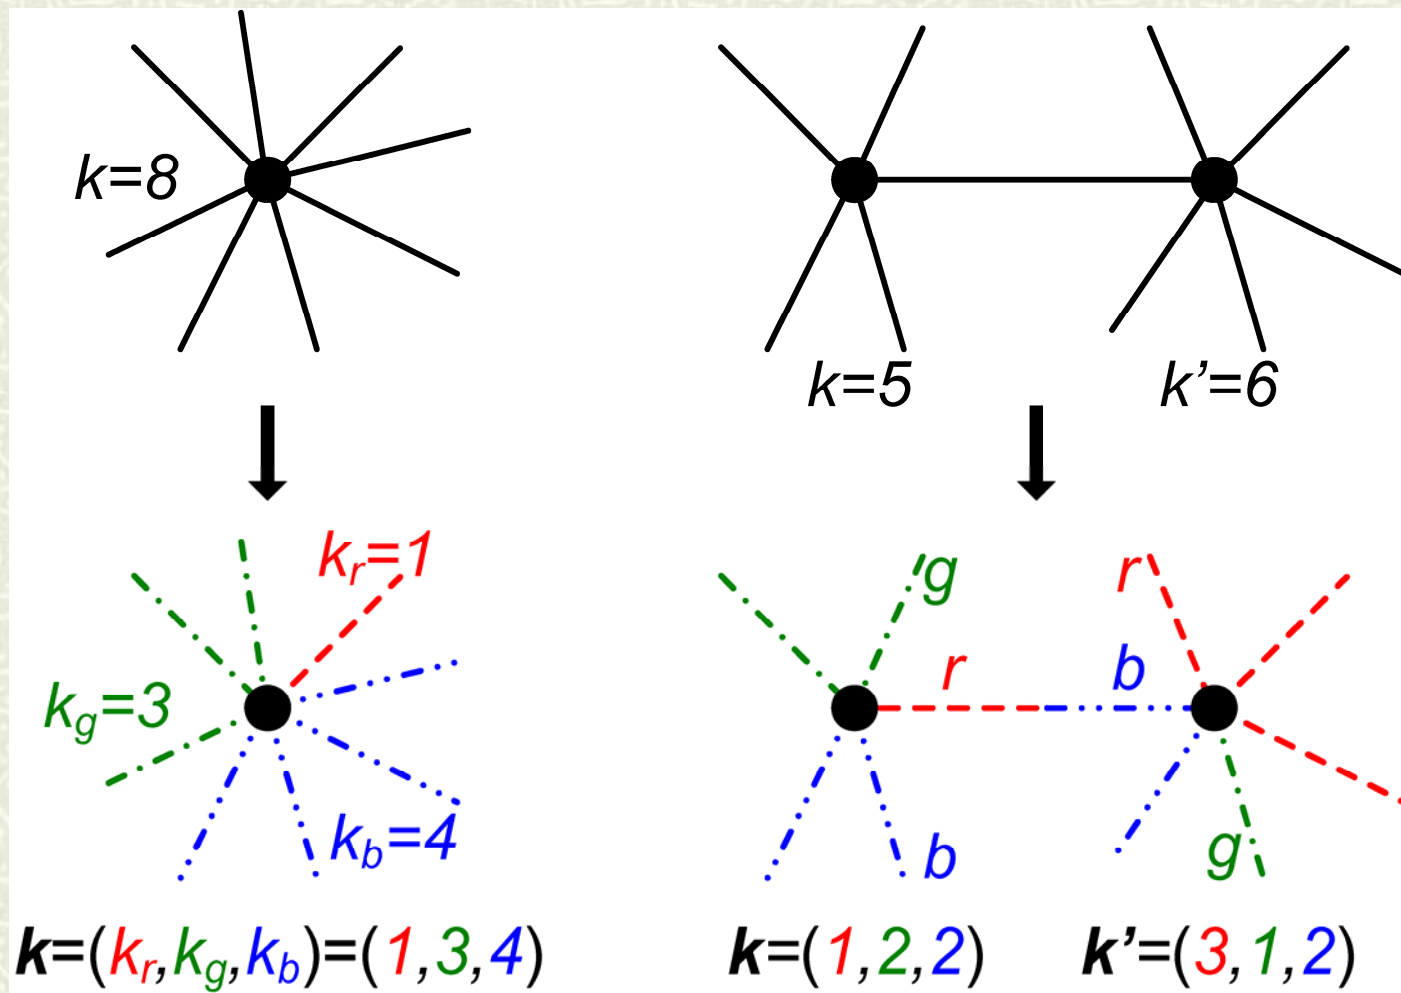
Finer adjustments

- ▣ Peering: peering links per time unit
 $c = 0.70$
- ▣ Sibling links:
 $\mu \approx 0$
- ▣ Multihoming:
 - ISP's average number of providers
 $\nu \approx 2$
 - non-ISP's average number of providers
 $m = 1.86$
- ▣ Analytic solution
 $\gamma = 2.1$
- ▣ Real Internet
 $\gamma = 2.1$



$$\gamma = 2 + \frac{1 - \mu}{1 + 2\nu + m\rho + 2c + \mu}$$

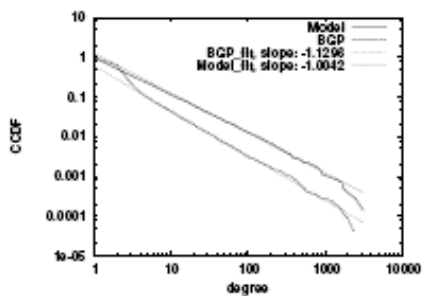
Annotated graphs



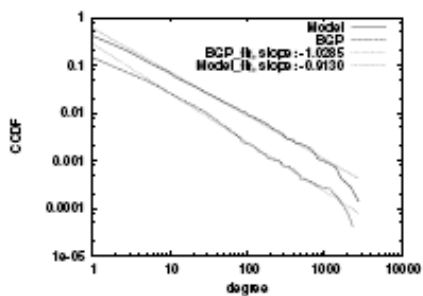
Model validation

- # Reproducing the joint degree distribution (JDD) of the AS Internet annotated with AS business relationships captures all its other properties in synthetically generated networks
 - # Simulate the model with all its parameters equal to their measured value and compare the JDDs in the modeled networks and the Internet
-

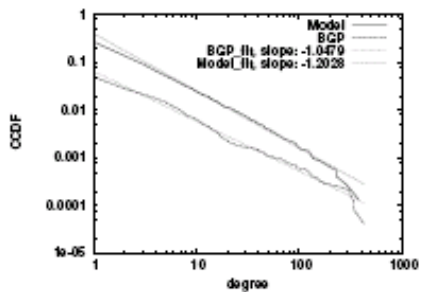
Validation results



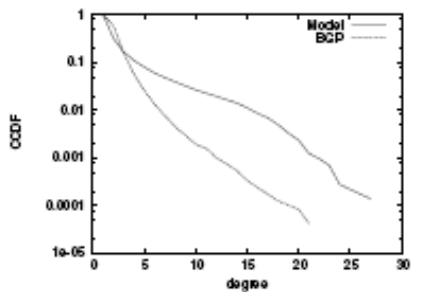
(a) Degree distribution of ASs (CCDF)



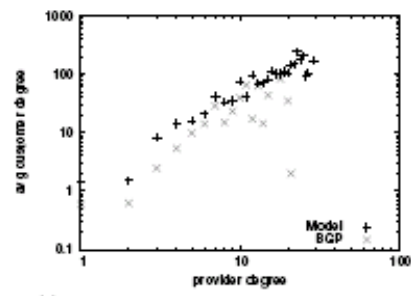
(b) Empirical customer distribution (CCDF)



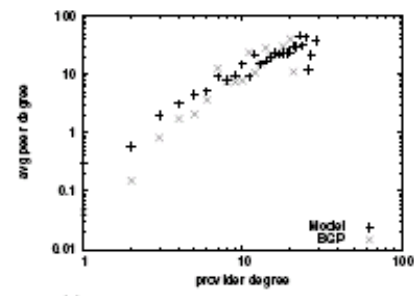
(c) Empirical peer distribution (CCDF)



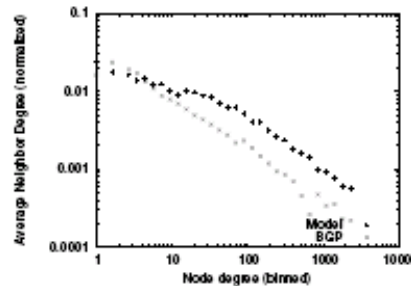
(d) Empirical provider distribution (CCDF)



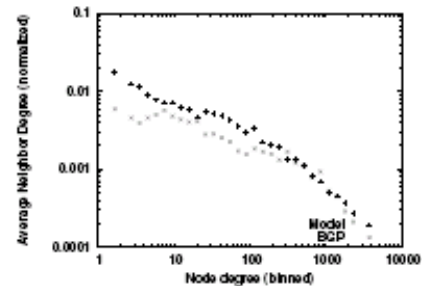
(e) ISP's provider vs. customer degrees



(f) ISP's provider vs. peer degrees



(g) ISP's degree vs. its customers' degree



(h) ISP's degree vs. its peers' degree

Conclusion

- # The Internet appears to evolve according to preferential attachment
 - Preferential attachment, with minor Internet-specific corrections, suffices to explain virtually all properties of the Internet AS-level topology and its evolution
 - # Most links are from customer to provider ASs
 - # Therefore to make a step forward and connect our model to “real economics,” one needs to explain how customers select their providers
 - # Popularity of providers, their “brand names,” may be a real explanation of preferential attachment in the Internet
-