Optimization-Based Approaches to Realistic Internet Topology Generation

David Alderson
with John Doyle, Walter Willinger,
Hongsuda Tangmunarunkit

June 11, 2002

Background

Modeling realistic Internet topology is important for several reasons...

- Design and evaluation of protocols
- Understanding current behavior
- Potential insights for future provisioning

Modeling efforts are relevant at multiple levels of network hierarchy (router vs. AS graph)

Several generation schemes are designed to match specific topological properties of interest.

- graph theoretic measures (connectivity, diameter, etc.)

- hierarchy

- power laws

One criticism of this approach has been that the generation mechanisms are disconnected from the experience of network engineers building real networks.

Our Approach
(what we believe)

- When building data networks, operators are solving an optimization problem (either explicit or implicit)

- To the extent that we can identify the objectives and constraints of this optimization, we can solve the optimization problem and (hopefully) generate realistic topologies

- There is much to be learned (in both directions) from a dialogue between theory and practice

Our Interest

- Identification of realistic optimization formulations for network design

- How various formulations affect resulting topologies

The ultimate objective of work in this direction would be to develop a topology generation scheme that is *realistic* in the following senses.

- The mechanisms that generate the topology are driven by the some of the same forces, such as cost (profit?) and traffic demand, that drive real network deployment.

- The *incremental decisions* made by the generation scheme are not incompatible with those of a network engineer.

- The *resulting topology* "looks like" the structures that are observed within real networking environments.

Outline

- Basic optimization formulation for network design
- The network access design problem
- Initial results and related work
- A strawman model for Internet topology
- Next steps

Directed graph $G = (N, A)$

L commodities, to be routed by known demands

c_{ij}^l = *flow cost per commodity* l *on arc* (i, j)

K cables, to be installed on arcs

σ_{ij}^k = *installation cost per cable* k *on arc* (i, j)

u_{ij}^k = *capacity of cable* k *on arc* (i, j)

For each commodity $l \in L$

$O(l)$ = *origins, node* $i \in O(l)$ *supplies* $o_i^l > 0$

$D(l)$ = *destinations, node* $i \in D(l)$ *demands* $d_i^l > 0$

$T(l)$ = *transshipment nodes*

Decision variables

x_{ij}^l = *continuous flow variable for commodity* l *on* (i, j)

y_{ij}^k = *integer design variable for cable* k *on* (i, j)

Problem:

simultaneously install cables and route flows at minimum cost.

$$\min_{x,y} \quad \sum_{l \in L} \sum_{(i,j) \in A} c_{ij}^l x_{ij}^l + \sum_{k \in K} \sum_{(i,j) \in A} \sigma_{ij}^k y_{ij}^k$$

$$s.t. \quad \sum_{j \in N^+(i)} x_{ij}^l - \sum_{j \in N^-(i)} x_{ji}^l = \begin{cases} o_i^l & i \in O(l) \\ -d_j^l & i \in D(l) \\ 0 & i \in T(l) \end{cases} \qquad i \in N, l \in L$$

$$\sum_{l \in L} x_{ij}^l \leq \sum_{k \in K} u_{ij}^k y_{ij}^k \qquad (i,j) \in A$$

$$x_{ij}^l \geq 0 \qquad (i,j) \in A, l \in L$$

$$y_{ij}^k \ \text{integer} \qquad (i,j) \in A, k \in K$$

This problem is notoriously hard because of the *capacity constraints*.

Consider the special case of minimum spanning tree problem. It is well-known that the uncapacitated version of this problem is solved easily using Greedy methods, while the capacitated version is NP-hard.

Solution methods for capacitated network design problems fall into a few groups:

1. *simplex-based cutting-plane methods*
2. *Lagrangian relaxation*
3. *heuristic approaches*

See the review by Gendron, Crainic, and Frangioni for details.

Caveats

Due to the difficult nature of these problems, they are generally solved approximately instead of exactly.

As a result, any topology that results from an optimization "solution" will be a function of

- *problem formulation (objective and constraints)*
- *problem data (parameter values)*
- *approximation technique used*

Previous Work

In many real-world topology design problems, the available cable types observe *economies of scale* in the cost per unit bandwidth. Optimization problems of this type are known as *Buy-at-Bulk* problems.

A special case of the capacitated network design problem occurs when the objective is to connect several source nodes to a single sink. Salman *et al* (1997) first proposed the buy-at-bulk cost structure for the single sink installation problem. This problem has been studied in the context of designing network access by Andrews and Zhang (1998) and Meyerson *et al* (2000, 2001).

Network Access Design (Andrews and Zhang)

K cable types:

u_k = *cable capacity*

σ_k = *installation cost*

δ_k = *per distance usage cost*, $\delta_k = \sigma_k / u_k$

economies of scale: $\sigma_1 < \sigma_2 < \cdots < \sigma_K$, $\delta_1 > \delta_2 > \cdots > \delta_K$

Total cost to send B traffic over a cable of type k installed on an arc of distance D

$$D \cdot (\sigma_k + B \cdot \delta_k)$$

Network Access Design (Andrews and Zhang)

Define α-*utilization* to mean the following:

$$d \cdot \delta_k < \alpha \cdot \sigma_k \ \rightarrow \ d \cdot \delta_{k-1} + \sigma_{k-1} < d \cdot \delta_k + \sigma_k$$

[Note: this "if" condition is really just $d < \alpha \cdot u_k$]

In words, having demand greater than $\alpha \cdot u_k$ means it is more economical to use cable type k over cable type $k-1$.

Assuming that all installed links were α-utilized, Andrews and Zhang showed that the optimal solution to the problem was a tree in which traffic from sources is routed through cables of increasing type. Their algorithm generates a tree that gives an $O(K^2)$ approximation, where K is the number of cable types.

Meyerson *et al*

Using the formulation of Andrews and Zhang and assuming identical and randomly placed demand nodes, Meyerson *et al* developed an incremental algorithm that provides a constant factor approximation to the optimal solution.

The algorithm (in its offline form) is extremely simple:

1. *Assign to each node cable type* k *with probability*

$$prob(k) = \min\{1, \gamma/u_k\}$$

2. *Keep the largest cable type assigned, discard all others*
3. *Connect node to the nearest node of greater cable type*

This algorithm can be easily extended to the case of multiple sinks and when nodes are connected incrementally.

Numerical Investigation

Use simulation to get a feel for the tradeoffs in parameters.

1. *Simulate for n=25,000 nodes for different values of* γ
2. *Look at node degree distribution (log(freq.) vs. degree)*

Cable	Capacity	Relative Cap.	u_k	σ_k	δ_k
DS0	64 kbps	1	1	1	1
DS1	1.544 Mbps	24.125	24	12	1/2
DS3	44.736 Mbps	672	600	150	1/4
OC3	155.520 Mbps	2430	2400	300	1/8
OC12	622.080 Mbps	9720	9600	600	1/16
OC48	2.488 Gbps	38880	38400	1200	1/32
OC192	9.953 Gbps	155520	153600	2400	1/64

Degree distribution for $\gamma = 1.0$

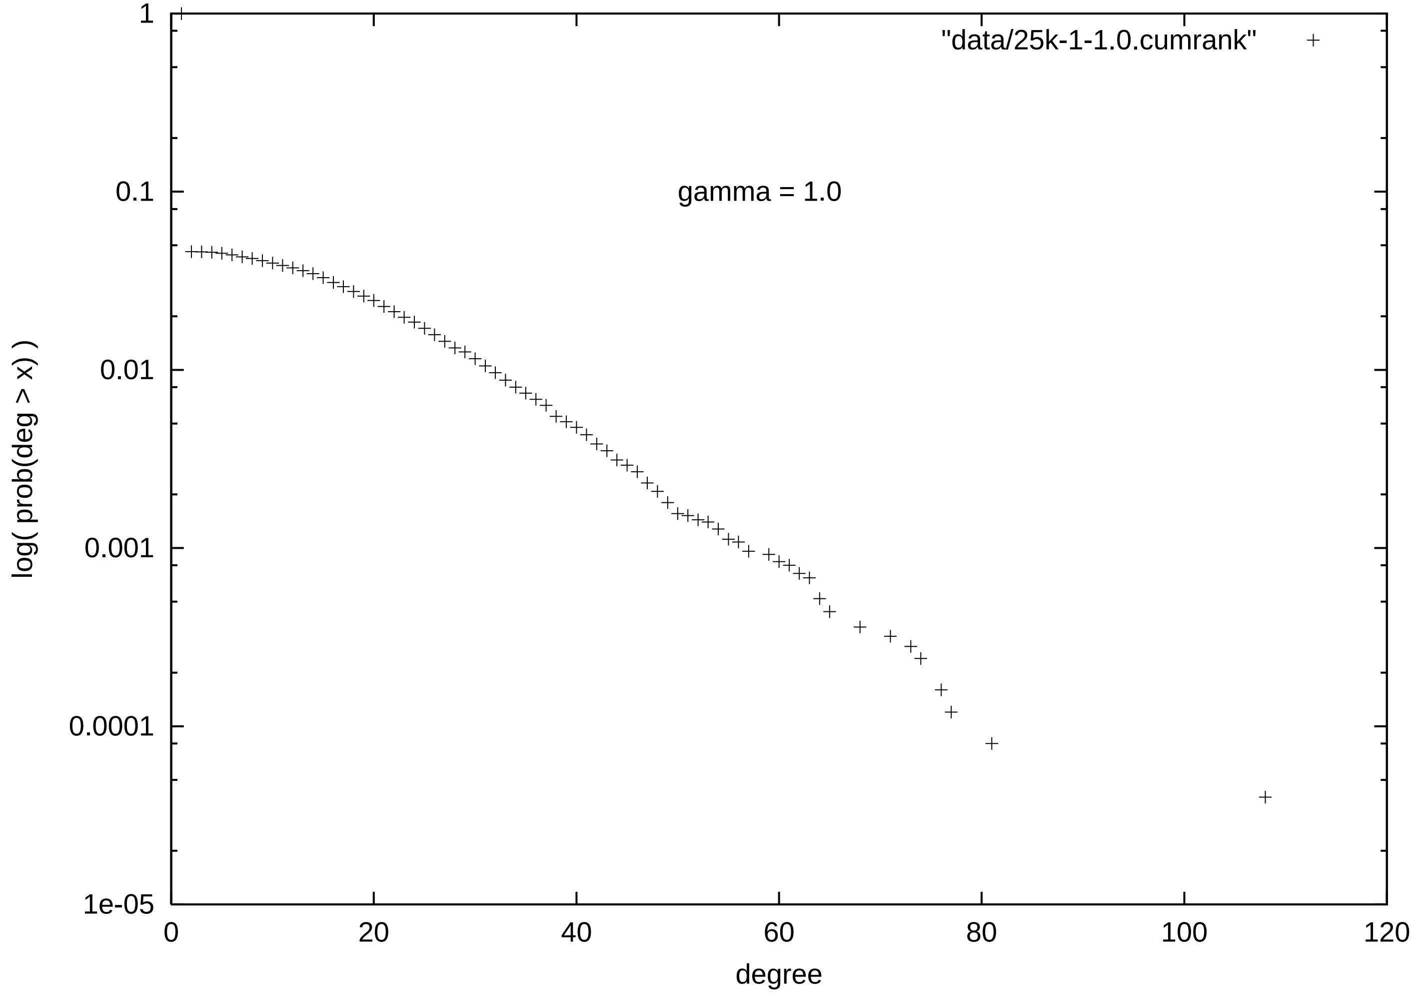

Degree distribution for $\gamma = 10.0$

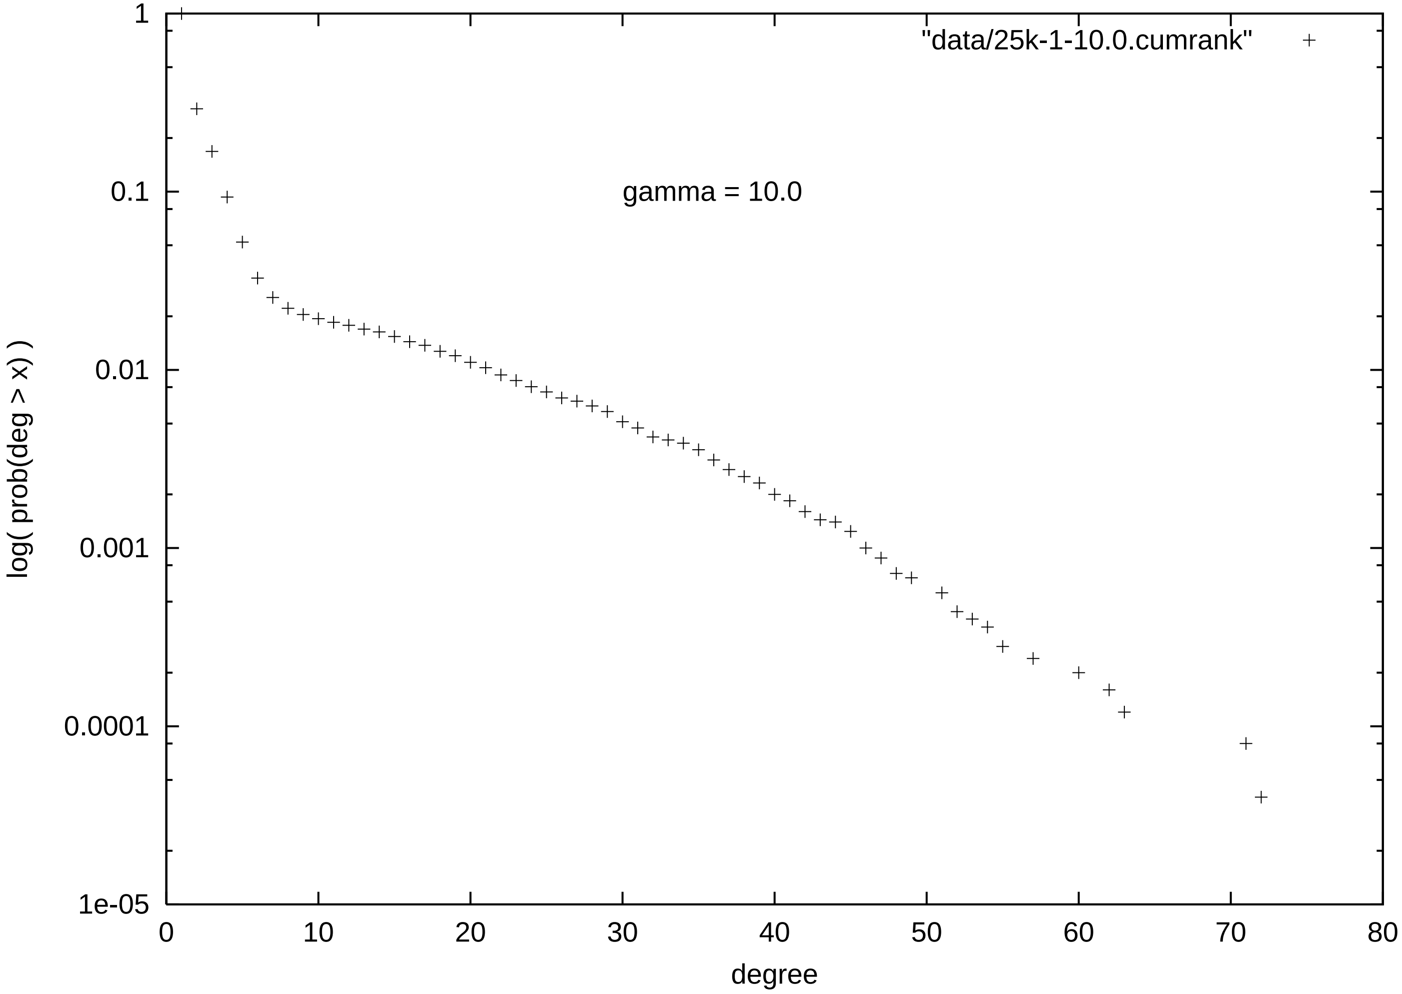

Degree distribution for $\gamma = 50.0$

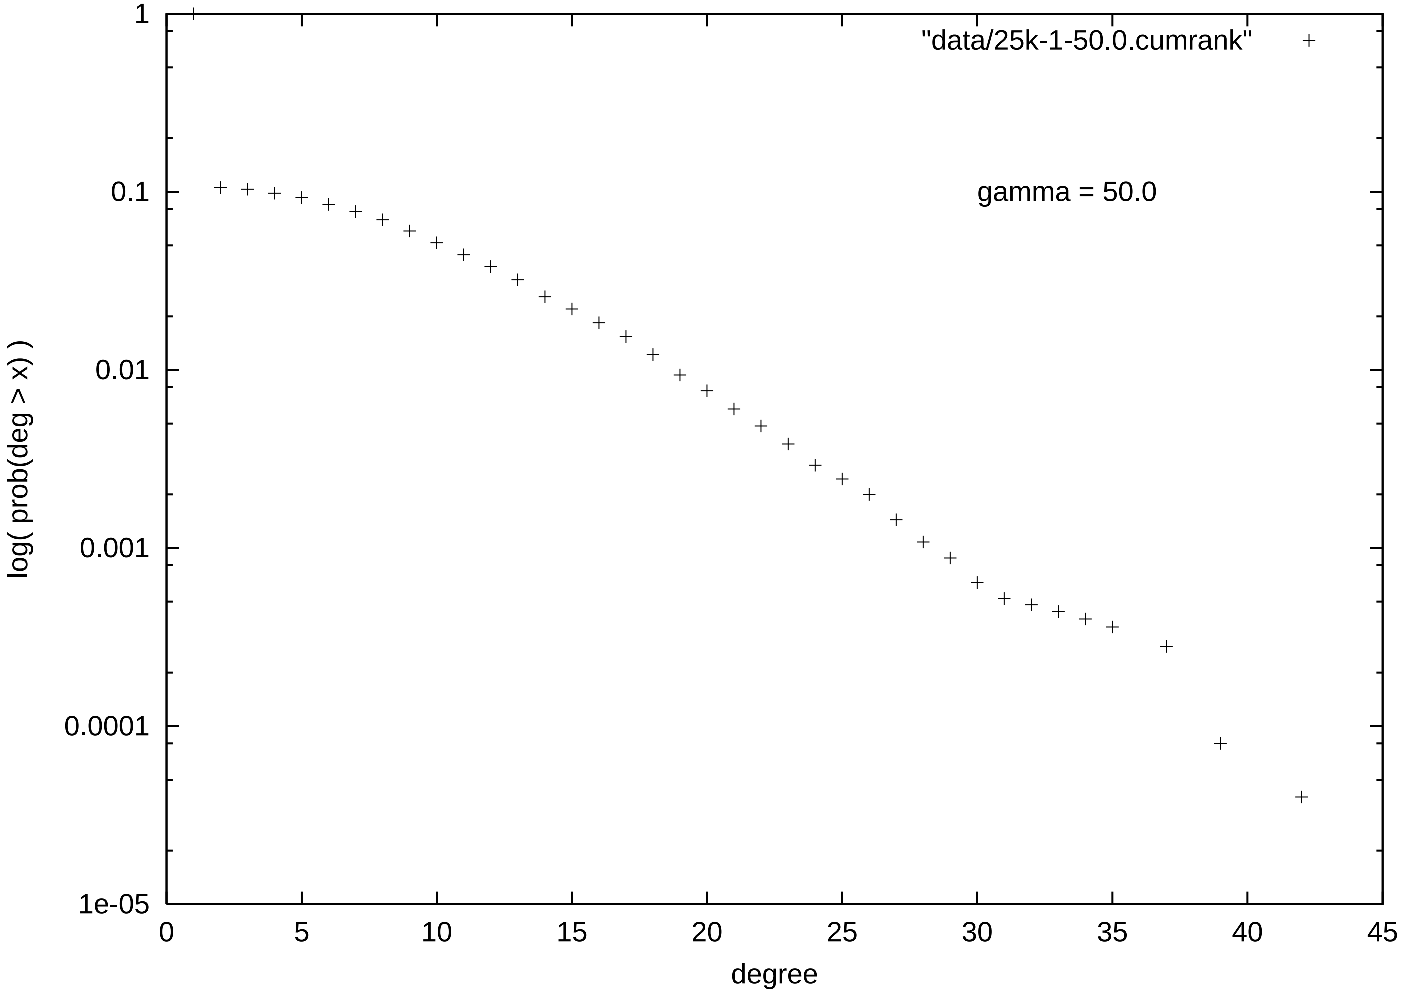

Degree distribution for $\gamma = 100.0$

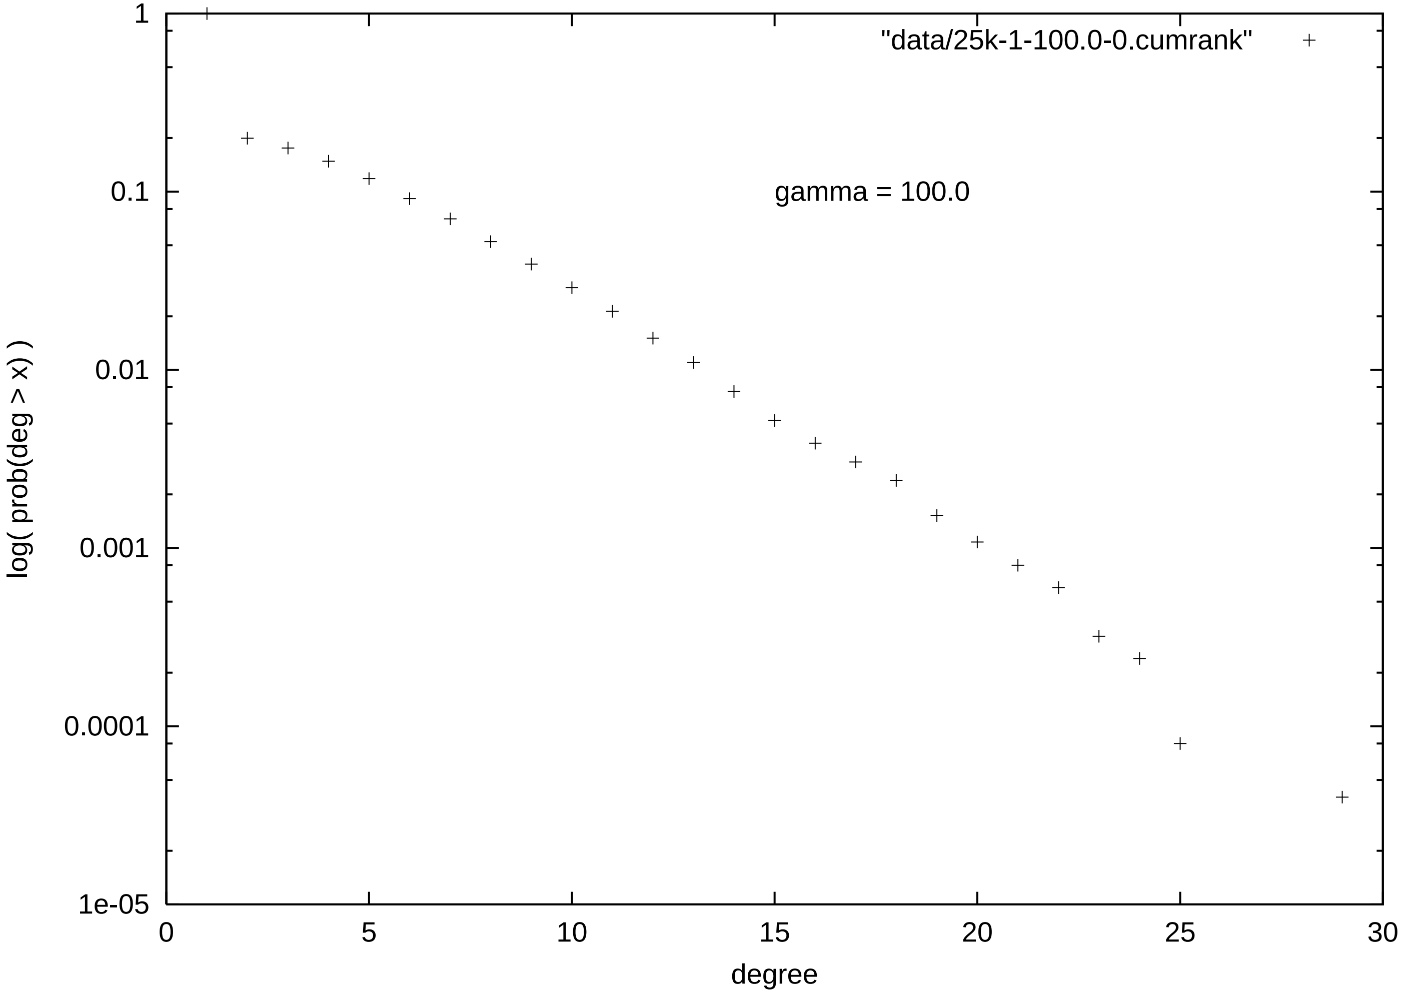

Degree distribution for $\gamma = 1.0, 10.0, 50.0, 100.0$

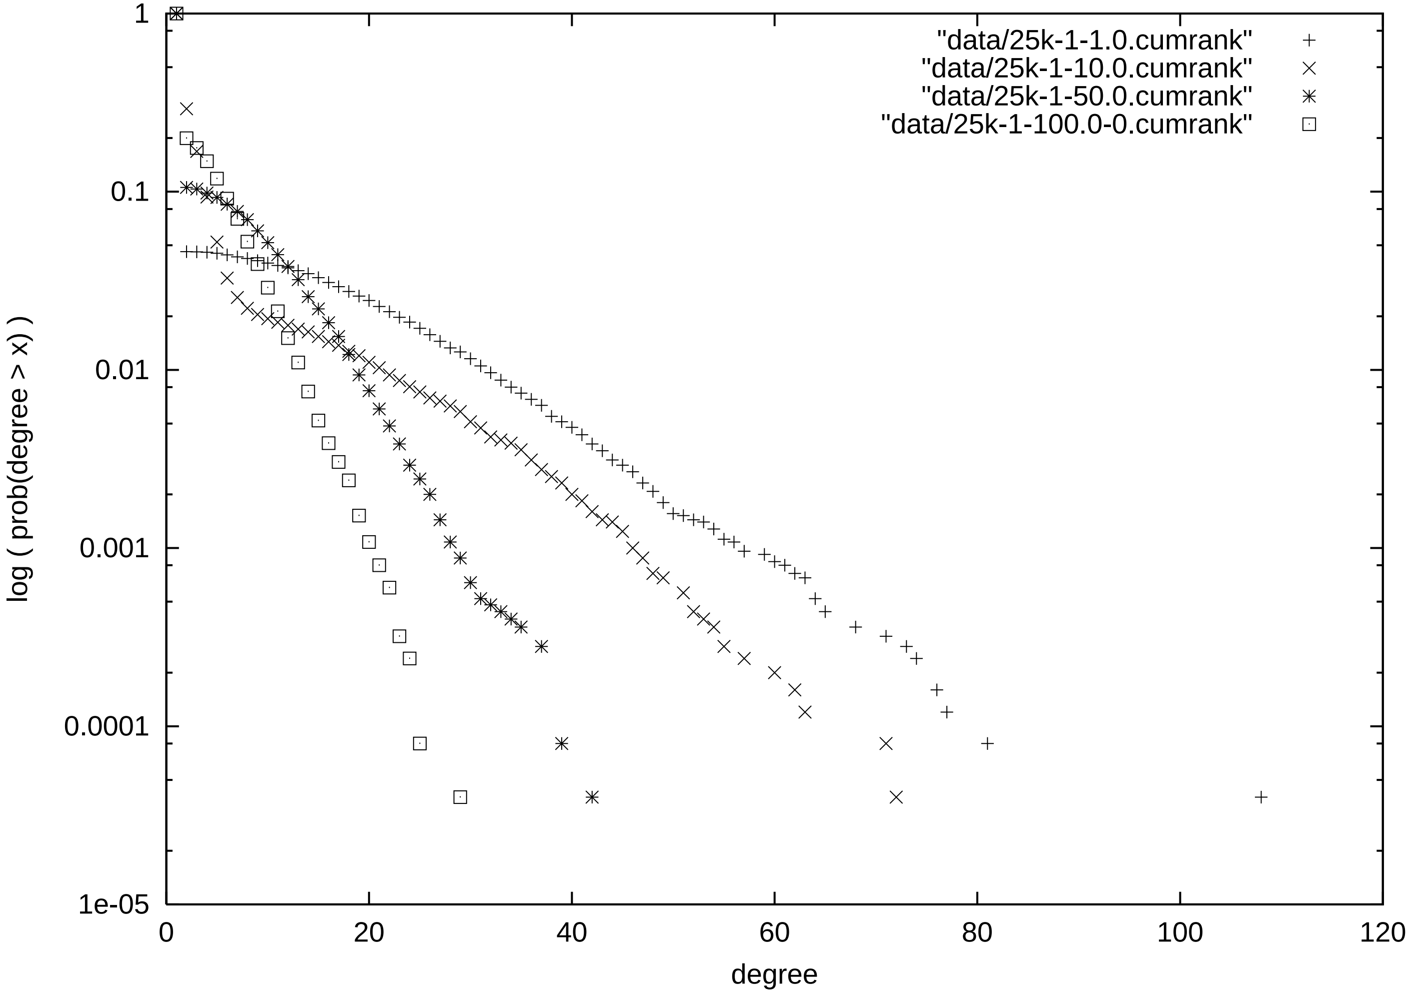

Probability of choosing cable type

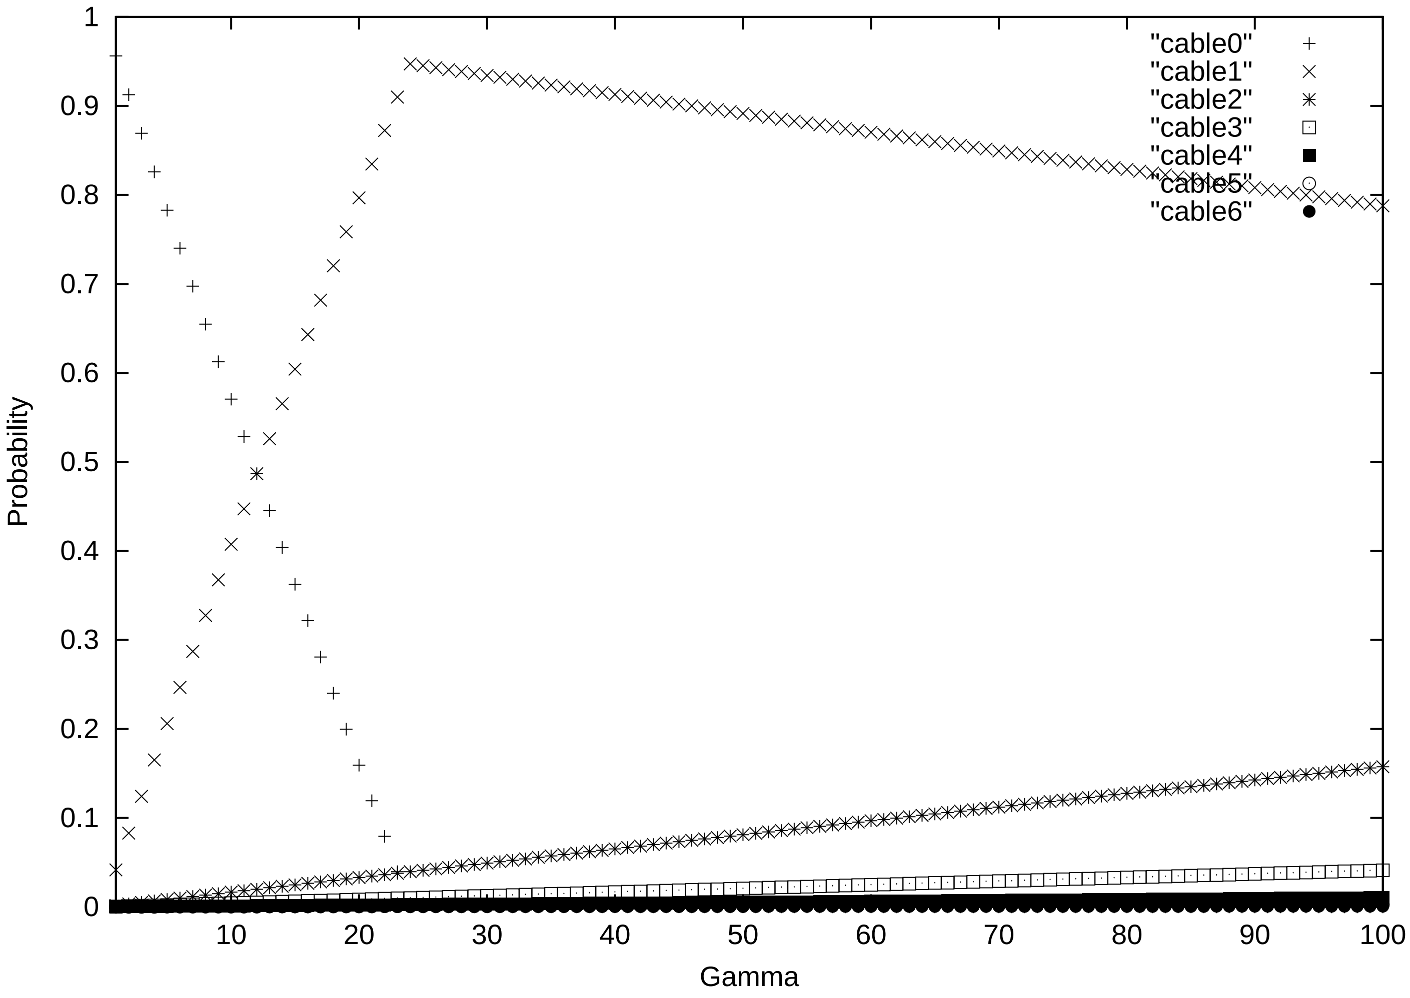

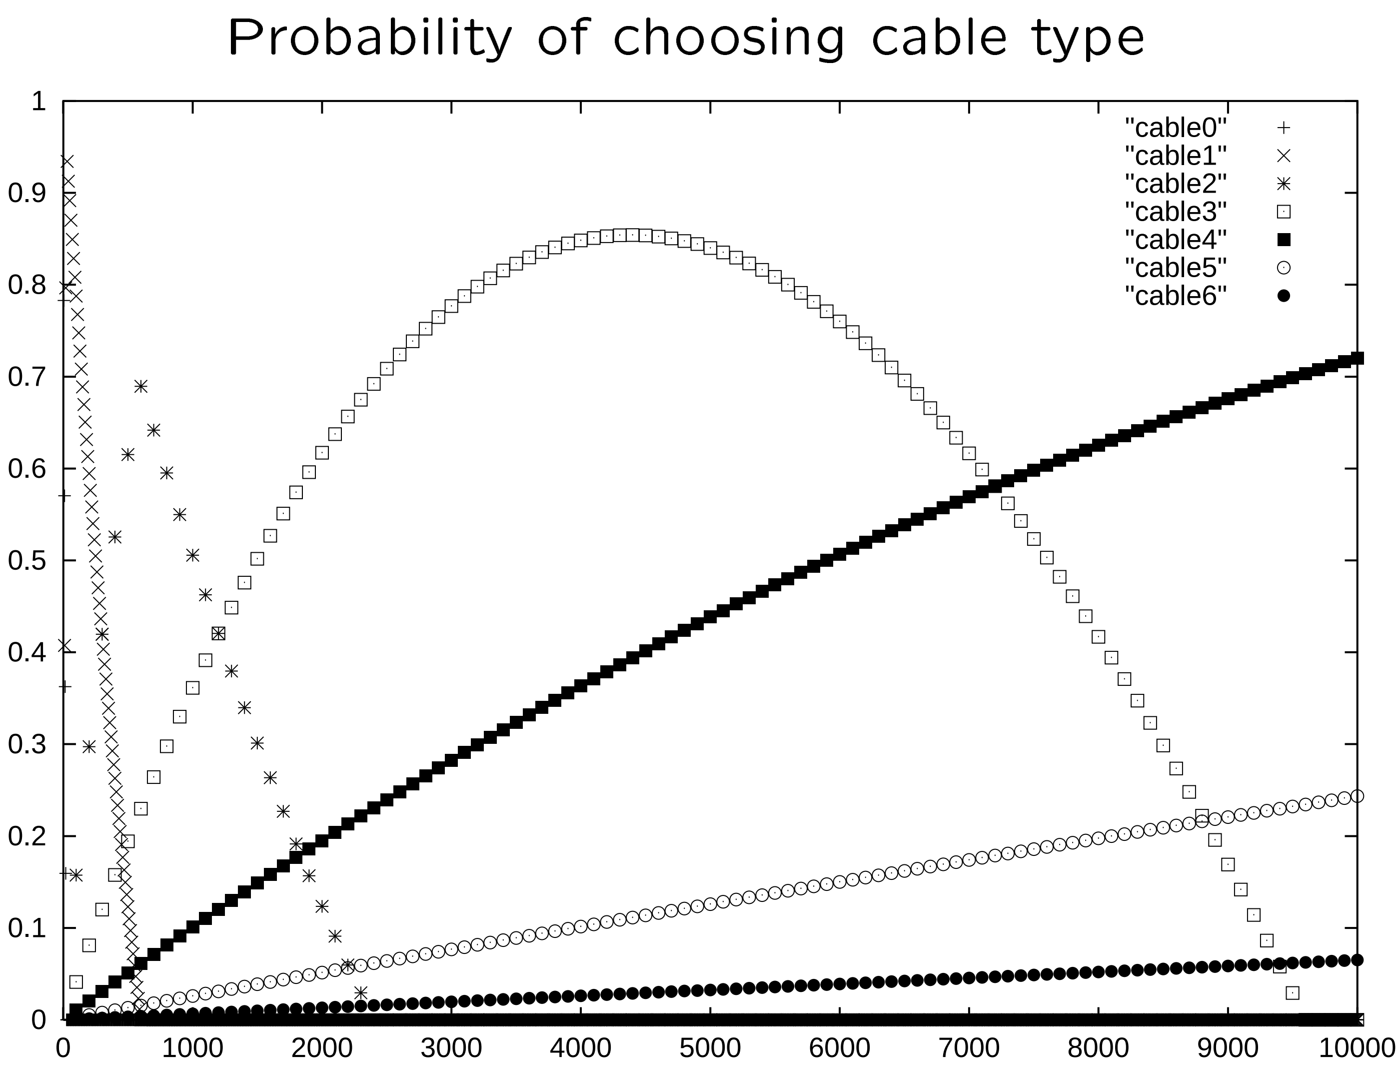
Probability of choosing cable type
"cable0"
"cable1"
"cable2"
"cable3"
"cable4"
"cable5"
"cable6"
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000

Initial Conclusions

1. A preliminary analysis of the topologies generated from the Meyerson algorithm indicates some evidence for an exponential distribution of node degree (not power laws), but we have not thoroughly searched the parameter space.

2. The parameter space is much more complicated than initially thought

- cable capacities (absolute and relative)

- γ tuning parameter to emphasize a region of the capacities

Related Work

Fabrikant, Koutsoupias, and Papadimitriou (2001) proposed a model for network access topology generation in the context of *heuristically optimized tradeoffs*.

$$\min_{j<i} \ \alpha \cdot dist(i,j) + h_j$$

$dist(i,j)$ = *distance between nodes* i *and* j
h_j = *measure of "centrality" (e.g. number of hops)*

Let $d(k)$ = the number of nodes with degree $\geq k$.
Basic result:
For $\alpha < 1/\sqrt{2}$, *resulting topology is a star*
For $\alpha = \Omega(\sqrt{n})$, $E[d(k)] < n^2 e^{-ck}$
For $\alpha \geq 4$ *and* $\alpha = O(\sqrt{n})$, $E[d(k)] = c(\frac{k}{n})^{-\beta}$

In Practice

Problems Solved by ISPs when building topology
- *Location and interconnectivity of metro core nodes*
- *Local access network design for metro area*
- *Interconnectivity of metro regions*

All of these problems can be formulated as optimization problems, given an understanding of
- *spatial distribution of demand points (geography)*
- *demand for traffic between points (population)*

This suggests the formulation of a grand optimization problem that generates a topology for a single ISP.

Generating ISP Topologies

The ability to generate "realistic, but fictitious" ISP topologies would be of great value

- *Supplement the lack of real ISP data*
- *Facilitate greater interaction with ISPs*
- *A framework for addressing broader technical and economic issues facing ISPs*
  - *Pricing*
  - *Provisioning*
  - *Peering*

This suggests yet another grand model in which individual ISP topologies are generated and then interconnected.
(Why not be bold?)

Ongoing Work

- *Identify the paramteres that are likely to lead to clustering (e.g. spatial distribution of demand, cable types)*
- *Explore additional optimization frameworks*
- *Understand, collaborate, leverage similar projects (e.g. RocketFuel project at U. Washington)*
- *Interaction/feedback with operators*
- *Economic framework*
- *Development of the grand model for the ISPs*