ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Spring Term 2016
Prof. Adrian Perrig

Operating Systems and Networks
Solution 9

Note: these solutions constitute material *supplemental* to the exercise sessions.

1 Dijkstra's Algorithm

A network between A and B is depicted in Figure 1. The numbers on the links correspond to the probability (multiplied by 100) that the link may fail. Link failures are independent from each other. How could you find the most reliable path from A to B? Hint: Use Dijkstra's algorithm.

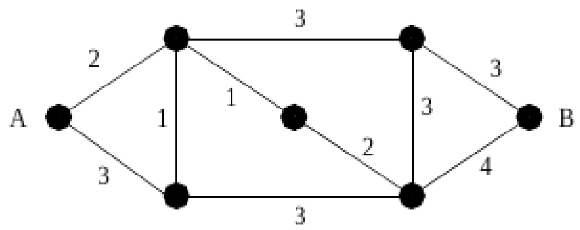

Figure 1: Network

Answer: Recall that Dijkstra's Algorithm computes the *shortest* paths in a network. We thus need to map our problem to an optimization problem (*shortest*) that respects the monotonicity of the Dijkstra operation (addition). Here is why: the combination of two paths of lengths ℓ_1 and ℓ_2 yields a new path of length $\ell = \ell_1 + \ell_2$, where $\ell > \ell_i$. For probabilities of link failures, this property does not hold. Two links with failure probabilities f_1 and f_2 do not simply yield a path with failure probability $f = f_1 + f_2$. Why? The probability that the composed path fails is computed as (1) link 1 fails and link 2 is up, or (2) link 1 is up and link 2 fails, or (3) both links fail.

It is thus easier to compute the probability of *availability* instead. Here, it holds that two links with availability probabilities a_1 and a_2 simply yield a path with availability probability $a = a_1 \cdot a_2$. Since we have that $a_i \leq 1$ for all i, the monotonicity property is preserved. We can thus apply Dijkstra with multiplication.

We first compute the availability probability for each link by $a_i = 100 - f_i$, as shown in Figure 2. Next, we run Dijkstra to obtain that the most reliable path from A to B is via C and D. The availability probability is 92.21, hence we obtain a failure probability of 7.79 for this path.[1]

[1] Note that this path corresponds to what one could consider the shortest paths from A to B if the numbers in the exercise were distances. This is a coincidence!

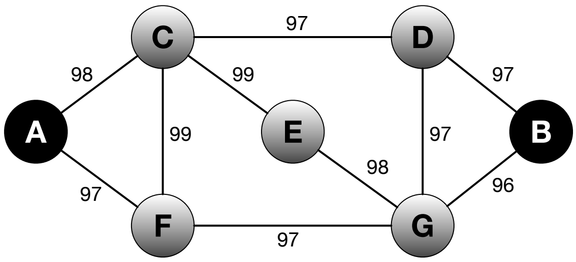

Figure 2: Network

2 Border Gateway Protocol

In BGP, ISPs announce "to interested parties" paths of networks which are reachable through those ISPs.

(1) Why don't residential ISPs (like Swisscom, UPC, Sunrise) send BGP announcements to their customers?

Answer: Customers of residential ISPs are *leafs* in the AS topology, meaning that there is only one way for customers to reach the internet (i.e., through their ISP). It makes no sense for BGP announcements to be sent to consumers because *all* traffic is always sent to the ISP first, and never any other destination.

(2) Would it make sense for home routers to send (to their ISP) BGP announcements advertising the private network address space of the customer?

Answer: No. First of all, private address space is not reachable directly, so the announcements would not benefit any BGP peers. Second, the ISP allocates an IP address to consumers, so they don't need any more information about what other networks are reachable behind that IP.

3 Hierarchical routing

IPv6 makes use of hierarchical routing to keep routing tables small. Use the data on `http://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xhtml` to find the region in which the following IPv6 addresses should exist:

(1) 2001:0200:0000:0000:beef:f00d:00c0:ffee

(2) 2001:1200:0000:0000:beef:f00d:00c0:ffee

(3) 2003:0000:0000:0000:beef:f00d:00c0:ffee

(4) 2001:4200:0000:0000:beef:f00d:00c0:ffee

(5) 2620:0000:0000:0000:beef:f00d:00c0:ffee

Answer:

(1) APNIC - Asia Pacific region

(2) LACNIC - Latin America

(3) RIPE - Europe

(4) Afrinic - Africa

(5) ARIN - North America