

Identifying Congestion Hotspots in MPLS Using Bayesian Networks

¹S.V. Kasmir Raja and ²P. Herbert Raj

¹SRM University, Chennai, India

²Department of Computer Science and Engineering, Alagappa University, Karaikudi, India

Abstract: Traffic Engineering (TE) broadly relates to optimization of the performance of a network. The overlay approach has been widely used by many service providers for Traffic Engineering (TE) in large Internet backbones. In the overlay approach, logical connections are set up between edge nodes to form a full mesh virtual network on top of the physical topology. IP routing is then run over the virtual network. Instead of overlaying IP routing over the logical virtual network, the integrated approach runs shortest path IP routing natively over the physical topology. Traffic engineering needs to determine the optimal routing of traffic over the existing network infrastructure by efficiently allocating resource in order to optimize traffic performance on an IP network. Traffic engineering objectives are achieved through carefully routing logical connections over the physical links. Common objectives of traffic engineering include balancing traffic distribution across the network and avoiding congestion hot spots. This study proposes a new approach called the Bayesian approach to avoid congestion hot spots without full mesh overlaying. This approach can be illustrated with a simple network, and then present a formal analysis of the Bayesian networks and a method for finding the congestion hot spots. Once, the congestion hot spots are identified then the traffic can be distributed, so that no link in the network is either over utilized or under utilized. With this Bayesian approach the quality of the routing can be improved and congestion can be avoided.

Key words: Traffic engineering, overlay model, integrated approach, congestion hot spots, bayesian network

INTRODUCTION

The unprecedented growth of the Internet has lead to a growing challenge among the ISPs to provide a good quality of service, achieve operational efficiencies and differentiate their service offerings. ISPs are rapidly deploying more network infrastructure and resources to handle the emerging applications and growing number of users. A routing specifies how to route the traffic between each origin-destination pair across a network. IP routing typically uses shortest-path computation with some simple metrics such as hop-count or delay. Although the simplicity of this approach allows IP routing to scale to very large networks, it does not make the best use of network resources (Awduche *et al.*, 1998).

In this study, a new approach is considered and that accomplishes traffic engineering objectives to find the congestion hot spots, without full-mesh overlaying concept and for achieving traffic engineering in the backbones. The formal analyses of the Bayesian networks propose a systematic method for finding the congestion likelihood in a network.

Enhancing the performance of an operational network, at both the traffic and the resource levels are the major objectives of traffic engineering (Awduche *et al.*,

2002). The goal of performance optimization of operational networks (Awduche *et al.*, 1998) is accomplished by routing traffic in a way to utilize network resources efficiently and reliably.

In large Internet backbones, service providers typically have to explicitly manage the traffic flows in order to optimize the use of network resources. This process is often referred to as traffic engineering.

Overlay approach: Currently, most large Internet backbones employ the so-called overlay approach for traffic engineering. With this approach service providers establish logical connections between the edge nodes of backbones and then overlay these logical connections onto the physical topology. Service providers can control the distribution of traffic over physical topology through carefully routing these logical connections over physical links. The optimal mapping between the logical connections and the physical links can be computed using a linear programming formulation.

While the overlay approach has been widely implemented on current Internet backbones, it suffers the so-called “N-Square” problem. As the size of the backbone network increases, the number of logical connections to be established will rise drastically, adding

considerable management complexity and messaging overheads. Second, while IP routing runs over such a fully meshed virtual network, each edge node has to establish routing peering with (N-1) other nodes. This poses a significant problem to current IP routers as most of them can not support a large number of peers. Note that multiple logical connections may go over the same physical link. Thus, the breakdown of a single physical link may cause multiple logical connections to fail and this will exaggerate the routing update load.

Integrated approach: Wang *et al.* (2001) proposed a new approach called Integrated approach that accomplishes traffic engineering objectives without full mesh overlaying. Instead of overlaying IP routing over the logical virtual network, the new approach runs shortest-path IP routing natively over the physical topology. It is theoretically proved that for any given traffic demands it is possible to select a set of link weights such that the shortest paths based on the selected link weights produce the same traffic distribution as that of the overlay approach with the assumption that traffic between the same source-destination pair can be split across multiple equal cost shortest paths, if exists.

Let us first illustrate with a simple example how the integrated approach works. Figure 1 shows a simple network topology, link capacities, and traffic demands. Each link has a capacity of 5 units and each demand needs bandwidth of 4 units. Although link capacities and traffic demands are unidirectional in IP networks, we assume they are bidirectional here for simplicity.

To meet the traffic engineering objectives, we need to place the demands over the links in a way that the traffic distribution is balanced and there is no congestion or hot spot in the network. The optimal routes can be calculated using a linear programming formulation (Yufei and Zheng, 1999).




Fig. 1: Topology and capacity

This Integrated approach has a number of advantages. First, it retains the simplicity of IP routing and requires little changes to the basic Internet architecture. Once the weights are calculated and set, the shortest-path routing protocol such as OSPF (Moy, 1998) can calculate the paths in the normal way, and packets are forwarded along the shortest paths. Second, it eliminates the “N-Square” problem all together and reduces managing overheads in setting up logical connections.

Common objectives of traffic engineering include balancing traffic distribution across the network and avoiding congestion hot spots. Shortest path chosen by IP routing does not always produce good network utilization. Poor utilization of network resources can be illustrated with the so-called fish problem. This leads to congestion hotspots in the network. The Bayesian approach is proposed here for finding the likelihood of congestion in the IP routing. Based upon the likelihood the traffic can be diverted through under utilized path or path having minimum likelihood for congestion. There by it is possible to improve the quality of the routing.

BAYESIAN NETWORKS

Graphical models are nothing but fusion of probability theory and graph theory. They provide a natural tool for dealing with two problems that occur throughout applied mathematics and engineering-uncertainty and complexity-and in particular they are playing an increasingly important role in the design and analysis of machine learning algorithms.

Representation: Probabilistic graphical models (Russell and Norving, 1995; Kevin, 2000) are graphs in which nodes represent random variables, and the arcs represent conditional independence assumptions. Hence, they provide a compact representation of joint probability distributions. Undirected graphical models, also called as Markov Random Fields (MRFs) or Markov networks, have a simple definition of independence: Two (set of) nodes A and B are conditionally independent given a third set, C, if all paths between the nodes in A and B are separated by a node in C. By contrast, directed graphical models also called Bayesian Networks or Belief Networks (BNs), have a more complicated notion of independence, which takes into account the directionality of the arcs.

For a directed model, we must specify the Conditional Probability Distribution (CPD) at each node. If the variables are discrete, this can be represented as a table (CPT), which lists the probability that the child node takes on each of its different values for each combination of values of its parents.

Table 1: CPD for the node “Cloudy”

P(C=F)	P(C=T)
0.5	0.5

Table 2: CPD for the node “Sprinkler”

C	P(S=F)	P(S=T)
F	0.5	0.5
T	0.9	0.1

Table 3: CPD for the node “Rain”

C	P(R=F)	P(R=T)
F	0.8	0.2
T	0.2	0.8

Table 4: CPD for the node “WetGrass”

C	R	P(W=F)	P(W=T)
F	F	1.0	0.0
T	F	0.1	0.9
F	T	0.1	0.9
T	T	0.01	0.99

Likelihood calculation: Let us consider the following example, in which all nodes are binary, i.e. have two possible values, which will be denoted by T (true) or F (false).

We see that the event “grass is wet” (W = true) has two possible causes: Either the water sprinkler is on (S = true) or it is raining (R = true). The strength of this relationship is shown in the Table 1-4. For example, we see that

$$\Pr(W=true | S=true, R=false) = 0.9$$

(second row) and hence,

$$\Pr(W=false | S=true, R=false) = 1 - 0.9 = 0.1,$$

since each row must sum to one. Since the C node has no parents, its CPT specifies the prior probability that it is cloudy (in this case, 0.5). (Think of C as representing the season: if it is a cloudy season, it is less likely that the sprinkler is on and more likely that the rain is on.)

The simplest conditional independence relationship encoded in a Bayesian network can be stated as follows: a node is independent of its ancestors given its parent, where the ancestor/parent relationship is with respect to some fixed topological ordering of the nodes.

By the chain rule of probability, the joint probability of all the nodes in the graph above is

$$P(C,S,R,W) = P(C) * P(S | C) * P(R | C,S) * P(W | C,S,R)$$

By using the conditional independence relationships, we can rewrite this as

$$P(C,S,R,W) = P(C) * P(S | C) * P(R | C) * P(W | C,S,R)$$

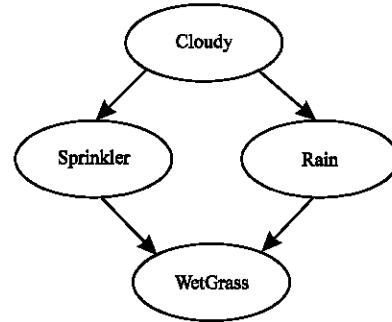


Fig. 2: Belief Network for event “grass is wet”

Where we were allowed to simplify the third term because R is independent of s given its parent C, and the last term because W is independent of C given its parents S and R.

The most common task to be solved using Bayesian networks is probabilistic inference. For example, suppose one observes the fact that the grass is wet. There are two possible causes for this: Either it is raining, or the sprinkler is on. Which is more likely? Then use Baye’s rule (Aji and McEliece, 2000; Kevin, 2000) to compute the posterior probability of each explanation (where 0 = false and 1 = true).

$$\text{Posterior} = \frac{\text{Likelihood} * \text{prior}}{\text{Marginal likelihood}}$$

or, in symbols,

$$P(R = r | e) = \frac{P(e | R = r) P(R = r)}{P(e)}$$

Where $P(R = r | e)$ denotes the probability that random variable R has value r given evidence e. The denominator is just a normalizing constant that ensures the posterior adds up to 1; it can be computed by summing up the numerator over all possible values of R, i.e.,

$$P(e) = P(R=0,e) + P(R=1,e) + \dots \dots \dots = \sum_r P(e | R=r) P(R=r)$$

This is called the marginal likelihood (since we marginalize out over R) and gives the prior probability of the evidence.

Applying Baye’s rule (Aji and McEliece, 2000; Kevin, 2000) to Fig. 2

$$\begin{aligned} P(S = 1 | W = 1) &= \frac{\Pr(S = 1, W = 1)}{\Pr(W = 1)} \\ &= \frac{\sum_{C,R} \Pr(C = c, S = 1, R = r, W = 1)}{\Pr(W = 1)} \\ &= \frac{0.2781}{0.6471} = 0.430 \end{aligned}$$

$$\begin{aligned}
 P(R = 1 | W = 1) &= \frac{\Pr(R = 1, W = 1)}{\Pr(W = 1)} \\
 &= \frac{\sum_{C,R} \Pr(C = c, S = s, R = 1, W = 1)}{\Pr(W = 1)} \\
 &= \frac{0.4581}{0.6471} = 0.703
 \end{aligned}$$

Where,

$$\begin{aligned}
 P(R = 1) &= \sum_{C,R,s} \Pr(C = c, S = s, R = r, W = 1) \\
 &= 0.6491
 \end{aligned}$$

is a normalizing constant, equal to the probability (likelihood) of the data. So, we see that it is more likely that the grass is wet because it is raining: The likelihood ratio is $0.7079/0.4298 = 1.647$.

The relationship between graphical models and bayes' rule: For complicated probabilistic models, computing the normalizing constant $P(e)$ is computationally intractable, either because there are an exponential number of (discrete) values of R to sum over, or because the integral over R cannot be solved in closed form (e.g., if R is a high-dimensional vector). Graphical models can help because they represent the joint probability distribution as a product of local terms, which can sometimes be exploited computationally (e.g., using dynamic programming or Gibbs sampling). Bayes nets (directed graphical models) are a natural way to represent many hierarchical Bayesian models.

To meet the traffic engineering objectives, demands have to be placed over the links in order to achieve balanced traffic distribution and to avoid congestion hot spots in the network. The Bayesian network method for finding likelihood can be applied for identifying the congestion hot spots. In the study, new Bayesian network method is proposed to measure the congestion likelihood in the integrated approach.

IDENTIFYING CONGESTION HOT SPOTS USING BAYESIAN NETWORKS

One approach to overcome the “traffic distribution problem” is to identify the congestion hot spots in the network. Optimal routes having minimum congestion likelihood can then be calculated. Instead of routing the demands over the congested routes we can select the routes that suit the current traffic demand and capacity of the network.

Let us first illustrate with a simple example how to calculate the congestion likelihood. Figure 3 shows a simple network topology and link capacities.




Fig. 3: Bayesian network for finding congestion likelihood

Table 5: CPD for the node “A”

A(A = 0)	P(A = 1)
0	1

Table 6: CPD for the node “D”

A	P(D = 0)	P(D = 1)
0	1	0
1	0.01	0.99

Table 7: CPD for the node “G”

D	P(G = 0)	P(G = 1)
0	1	0
1	0.9	0.1

Table 8: CPD for the node “C”

D	P(C = 0)	P(C = 1)
0	1	0
1	0.1	0.9

Table 9: CPD for the node “E”

G	C	P(E = 0)	P(E = 1)
0	0	1.0	0.0
0	1	0.01	0.99
1	0	0.2	0.8
1	1	0.01	0.99

Let us consider the following example, in which all nodes are binary, i.e. have two possible values, which will be denoted by 1 (true) or 0 (false).

Event “A is the source and E is the destination” (E = 1) has two possible causes: either the demand is routed through node G (G = 1) or through node C (C = 1). The strength of this relationship is shown in the Table 5-8. By applying the Baye’s rule to Fig. 3, the calculation of the congestion likelihood from source A to destination E through two different routes ADGE and ADCE as follows.

In the Table 5, probability of the node present is 1 (ie A=1). The CPD of node D,G,C and E presented in Table 6-9, respectively.

For example, from Table 9, we see that,

$$\Pr(E=1 | G=0, C=1) = 0.99$$

(second row), and hence,
Where,

$$\Pr(E=0 | G = 0, C = 1) = 1 - 0.99 = 0.01,$$

since, each row must sum to one.
Congestion likelihood of the route ADCE,

$$\begin{aligned} P(C = 1 | E = 1) &= \frac{\Pr(C = 1, E = 1)}{\Pr(E = 1)} \\ &= \frac{\sum_{A,D,G} \Pr(A = a, D = d, C = 1, G = g, E = 1)}{\Pr(E = 1)} \\ &= 0.9911 \end{aligned}$$

Congestion likelihood of the route ADGE,

$$\begin{aligned} P(G = 1 | E = 1) &= \frac{\Pr(G = 1, E = 1)}{\Pr(E = 1)} \\ &= \frac{\sum_{A,D,C} \Pr(A = a, D = d, C = 1, G = 1, E = 1)}{\Pr(E = 1)} \\ &= 0.10801 \end{aligned}$$

Where,

$$\begin{aligned} \Pr(E = 1) &= \sum_{A,D,C,G} \Pr(A = a, D = d, C = c, G = g, E = 1) \\ &= 0.89001 \end{aligned}$$

From the above calculations, the route ADCE is more congested when compared to that of the other route ADGE. So, by calculating the likelihood, the congestion hot spots can be identified in a network.

CONCLUSION

The new Bayesian approach is proposed for achieving traffic engineering in the backbones. Instead of relying on the mapping of logical connections of physical links to manage traffic flows in the network, we

run IP routing natively over the physical topology and control the distribution of traffic flows through setting appropriate link weights for shortest path routing.

Common objectives of traffic engineering include balancing traffic distribution across the network and avoiding congestion hot spots. Shortest path chosen by IP routing does not always produce good network utilization. Poor utilization of network resources can be illustrated with the so-called fish problem. This leads to congestion hotspots in the network. We proposed a Bayesian approach for finding the likelihood of congestion in the IP routing. Based upon the likelihood traffic can be diverted through the load of the under utilized path or path having minimum likelihood for congestion. For any set of feasible routes, the congestion likelihood can be calculated with the help of the Bayesian network approach. So, it is possible to improve the quality of the routing.

REFERENCES

- Aji, S.M. and R.J. McEliece, 2000. The generalized distributive law. *IEEE. Trans. Inform. Theory*, 46: 325.
- Awduche, D., J. Malcolm, M. O'Dell and J. McManus, 1998. Requirements for traffic engineering over MPLS, IETF. Internet draft, draft-awduche-mpls-traffic-engg-00.txt.
- Awduche, D. *et al.*, 2002. Overview and principles of Internet Traffic Engineering, IETF. Internet draft, draft-ietf-tewg-principles-02.txt.
- Kevin Murphy, 2000. A brief introduction to Baye's rule. Web: <http://www.cs.ubc.ca/~murphyk/Bayes/bayes.html>. <http://www.cs.ubc.ca/~murphyk/Bayes/bayesrule.html>.
- Moy, J., 1998. OSPF version 2, Internet rfc, IETF. Technical report RFC 2328.
- Russell, S. and P. Norvig, 1995. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ.
- Yufei Wang and Zheng Wang, 1999. Explicit routing algorithms for internet traffic engineering, Proceedings of ICCCN.
- Yufei Wang, Zheng Wang and Leah Zhang, 2001. Internet traffic engineering without full mesh overlaying, Infocom.