

## Opti Routing: An Optimized Routing Based On Traffic Monitoring

P. Madhukrishna M.Tech,  
CSE Department,  
KVSRTC College, Kurnool.

B. Mahesh Assoc.Professor,  
CSE Department,  
KVSRTC College, Kurnool.

### **Abstract—**

Monitoring the traffic at one or more points in a network is of interest to network operators for reasons of traffic accounting, debugging and traffic engineering. Previous research in the area has focused on deriving a placement of monitors across the network toward the end of maximizing the monitoring utility of the network operator for a given traffic routing. Here traffic characteristics and measurement objectives can dynamically change over time, rendering a previously optimal placement of monitors suboptimal. It is not feasible to dynamically redeploy/reconfigure measurement infrastructure to cater to such evolving measurement requirements. can overcome this problem by strategically routing traffic subpopulations over fixed monitors. We refer to this approach as *OptiRouting*. The main challenge for *OptiRouting* is to efficiently utilizing bandwidth resources and meeting quality-of-service. A fundamental feature of intradomain routing, which makes *OptiRouting* feasible, is that intradomain routing is often specified for aggregate flows. *OptiRouting* can therefore differentially route components of an aggregate flow while ensuring that the aggregate placement is compliant to original traffic engineering objectives. In this paper, we present a theoretical framework for *OptiRouting*. Furthermore, as proofs of concept, we present synthetic and practical monitoring applications to showcase the utility enhancement achieved with *OptiRouting*.

**Index Terms—**Anomaly detection, intradomain routing, network management, traffic engineering, traffic measurements.

### **I. INTRODUCTION**

Several past research efforts have focused on the optimal deployment of monitoring infrastructure in operational networks for accurate and efficient measurement of network traffic. Such deployment involves both monitoring infrastructure placement as well as configuration decisions. An example of the former includes choosing the interfaces at which to install DAG cards, and the latter includes tuning the sampling rate and sampling scheme of the DAG cards.

The optimal placement and configuration of monitoring infrastructure for a specific measurement

objective typically assumes a priori knowledge about the traffic characteristics. Furthermore, these are typically performed at longer timescales to allow provisioning of required physical resources. However, traffic characteristics and measurement objectives may evolve dynamically, potentially rendering a previously determined solution suboptimal.

We propose a new approach called *OptiRouting* to address this limitation. *OptiRouting* forwards network traffic across routes where it can be best monitored. Our approach is complementary to the well investigated monitor placement problem [1]–[3] that takes traffic routing as an input and decides where to place monitors to optimize measurement objectives; *OptiRouting* takes monitor deployment as an input and decides how to route traffic to optimize measurement objectives. Since routing is dynamic in nature (a routing decision is made for every packet at every router), *OptiRouting* can conceptually adjust to changing traffic patterns and measurement objectives. In this paper, our focus is on the overall monitoring utility, defined as a weighted sum of the monitoring achieved over all flows.

The main challenge for *OptiRouting* is to work within the constraints of existing intradomain traffic engineering (TE) operations that are geared for efficiently utilizing bandwidth resources, or meeting quality-of-service (QoS) constraints, or both. This paper presents a framework for *OptiRouting* that allows rerouting traffic toward the end of optimizing an ISP's measurement objectives while being compliant to TE constraints. Our framework is generic and can be leveraged for a wide variety of measurement scenarios. We highlight a few examples as follows.

- A simple scenario involves routers implementing uniform sampling or an approximation of it, with network operators being interested in monitoring a subset of the traffic. *OptiRouting* can be used to make important traffic traverse routes that maximize their overall sampling rate.
- Networks might implement heterogeneous sampling algorithms, each optimized for certain kinds of traffic subpopulations. For instance, some routers can implement sophisticated algorithms to give accurate

flow size estimates of medium sized flows that otherwise would not have been captured by uniform sampling. OptiRouting can then route traffic subpopulations that might have medium sized flows across such routers. A network can have different active and passive measurement infrastructure and algorithms deployed, and OptiRouting.

### 1.1 Can direct traffic across paths with greater measurement potential.

OptiRouting can be used to conserve measurement resources. For instance, all packets belonging to a certain traffic subpopulation can be conjointly routed to avoid maintaining states across different paths. Similarly, if the state at a node is maintained using probabilistic data structures (such as sketches), OptiRouting can enhance the accuracy of such structures by selecting the traffic that traverses the node.

This paper presents a general routing framework for OptiRouting, assuming the presence of special forwarding mechanisms. We present three flavors of OptiRouting, each of which works with a different set of compliancy constraints, and we discuss two applications as proofs of concept. These OptiRouting applications illustrate the significant improvement achieved by this additional degree of freedom in tuning how and where traffic is monitored.


This paper is an extended version of our previous work [4], which we believe to be the first attempt to leverage routing as a degree of freedom for monitoring traffic. The present work extends upon [4] as follows.

- The results in [4] indicated that the performance of OptiRouting is sensitive to the number of paths present between pairs of nodes. It is the relative difference in measurement capacity across such paths between a pair of nodes that is leveraged by OptiRouting to improve monitoring performance. Whereas [4] showed significant performance gains for OptiRouting, the choice of experimental networks was restricted to networks with a very low number of paths present between node pairs. This paper reports the results for a more realistic set of networks (higher average degree), contributing to a more realistic performance evaluation of OptiRouting.
- The fundamental idea behind OptiRouting is to divide traffic aggregates into subpopulations and then differentially route the traffic subpopulations based on the monitoring capacity of available routes and the relative measurement importance of the traffic subpopulations. It was observed in [4] that the way traffic aggregates are decomposed into multiple subpopulations has an impact on OptiRouting performance. This paper extends upon [4] by introducing additional and more involved

decomposition methods than those presented in [4], resulting in improved OptiRouting performance.

- We also take a closer look at the solution computation times of OptiRouting problems and their scalability in Section IV-A-VI. We present an approximation algorithm that allows one to tradeoff OptiRouting performance for faster computation times.
- Finally, in Section VI, we discuss issues encountered in deploying OptiRouting solutions in real networks and dynamic environments where both network applications and measurement objectives may keep changing.

The rest of this paper is organized as follows. We present an overview of OptiRouting in Section II. Section III details the OptiRouting framework. Our example monitoring applications and a detailed performance evaluation are given in Section IV. Section V presents related work. We conclude in Section VI.



**Fig. 1. Illustration of using routing to focus on a traffic subpopulation. In the above example, router has special sampling of interest to us. To apply this sampling on Flowset 2, we can route through router A, while (b) violating, or (c) being compliant to TE policy. (a) Original. (b) Violating. (c) Compliant.**

## II. MEASURROUTING OVERVIEW

As mentioned in Section I, OptiRouting must be cognizant of any implications that rerouting traffic has on TE policy. They are three fundamental ways in which OptiRouting enhances traffic monitoring utility without violating TE policy.

### 2.1 TE policy is usually defined for aggregated flows:

On the other hand, traffic measurement usually deals with a finer level of granularity. For instance, we often define a flow based upon the five-tuple. For measurement purposes. Common intradomain protocols (IGPs) like OSPF [5] and IS-IS [6] use link weights to specify the placement of traffic for each origin–destination (OD) pair (possibly consisting of millions of flows). The TE policy is oblivious of how constituent flows of

an OD pair are routed as long as the aggregate placement is preserved. It is possible to specify traffic subpopulations that are distinguishable from a measurement perspective but are indistinguishable from a TE perspective. OptiRouting can, therefore, route our fine grained measurement traffic subpopulations without disrupting the aggregate routing. The example depicted in Fig. 1 illustrates this argument. It shows four traffic subpopulations and that have the same ingress and egress nodes. Suppose that and are of equal size. Router has some dedicated monitoring equipment, and it is important for the network operator to monitor.

Our TE policy is to minimize the maximum link utilization. Fig. 1(a) depicts the original routing that obeys the TE policy. Fig. 1(b) represents a routing that violates the TE policy in order to route through route. However, if the traffic subpopulations are routed as in Fig. 1(c) is allowed to pass through the dedicated monitoring equipment, and the routing is indistinguishable from the original from the perspective of our TE policy. It is important to note that the aggregate traffic must span multiple paths in order for OptiRouting to be useful in this way. If the aggregate traffic traverses a single path, then no opportunity exists to differentially route sub-sets of the traffic.

- The second way in which OptiRouting is useful stems from the definition of TE objectives. TE objectives may be oblivious to the exact placement of aggregate traffic and only take cognizance of summary metrics such as the maximum link utilization across the network. An aggregate routing that is slightly different from the original routing may still yield the same value of the summary metric. Suppose and pertain to two different OD pairs in Fig. 1(a). Then, the new routing depicted by Fig. 1(c) changes the aggregate traffic placement discussed above. However, from a TE perspective, the total link utilization of all links remains the same.
- Finally, a network operator can specify a certain permissible level of TE policy violations. Such a specification would enable a tradeoff between the advantage derived from OptiRouting and adherence to TE policy. For instance, if the network operator is willing to allow a 33% increase in the maximum link utilization, the routing in Fig. 1(b) becomes a compliant solution.

The above discussion deals with the requirement that OptiRouting must operate within the confines of the TE policy. The other equally important challenge is that any OptiRouting solution should be physically realizable according to the constraints of the

underlying forwarding mechanisms. For instance, in order to selectively route a certain traffic subpopulation, the capability must exist to execute the requisite forwarding. This introduces a host of issues. It would require state to be maintained for all traffic subpopulations and might impose limits on the cardinality or the membership of such traffic subpopulations. Other concerns may stem from the exact routing protocols used to implement OptiRouting. For instance, a routing protocol may impose a constraint that traffic between a pair of nodes may only traverse paths that are along shortest paths with respect to certain link weights.

### III. MEASURING FRAMEWORK

We now present a formal framework for OptiRouting in the context of a centralized architecture. A centralized architecture refers to the case where the algorithm deciding how distributed nodes will route packets using OptiRouting has global information of: 1) the TE policy; 2) the topology and monitoring infrastructure deployment; and 3) the size and importance of traffic subpopulations.

A *macro-flowset* represents a set of flows for which an aggregate routing placement is given. In the context of intradomain IP routing, a macro-flowset comprises all flows between an OD pair. For MPLS networks, macro-flowsets can be defined as all flows between an ingress-egress pair in the same QoS class. Our only requirement is that flows in a macro-flowset have the same ingress and egress nodes. In this paper, we consider all flows between an OD pair to constitute a single macro-flowset.

#### Algorithm:

```

1:  $\Psi_{x:y \in Y_x} = \emptyset$ 
2: for all  $(i, j) \in E$  do
3:   if  $\Gamma_{x:y \in Y_x}$ 
4:      $ij > 0$  then
5:        $\Psi_{x:y \in Y_x}$ 
6:        $\leftarrow \Psi_{x:y \in Y_x} \cup \{(i, j)\}$ 
7:     end if
8:   end for
9:  $\hat{E} \leftarrow E / \Psi_{x:y \in Y_x}$ 
10:  $\{A \text{ specific order of choosing links in } \hat{E} \text{ may be specified for the following part}\}$ 
11: for all  $(i, j) \in \hat{E}$  do
12:   if Induced graph of  $\Psi_{x:y \in Y_x} \cup \{(i, j)\}$  is acyclic then
13:      $\Psi_{x:y \in Y_x}$ 
14:      $\leftarrow \Psi_{x:y \in Y_x}$ 

```

```

U {(i, j)}
12: end if
13: ^E ← ^ E/(i, j)
14: end for
    
```

### IV. PERFORMANCE EVALUATION

This section evaluates the performance of OptiRouting for specific monitoring applications. A OptiRouting application can be defined by specifying the sampling resolution function( $\beta$ ), and its constituents i.e., link sampling characteristics( $\{S\} (i,j) \in E$ ) and micro-flowset sampling utilities ( $\{I\} y \in \theta$ ).

I proceed to define and study two OptiRouting applications in § IV-A and § IV-B. For both applications we consider the utilization of the most congested link as our TE metric, i.e.,  $\sigma\Gamma$  and  $\sigma\gamma$  represent the maximum link utilization resulting from the original and micro-flowset routing, respectively. We also have a common definition of the link sampling characteristics across both our applications. The sampling characteristic

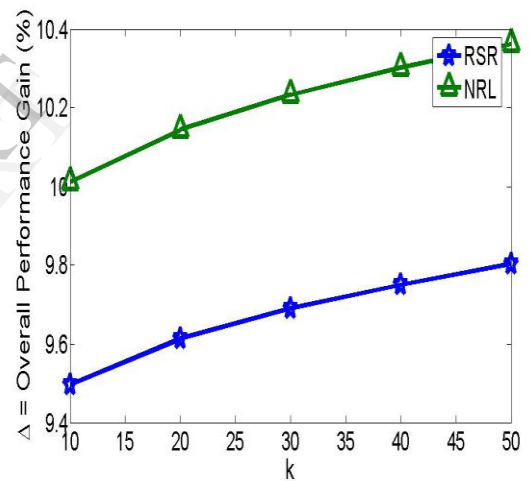

of a link  $(i, j)$ ,  $S(i, j)$  is equal to  $p_{ij} \in [0, 1]$ , where  $p_{ij}$  represents the known sampling rate of link  $(i, j)$ . We have a set of flows  $F$ . Each flow  $f \in F$  has an associated ingress node  $inf \in V$  and egress node  $outf \in V$ . I can represent the traffic demand of flow  $f$  by  $bf$ , and the importance or utility of sampling it by  $if$

We define  $k$  to be the total number of micro-flowsets for each macro flowset. We use  $v_y \in \theta$  to represent the set of flows that belong to the micro-flowset  $y$ . It follows that the aggregate traffic demand for macro flowset  $x$  is given by  $\Phi_x = \sum_{f \in F_x} bf$ . Most IP networks use link state protocols such as OSPF [4] and IS-IS [5] for intradomain routing. In such networks, every link is assigned a cost and traffic between any two nodes is routed along minimum cost paths. Setting link weights is the primary tool used by network operators to control network load distribution and to accomplish TE objectives. I use the popular local search Meta heuristic in [9] to optimize link weights with respect

to our aggregate traffic demands  $\{\Phi\} x \in \Theta$ . The optimized link weights are then used to derive our original routing  $\{\Gamma\} x \in \Theta (i, j) \in E$ .

Our applications are differentiated on the basis of the set of flows  $F$ , and how we assign the sampling importance  $if$  and the traffic demand  $bf$  of each flow  $f \in F$ . For both our applications we can consider the importance of a flow  $f$ ,  $if$ ,

to be the points we earn if we were to sample a byte for that flow. I wish to maximize the the total number of points earned, by routing our traffic across the given topology. This total number of points is given by the following:



The OptiRouting formulation requires us to specify the sampling

utility function for each micro-flowset. Towards this end we define the sampling utility function as  $I_y \in \theta = \sum_{f \in v_y} if bf$ . Thus, the sampling utility of a micro-flowset is the sum of sampling utilities of its flows weighted by the flow sizes.

Note that according to our definition  $\beta = \Delta MR$ . Therefore, for a given flows to micro-flowset assignment, maximizing  $\beta$  is equivalent to maximizing  $\Delta MR$  and  $\Delta$ .

## V. RELATED WORK

Earlier work in the area of traffic monitoring has focused on

1) inferring characteristics of original traffic from sampled traffic, 2) investigating and improving the effect of oblivious sampling on monitoring certain traffic subpopulations, and 3) placing monitor agents at certain strategic network locations. We summarize existing work in these three areas. Claffy *et al.* [16] compared various sampling approaches at both packet-based and time based granularities [16]. Several other research efforts aim to improve estimation of “heavy-hitter” traffic volume, flow-size distributions, traffic matrices, or flow durations [17–19] [20–24]. Recent work has demonstrated that conventional sampling techniques can obscure statistics needed to detect traffic anomalies [25] or execute certain anomaly detection algorithms [26]. All these previous works highlight the importance of being able to focus on specific traffic subpopulations. [27] proposes ways to focus monitoring budget on a specific traffic subpopulation by defining individual bins based on one or more tuples and allocating sampling budget to each bin. The traffic belonging to individual bins are identified using a counting bloom filter. There exists other proposals [28, 29] that also define the traffic subpopulation in a flexible manner. All of the above mentioned works are orthogonal in nature to our proposal as their work focuses on improving monitoring at one monitor, while our work tries to route traffic to make best use of these monitors. The closest research efforts to ours are those presented in [1–3, 30, 31], which aim to achieve effective coordination across multiple traffic monitors to improve network-wide flow monitoring.

## VI. FUTURE DIRECTIONS

Future work also involves deploying OptiRouting in a realistic network substrate that is programmable. Specifically, we intend to implement OptiRouting over Open Flow [7]. The controller in such a system will feature an online version of the OptiRouting algorithm that takes into account forecasting information of flows to determine an approximate solution at runtime. We would like to study the effect of additional constraints imposed by such realistic network substrates. We believe that our framework provides a firm foundation for routing assisted traffic engineering. We speculate that, in conjunction with programmable routing agents, our framework will stimulate further OptiRouting applications.

## REFERENCES

- [1] C. Chaudet, E. Fleury, I. G. Lassous, H. Rivano, and M.-E. Voge, “Op-timal positioning of active and passive monitoring devices,” in *Proc. ACM CoNEXT*, Toulouse, France, Oct. 2005, pp. 71–82.
- [2] K. Suh, Y. Guo, J. Kurose, and D. Towsley, “Locating network monitors: Complexity, heuristics and coverage,” in *Proc. IEEE INFOCOM*, Miami, FL, Mar. 2005, vol. 1, pp. 351–361.
- [3] G. R. Cantieni, G. Iannaccone, C. Barakat, C. Diot, and P. Thiran, “Re-formulating the monitor placement problem: Optimal network-wide sampling,” in *Proc. ACM CoNEXT*, Lisboa, Portugal, Dec. 2006, Ar-ticle no. 5.
- [4] S. Raza, G. Huang, C.-N. Chuah, S. Seetharaman, and J. P. Singh, “MeasuRouting: A framework for routing assisted traffic monitoring,” in *Proc. IEEE INFOCOM*, San Deigo, CA, Mar. 2010, pp. 1–9.
- [5] “OSPF,” The Internet Society, RFC 2328, 1998 [Online]. Available: <http://tools.ietf.org/html/rfc2328>
- [6] “IS-IS,” The Internet Society, RFC 1142, 1990 [Online]. Available: <http://tools.ietf.org/html/rfc1142>
- [7] C. Wiseman, J. Turner, M. Becchi, P. Crowley, J. DeHart, M. Hait-jema, S. James, F. Kuhns, J. Lu, J. Parwatikar, R. Patney, M. Wilson, K. Wong, and D. Zar, “A remotely accessible network processor-based router for network experimentation,” in *Proc. ACM/IEEE ANCS*, San Jose, CA, Nov. 2008, pp. 20–29.
- [8] “The OpenFlow Switch Consortium,” Stanford University, Stanford, CA [Online]. Available: <http://www.openflowswitch.org>
- [9] R. Morris, E. Kohler, J. Jannotti, and M. F. Kaashoek, “The Click modular router,” in *Proc. ACM SOSP*, Charleston, SC, Dec. 1999, pp. 217–231.
- [10] B. Fortz and M. Thorup, “Internet traffic engineering by optimizing OSPF weights,” in *Proc. IEEE INFOCOM*, Tel-Aviv, Isreal, Mar. 2000, vol. 2, pp. 519–528.
- [11] A. Medina, N. Taft, S. Battacharya, C. Diot, and K. Salamatian, “Traffic matrix estimation: Existing techniques and new directions,” in *Proc. ACM SIGCOMM*, Pittsburgh, PA, Aug. 2002, pp. 161–174.
- [12] “BRITE,” Boston University, Boston, MA [Online]. Available: <http://www.cs.bu.edu/BRITE/>
- [13] N. Spring, R. Mahajan, and T. Anderson, “The causes of path inflation,” in *Proc. ACM SIGCOMM*, Karlsruhe, Germany, Aug. 2003, pp. 113–124.
- [14] S. Dharmapurikar, P. Krishnamurthy, T. Sproull, and J. Lockwood, “Deep packet inspection using parallel bloom filters,” in *IEEE Mic*, 2003.
- [15] Y. Ohara, S. Imahori, and R. V Meter, “MARA: Maximum alternative routing algorithm,” in F. Yu, R. H. Katz, and T. V. Lakshman, “Gigabit rate packet pattern-matching using TCAM,” in *Proc. IEEE ICNP*, Berlin, Germany, Oct. 2004, pp. 174–183.
- [16] S. Dharmapurikar, P. Krishnamurthy, T. Sproull, and J. Lockwood, “Deep packet inspection using parallel bloom filters,” *IEEE Micro*, vol. 24, no. 1, pp. 44–51, Jan.–Feb. 2004.
- [17] “The Internet2 network,” 2009 [Online]. Available: <http://www.in-ternet2.edu>
- [18] K. C. Claffy, G. C. Polyzos, and H.-W. Braun, “Application

of sam-pling methodologies to network traffic characterization,” in *Proc. ACM SIGCOMM*, San Francisco, CA, Sep. 1993, pp. 194–203.

- [19] B.-Y. Choi and S. Bhattacharyya, “On the accuracy and overhead of Cisco sampled netflow,” in *Proc. ACM SIGMETRICS LSNI*, Banff, Canada, Jun. 2005, pp. 18–23.
- [20] N. Duffield and C. Lund, “Predicting resource usage and estimation accuracy in an IP flow measurement collection infrastructure,” in *Proc. ACM SIGCOMM IMC*, Miami Beach, FL, 2003, pp. 179–191.
- [21] C. Estan, K. Keys, D. Moore, and G. Varghese, “Building a better Netflow,” in *Proc. ACM SIGCOMM*, Portland, OR, Sep. 2004, pp. 245–256.
- [22] C. Estan and G. Varghese, “New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice,” *Trans. Comput. Syst.*, vol. 21, no. 3, pp. 270–313, 2003.
- [23] N. Hohn and D. Veitch, “Inverting sampled traffic,” in *Proc. ACM SIG-COMM IMC*, Miami, FL, Oct. 2003, pp. 222–233.
- [24] R. Kompella and C. Estan, “The power of slicing in Internet flow mea-surement,” in *Proc. ACM SIGCOMM IMC*, Berkeley, CA, Oct. 2005, p. 9.
- [25] A. Kumar, M. Sung, J. Xu, and J. Wang, “Data streaming algorithms for efficient and accurate estimation of flow size distribution,” in *Proc. ACM SIGMETRICS*, New York, NY, Jun. 2004, pp. 177–188.
- [26] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, “An information-the-oretic approach to traffic matrix estimation,” in *Proc. ACM SIGCOMM*, Karlsruhe, Germany, Aug. 2003, pp. 301–312.
- [27] X. Li, F. Bian, M. Crovella, C. Diot, R. Govindan, G. Iannaccone, and A. Lakhina, “Detection and identification of network anomalies using sketch subspaces,” in *Proc. ACM SIGCOMM IMC*, Rio de Janeiro, Brazil, Oct. 2006, pp. 147–152.
- [28] J. Mai, C.-N. Chuah, A. Sridharan, T. Ye, and H. Zang, “Is sampled data sufficient for anomaly detection?,” in *Proc. ACM SIGCOMM IMC*, Rio de Janeiro, Brazil, Oct. 2006, pp. 165–176.
- [29] A. Ramachandran, S. Seetharaman, N. Feamster, and V. Vazirani, “Fast monitoring of traffic subpopulations,” in *Proc. ACM SIGCOMM IMC*, Vouliagmeni, Greece, Oct. 2008, pp. 257–270.
- [30] H. V. Madhyastha and B. Krishnamurthy, “A generic language for ap-plication-specific flow sampling,” *Comput. Commun. Rev.*, vol. 38, no. 2, pp. 5–16, 2008.
- [31] L. Yuan, C.-N. Chuah, and P. Mohapatra, “ProgME: Towards pro-grammable network measurement,” in *Proc. ACM SIGCOMM*, Kyoto, Japan, Aug. 2007, pp. 97–108.
- [32] M. R. Sharma and J. W. Byers, “Scalable coordination techniques for distributed network monitoring,” in *Proc. PAM*, Boston, MA, Apr. 2005, pp. 349–352.
- [33] V. Sekar, M. K. Reiter, W. Willinger, H. Zhang, R. R. Kompella, and D. G. Andersen, “CSAMP: A system for network-wide flow monitoring,” in *Proc. USENIX NSDI*, San Francisco, CA, Apr. 2008, pp. 233–246.
- [34] G. Huang, S. Raza, S. Seetharaman, and C.-N. Chuah, “Dynamic mea-surement-aware routing in practice,” *IEEE Network*, vol. 25, no. 3, pp. 29–34, May-Jun. 2011.

- [35] J. R. D. Xu and M. Chiang, “Link-state routing with hop-by-hop for-warding can achieve optimal traffic engineering,” in *Proc.*

*IEEEIN-FOCOM*,2008.

## ABOUT THE AUTHORS



**B. Mahesh**, received his M.Tech in CSE ,from JNTUA, Anantapur, India.He attended 2 International conferences and 1 National conference.Present he

is working as an Assoc.Professor at KVSRTI, Kurnool, India. He is interested in Network Security and Cloud Computing.



**P. Madhu Krishna**, received his B.Tech degree in Computer Science and Engineering from INTEL Engineering college, Anantapur, India, in 2011. Currently pursuing M.Tech in computer science and Engineering at KVSRI Institute of Technology, Kurnool, India.