

Optimal NetFlow Deployment in IP Networks

Hui Zang¹ and Antonio Nucci^{2*}

¹ Sprint Advanced Technology Laboratories

1 Adrian Court, Burlingame, CA, USA

hzang@sprintlabs.com

² Narus, Inc.

500 Logue Avenue, Mountain View, CA, USA

anucci@narus.com

Abstract. This paper investigates the problem of deploying NetFlow with optimized coverage and cost in an IP network. Deploying a network-wide monitoring infrastructure in operational networks is necessary for practical reasons and Cisco NetFlow is a promising solution. However, several cost factors are associated with enabling NetFlow given the current conditions in such a network. We argue that enabling NetFlow to cover a major portion of traffic instead of the entire traffic will achieve significant cost savings while at the same time give operators enough monitoring capabilities. Therefore we aim to solve the Optimal NetFlow Location Problem (ONLP) for a given coverage ratio. We analyze various cost factors to enabling NetFlow in such a network. We model the problem as an Integer Linear Program (ILP). Although we are able to obtain optimal solutions for Sprint's North America Network by solving the ILP, two heuristic algorithms, Max-Plus (MP) and Least-Minus (LM), are developed to cope with larger-sized problems, given the NP-hard nature of the ONLP problem. The performance of the ILP and heuristics is demonstrated by numerical results and the LM heuristic is able to achieve sub-optimal solutions within 1~2% difference from the optimal solutions in a mixed router environment. It is observed that we can achieve 55% cost savings by covering 95% instead of 100% of the network traffic. The problem and the proposed methodology can be generalized to optimal deployment of new services and features in any types of networks.

Keywords: NetFlow, Integer Linear Programming, Optimal Placement

1 Introduction

Operating a large IP network without a detailed, network-wide knowledge of the traffic demands is challenging. An accurate view of the traffic demands is crucial for a number of important tasks, such as failure diagnosis, capacity planning and forecasting, routing and load-balancing policy optimization, attack identification, etc. It is obvious to operators now that network monitoring and traffic measurement is a *necessity* and Cisco's

* This work was done while Antonio Nucci was at Sprint Advanced Technology Laboratories.

NetFlow [1] emerges as a viable solution to this problem. NetFlow has received attention from both industry and academic researchers. For example, NetFlow data has been used to examine the accuracy of traffic matrix estimation techniques [2]. The prior work on NetFlow has been focusing on performance issues. Reference [3] compares NetFlow to SNMP and packet-level data collection, while [4] proposes new sampling techniques that can be used by NetFlow. In this paper, we study the issues in the deployment of NetFlow.

NetFlow is a set of features available on Cisco routers and other switching devices that provide network operators with access to IP flow information from their data networks [1]. The NetFlow infrastructure consists of two main components: NetFlow Data Export (NDE) and NetFlow Collector (NFC). The NDE is a module configured on routers and captures each IP flow traversing a router.³ When a timer expires or the NetFlow cache becomes full, IP flow statistics, such as number of IP flows, number of packets and bytes associated to each flow, source/destination AS numbers, source/destination prefix masks, etc, are exported to a NFC as UDP packets.

IP networks usually contain a large diversity of routers and not all interfaces on all routers can run NetFlow. Although NetFlow can be configured at per-interface basis, NetFlow-supporting capability is determined by the linecard and the router. There are three types of linecards in terms of NetFlow support: 1) linecards that support NetFlow in most traffic conditions, 2) linecards that do not support NetFlow, and 3) linecards that support NetFlow in only certain (light) traffic conditions. Care must be taken for type 3) linecards since turning on NetFlow could potentially impact linecard's performance on packet forwarding, i.e. cause losses and large latency, or generate inaccurate flow statistics. Linecards of types 2) and 3) can usually be upgraded to a newer configuration to support NetFlow.

Enabling NetFlow at specific router interfaces is not enough. The IP flow statistics exported by NDE modules at each router must be collected by NFCs. Operators process all the data stored in NFCs to gather the information they need.⁴ There are two problems when NFCs are considered. First, only a limited number of routers can be served by the same NFC. Second, carriers prefer to collocate NFCs with the NDEs that they serve to avoid the flooding of large amount of information over long-haul IP links.

Therefore, in order to enable NetFlow and utilize the data properly, operators need to identify: 1) a proper configuration for each router enabled to support NetFlow (NDE); and 2) a proper location for each NetFlow Collector (NFC).

The goal of this paper is to provide a methodology and a set of recommendations to optimizing the NetFlow deployment process. More precisely we are interested in identifying which routers and which linecards on routers should be NetFlow-enabled, such to cover a major portion of network traffic while minimizing the total capital investment required. We refer to the problem of covering a given fraction of traffic on the selected routers while minimizing the total cost as the *Optimal NetFlow Location Problem (ONLP)*. The solution to this problem will assist an operator in two situations: i) For an operator who has decided to deploy NetFlow, identify the proper locations of routers to enable NetFlow

³ An IP flow is identified as the combination of seven fields as Source and Destination IP addresses, Source and Destination Port numbers, IP protocol type, ToS bytes and Input Logical Interface (ifIndex).

⁴ NetFlow Data Analyzer (NDA) is a NetFlow-specific traffic analysis tool that enables the users/operators to retrieve, display and analyze NetFlow data collected from several NFC modules.

is covered by NetFlow:

$$\sum_{i:R_i \in \mathcal{R}'} \sum_{j:I_j^{R_i} \in \mathcal{I}^{R_i}} t_{i,j} \geq D \times \sum_{1 \leq i \leq N} \sum_{1 \leq j \leq S_i} t_{i,j},$$

while at the same time, minimizing

$$\sum_{i:R_i \in \mathcal{R}'} F(R_i, \mathcal{I}^{R_i}) + \sum_{1 \leq j \leq L} C(|\mathcal{R}' \cap \mathcal{P}_j|),$$

where $|\cdot|$ denotes the cardinality of a set.

We formulate the *Optimal NetFlow Location Problem* (ONLP) as an Integer Linear Program (ILP). Different constraints may be applied to different routers. We consider Cisco GSR routers [5] and 7500 routers [6] in this exercise. In [7], we discuss the details of their capability in supporting NetFlow and we also set up a testbed to study the impact of NetFlow on 7500 routers and determine the need and cost of upgrading a 7500 linecard. Although totally different methods are applied to obtain cost figures for both families of routers, from the modeling perspective, the main differences between both families of routers are the following. First, when upgrading a 7500 linecard, only the processor and memory are upgraded and the interfaces on the linecard remain unchanged, while the entire linecard is replaced when upgrading a GSR linecard which implies that the number of interfaces on the linecard may change with the upgrade. Second, a router consists of a Route Switch Processor (RSP) and a number of linecards. When upgrading a 7500 router's linecards, sometimes we need to upgrade the RSP on this router as well. However, when upgrading a GSR router's linecards, we do not need to upgrade the GSR's RSP because most of the processing is done by the linecards. These differences will be reflected in the ILP formulation.

2.1 Notation

Let \mathcal{G}_{7500} and \mathcal{G}_{GSR} be the set of all 7500 and GSR routers, respectively. Let $\mathcal{P} = \{1, 2, \dots, L\}$ be the set of all PoPs in the network and \mathcal{P}_i represent the set of routers belonging to PoP i . A router is present in one and only one PoP. For a router g , let $\mathcal{S}(g)$ be the set of slots on router g , whose cardinality is denoted by $|\mathcal{S}(g)|$. Let $t(g, s)$ be the traffic processed at slot s on router g . We define the specific notations for 7500 routers, GSRs, collectors and traffic coverage respectively.

7500 routers Let $c(g)$ be the *minimal* cost to upgrade the current configuration of router g to one that supports NetFlow. $c(g) = 0$ if the current one supports NetFlow. Binary parameter $r(g) = 1$ if such an upgrade is available, and $r(g) = 0$ otherwise. Let $c(g, s)$ be the *minimal* cost to upgrade the current configuration at slot s , router g to one that supports NetFlow. $c(g, s) = 0$ if the current configuration supports NetFlow. Binary parameter $r(g, s) = 1$ if such an upgrade is available, and $r(g, s) = 0$ otherwise.

GSR routers Let \mathcal{T} be the set of all linecard types present on the routers in \mathcal{G}_{GSR} . For each $g \in \mathcal{G}_{GSR}$ we define $\mathcal{T}(g)$ as the set of linecard types present on router g . Each

Constraint (1) links the variables γ associated to each router with variables η associated to each slot. The left inequality in (1) forces $\gamma(g)$ to be 0 if none of its slots has been selected to run NetFlow. The right inequality in (1) forces $\gamma(g)$ to be 1 if one or more of its slots have been selected to run NetFlow.

- Relationship between $r(g)$ and $\gamma(g)$, and $r(g, s)$ and $\eta(g, s)$:

$$r(g) \geq \gamma(g) \quad \forall g \in \mathcal{G}_{7500} \quad (2)$$

$$r(g, s) \geq \eta(g, s) \quad \forall g \in \mathcal{G}_{7500} \cup \mathcal{G}_{GSR}, \forall s \in \mathcal{S}(g) \quad (3)$$

Constraints (2) and (3) guarantee that a router/slot can be selected to have NetFlow enabled only if its current configuration supports NetFlow or it can be upgraded to another configuration that supports NetFlow.

- Number of interfaces on GSR routers:

$$a_g(t)\nu_g(t) \geq \sum_{s \in \mathcal{V}_g(t)} \eta(g, s)p_{g,s}(t) \quad \forall g \in \mathcal{G}_{GSR}, \forall t \in \mathcal{T}(g) \quad (4)$$

Constraint (4) guarantees that we invest in the minimum number of linecards necessary according to the selection we made. For example, if router g has two linecards of type t with one port being used on each, and the upgraded version of linecard type t has four ports available, then Constraint (4) implies that only one upgraded version of linecard type t is necessary, i.e. $\nu_g(t) \geq 1$. When the total cost is minimized by the objective function, $\nu_g(t)$ will be forced to be 1.

- Fraction of the total traffic to be covered by enabling NetFlow on specific routers and slots:

$$\sum_{g \in \mathcal{G}_{7500}} \sum_{s \in \mathcal{S}(g)} t(g, s)\eta(g, s) + \sum_{g \in \mathcal{G}_{GSR}} \sum_{s \in \mathcal{S}(g)} t(g, s)\eta(g, s) \geq D \times \sum_{g \in \mathcal{G}_{7500} \cup \mathcal{G}_{GSR}} \sum_{s \in \mathcal{S}(g)} t(g, s) \quad (5)$$

Constraint (5) ensures that the final solution selected must cover at least a D fraction of the total traffic. It is clear that the larger D is, the larger will be the number of slots enabled to support NetFlow and the associated deployment cost.

- The number of collectors needed per PoP:

$$N \times NC_i \geq \sum_{g \in \mathcal{P}_i} \gamma(g) \geq NC_i \quad \forall i \in \mathcal{P} \quad (6)$$

Constraint (6) ensures that for any PoP, if there are routers with NetFlow enabled, the number of collectors in this PoP will be sufficient to cover all these routers. At the same time, no collectors should be placed at any given PoP where no router is enabled with NetFlow.

Algorithm 1 Heuristic I - Max-Plus (MP)

1.0 Initialize $T_{covered} = 0$, and $C_{total} = 0$. Set

$$T_{remaining} = T_{total} \times D - T_{covered} \quad (7)$$

1.1 Examine all slots without NetFlow enabled. For each slot s on router g at PoP p , calculate $C_{collector}(g, s)$, as the additional collector cost at PoP p if slot s were to be selected to enable NetFlow.

$$C_{collector}(g, s) = \begin{cases} 0 & \text{if router } g \text{ has NetFlow on, or if collectors} \\ & \text{at PoP } p \text{ can support one more router} \\ C & \text{otherwise} \end{cases} \quad (8)$$

$$CostPerBit(g, s) = (c(g, s) + C_{collector}(g, s)) / \text{Min}(t(g, s), T_{remaining}) \quad (9)$$

1.2 Enable NetFlow on slot s at router g with the smallest $CostPerBit(g, s)$. Set $T_{covered} = T_{covered} + t(g, s)$, and $C_{total} = C_{total} + c(g, s) + C_{collector}(g, s)$. Update $T_{remaining}$ by Eqn. (7).

1.3 Repeat Steps 1.1 through 1.2 until $T_{remaining} \leq 0$ and return.

4 Numerical Results

In this section, we present numerical results obtained by applying the ILP model and the heuristics on Sprint's North America IP backbone network (SNAIB-NET) with real traffic. We consider traffic carried on all links between *gateway (GW) routers* and *backbone (BB) routers*. We choose to enable NetFlow on gateway routers because it is more cost-effective to upgrade gateway routers than backbone routers as we found out by going through the router configurations [7].

4.1 Platform and Speed

We solve the ILP models using CPLEX [9] running on a 2.4 GHz Xeon processor with 1 GB RAM space. The time it takes to solve the ILP models for SNAIB-NET gateway routers ranges from a few seconds to 30 minutes. Note that we solved for several hundred of routers which is a subset of SNAIB-NET. Therefore, for networks of sizes less than hundreds of routers, it is feasible to use the ILP model to find an optimal solution for ONLP. The heuristics runs much faster - it takes sub-seconds to seconds for each heuristic to solve the problem for all coverage ratios.

4.2 The ILP Model and the Heuristics

In this subsection, we present the solutions from the ILP and two heuristics and compare the performance achieved by the heuristics with the optimal solution obtained from the ILP model. Figure 1(a) shows the normalized cost obtained from solving the ILP model and the two heuristics to achieve different coverage ratios from 50% to 100%. The costs are normalized by the cost required to provide 100% coverage, which is the same from all three methods. We notice that the cost to achieve 95% coverage is only about 45% of the cost that is required for 100% coverage. In Fig. 1(b), we plot the relative difference, i.e., the cost difference normalized by the optimal value between the results obtained by each heuristic and those obtained by solving the ILP. We can see that the two heuristics perform differently in terms of optimality. LM performs significantly better than MP. At 50% coverage, the solution from MP is 7% higher than the optimal solution while the

5 Conclusions

In this paper, we studied the optimization problem for NetFlow deployment in an IP network. Specifically, we considered a partial NetFlow deployment to achieve the lowest cost for a given coverage ratio, which is the Optimal NetFlow Location Problem (ONLP).

We developed an ILP model and two heuristic algorithms to select routers and slots to support NetFlow and the associated configurations such that a certain amount of network traffic is covered at a minimum cost.

We solved ONLP for Sprint's IP backbone network in north America. We presented numerical results from applying the ILP model and two heuristics. We demonstrated that, it is possible to achieve significant cost savings by adopting a partial NetFlow deployment strategy, i.e., to cover a major portion of the network traffic instead of the entire traffic. A good coverage ratio is suggested as 95%, with 55% cost reduction.

Although our discussion was focused on Cisco NetFlow and the results were collected from Sprint's operational IP backbone network only, the results can be referenced in similar practices and the methodology proposed can be extended and applied to a wide variety of network location problems to enable different features and services. Besides NetFlow from Cisco, other vendors also support similar flow-based monitoring services, such as sFlow [10], and our methodology can be applied to the deployment of sFlow as well. In addition, as ongoing work, we are extending our approach to network monitoring functions of finer granularity such as packet trace collection.

Acknowledgment

We thank Travis Dawson and Beng-Ong Lee at Sprint ATL for their support in the 7500 router testing and answers to our various NetFlow-related questions.

References

1. "NetFlow Services Solutions Guide," Cisco white paper.
2. A. Soule, A. Nucci, R. Cruz, E. Leonardi and N. Taft, "How to Identify and Estimate the Largest Traffic Elements in a Dynamic Environment", *Proceedings of ACM Sigmetrics*, New York, USA, July 2004.
3. R. Sommer and A. Feldmann, "NetFlow: Information Loss or Win?" *Proceedings of Internet Measurement Workshop*, Marseille, France, Nov. 2002.
4. C. Estan, K. Keys, D. Moore, and G. Varghese, "Building a Better NetFlow," *Proceedings of ACM Sigcomm*, Portland, OR, USA, August 2004.
5. "Cisco 12000 Series Router," Cisco white paper.
6. "Cisco 7500 Series Router," Cisco white paper.
7. H. Zang, and A. Nucci, "Optimal NetFlow Deployment in IP Networks," *Sprint ATL Research Report RR05-ATL-061624*, June 2005.
8. M. R. Garey and D. S. Johnson, "Computers and Intractability, A Guide to the Theory of NP-Completeness," Bell Telephone Laboratories, Inc., 1979.
9. <http://www.ilog.com/products/cplex>.
10. <http://www.sflow.org>.