Dynamic Traffic Management mechanism for active optimization of ISP costs

Zbigniew Duliński
Jagiellonian University
Kraków, Poland
Email: dulinski@th.if.uj.edu.pl

Rafał Stankiewicz
AGH University of Science and Technology
Kraków, Poland
Email: rafal.stankiewicz@agh.edu.pl

***Abstract*—This paper proposes a dynamic traffic management mechanism which addresses the problem of cost-optimal distribution of inevitable inter-domain traffic on multiple links, and enables dynamic overlay traffic management and adaptation to constantly changing traffic patterns to meet cost optimization goals. The mechanism allows for the management of traffic generated by overlay applications such as content delivery platforms, online storage systems, or P2P.**

I. Introduction

A very large share of Internet traffic is currently generated by overlay applications. A few years ago P2P traffic dominated. Nowadays, content delivery platforms and online storage systems generate the majority of traffic [1] but P2P traffic is still important. Since P2P traffic became significant, a lot of attention has been devoted to the optimization of overlay traffic distribution on the Internet. Various overlay traffic management solutions have been proposed. They are classified in [2]. One of the most important issues form the ISPs' point of view is the minimization of inter-domain traffic and its related cost. Proposed solutions focus mainly on traffic localization, which is achieved by influencing on overlay application behavior in such a way that it exchanges more traffic within a single autonomous system instead of transferring it between remote domains [3], [4]. Some work has also been devoted to the management of inter-domain traffic by the appropriate selection of remote content source, e.g., [5]. As a result, the traffic may be transferred through cheaper links. It is important to note that all the aforementioned overlay traffic management solutions use static ISP policy, i.e., valid for a relatively long timescale. Development of new overlay traffic management for cloud-based applications is currently a research subject, e.g., in the SmartenIT project [6].

In the current paper, we present a dynamic traffic management (DTM) mechanism. It is innovative in at least two ways: (1) it addresses the problem of cost-optimal distribution of inevitable inter-domain traffic on multiple links, and (2) it enables dynamic overlay traffic management and adaptation to constantly changing traffic patterns to meet cost optimization goals. The mechanism may react on scales as short as minutes. The ISP policy may change frequently adapting to the current situation. One of the basic assumptions for the mechanism is that the desired content available in content delivery platforms, online storage systems, or P2P applications has multiple replicas in the global network and can be downloaded from one of them. An ISP can influence overlay applications' decision on the selection of the source so that the related traffic will be transferred by a link preferred by the ISP.

II. General concept of DTM

DTM is a generic concept addressing the problem of minimization of the ISP's costs related to the amount of traffic in the network. The cost may reflect a monetary cost for inter-domain traffic, cost of energy consumption by network nodes or devices, or other monetary or non-monetary costs related to the amount of traffic (e.g., utilization of resources, congestion, etc.). The dynamic management performs optimization to find a best solution over criteria and conditions defined by the ISP. It is able to react to changing conditions on a short time scale. This paper discuss practical implementation of DTM concept enabling optimization of cost of inter-domain traffic generated by overlay applications.

III. DTM applicability for contemporary network scenarios

We assume that the ISP operates an autonomous system (AS) with at least two inter-domain links (i.e., has a multi-homed AS). ISP customers (either individual or institutional) use an overlay application which distributes content. The same content may be provided by multiple sources: servers, peers (in P2P application), clouds, data centers (DC), etc., located in distinct autonomous systems. The routing paths from various content sources to the destination in ISP's domain may be different and pass through different transit AS-es and inter-domain links. An ISP's customer application may obtain the desired content via different inter-domain links. The cost related to the download of the content may depend on the inter-domain link used.

DTM mechanism is assumed to run on a centralized ISP-owned server. It is deployed in the ISP domain and is under full control of the ISP. In general, this server is conceptually a kind of *Oracle* server [7]. For example, it might be an implementation of an ALTO server [4]. The server is responsible for gathering traffic information from the local ISP domain, exchanging some information with similar servers in other domains operated by trusted ISPs (e.g., using an inter-ALTO protocol), and providing guidelines to overlay applications.

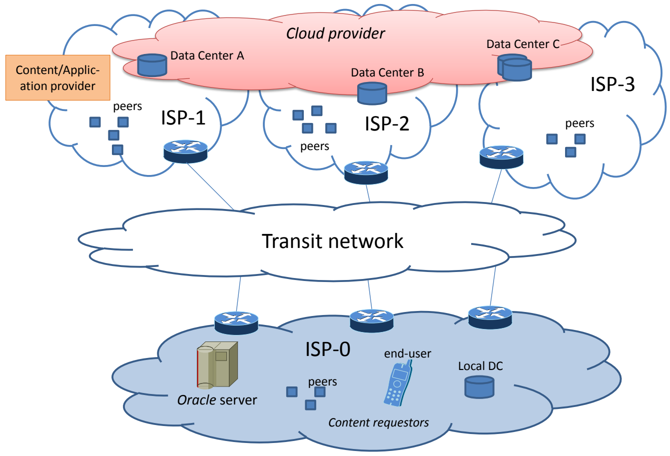

Fig. 1. Network model for the use-cases description

Traffic on each inter-domain link is classified as *manageable* or *non-manageable*. The latter cannot be influenced by DTM. Moreover, even the cooperative application follows the ISP's suggestion with some probability, referred to as a *success rate*. To sum up, the capability of DTM to influence the total cost of inter-domain traffic depends on two factors: the share of manageable traffic in the total traffic, and the success rate.

In the rest of this section we describe use-cases for the DTM mechanism. Figure 1 is an illustration of the use-cases. ISP-0 deployed an *Oracle* with the DTM mechanism. This is concerned with influencing the selection of the source of content to be downloaded by local peers, end-users, or DCs.

A. Peer-to-peer use-case

In this use-case we consider a P2P application. All, or parts of, the desired content are possessed by peers in the ISP-1, ISP-2 and ISP-3 domains. The peer in the ISP-0 domain needs the content. The communication follows this sequence:

1) The peer in the ISP-0 domain sends the list of potential communication partners to the *Oracle* server;
2) The *Oracle* sorts the peer list using DTM;
3) The *Oracle* sends the sorted list to the requesting peer;
4) The peer connects with peers in other domains in accordance with the received sorted peer list;
5) The downstream traffic goes via links preferred by ISP-0.

B. Content download from data centers

In this use-case we assume a cloud based service. The required content may be provided from a few DCs. The routing paths from various DCs to the end-user device may cross different inter-domain links. The end-user requests content from a content/application provider and has no knowledge of the multiple sources of the content. Thus, if the ISP wants to influence the decision on content source selection, it should cooperate with one of the parties involved in content delivery, e.g., the content/application provider or a cloud provider (DCs owner). The communication sequence may be as follows:

1) The end-user localized in the ISP-0 domain sends a request for the selected content to the content provider;

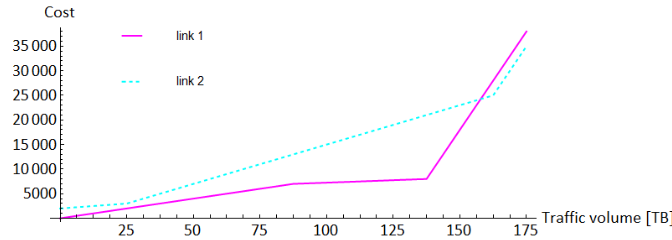

Fig. 2. Cost functions used for accounting cost of inter-domain traffic

2) The content/application provider redirects the request to the cloud supervisor which knows potential content localization (selected DCs);
3) The supervisor contacts with the *Oracle* server in the ISP-0 domain. The list of available DCs is sent;
4) The *Oracle* sorts the list using DTM;
5) The *Oracle* sends the list back to the cloud supervisor;
6) The cloud supervisor begins content provisioning (the end-user establishes a connection with a selected DC);
7) The related traffic traverses selected inter-domain links chosen by the dynamic traffic management procedure.

IV. PRACTICAL IMPLEMENTATION OF DTM

We have designed a practical implementation of the DTM mechanism. It consists of the following main components explained in the example:

- optimization algorithm used to find the optimal solution;
- compensation algorithm constantly influencing the traffic distribution to achieve the optimal solution at the end of the accounting period;
- interface to the application which translates traffic compensation decisions into application-specific recommendations for traffic management.

The tariff used on inter-domains link may use different cost functions. In general, the cost function is defined as $c_i = f_i(B_i)$ where c_i is a cost calculated for link i, $f_i()$ is a cost function used on that link, and B_i is measured traffic at link i. The cost function might be non-linear. It may be based on traffic volume or 95$^{\text{th}}$ percentile. The total cost of inter-domain traffic is calculated as $C = \sum_{i=1}^{N^v} c_i^v + \sum_{i=1}^{N^p} c_i^p$ where N^v and N^p are the number of links charged according to traffic volume and percentile rules, respectively. Costs calculated on respective links are denoted as c_i^v and c_i^p.

For simplicity's sake, we assume that the AS has two inter-domain links, both charged for the traffic volume. Cost functions used on those links are presented in Fig. 2.

A. The optimization algorithm

The optimization algorithm finds a best solution to be achieved at the end of an accounting period. Using cost functions and information from the previous period, it makes a prediction for the next one and tries to find a traffic distribution in which the cost is minimized. DTM strives to influence the traffic distribution in such a way that the cost at the end of the current accounting period will be close to the optimal solution.

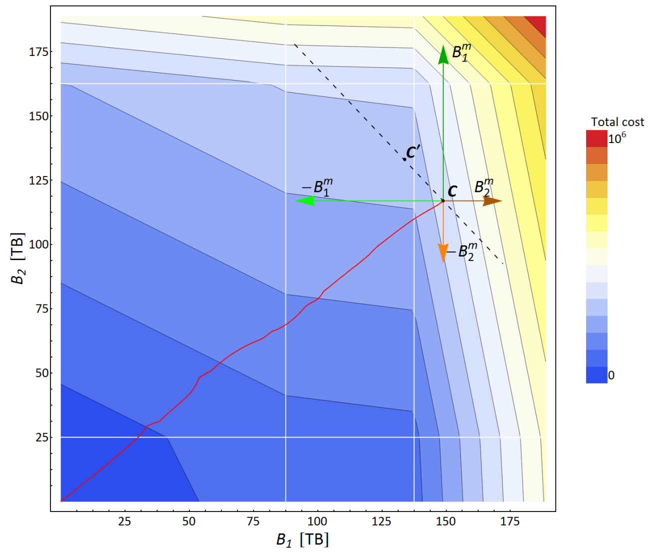

Fig. 3. A cost map as a function of traffic volume on both inter-domain links and cost optimization potential

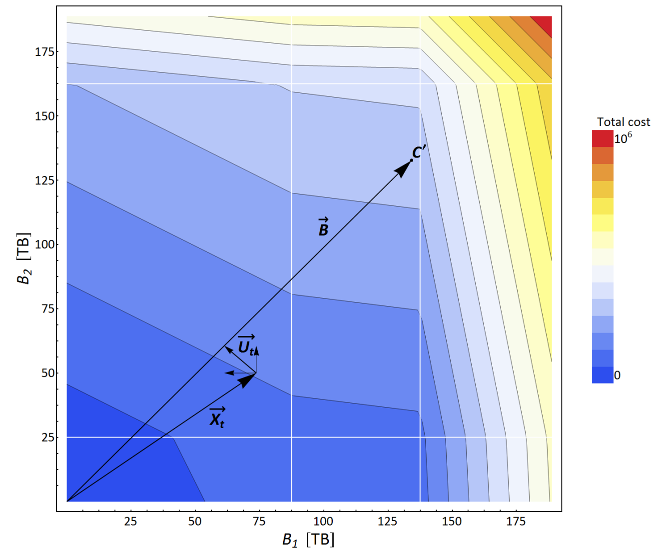

Fig. 4. Illustration of a traffic compensation mechanism

The plot in Figure 3 shows a cost map related to traffic volumes at both inter-domain links. The X-axis and Y-axis represent traffic volumes B_1 and B_2, respectively. The contour plot shows total monetary cost calculated for B_1 and B_2 traffic using respective cost functions. Point C shows traffic volumes on both links as well as represents a total cost for inter-domain traffic paid in the previous accounting period. The red curve shows how the cumulative traffic volume grew on both links during the previous accounting period. The manageable traffic at link 1 is denoted as B_1^m. If all this traffic were downloaded from sources available via link 2 instead of link 1, the traffic volume on link 1 could be decreased by B_1^m while the traffic volume at link 2 would be increased by the same value (the green arrows in Fig. 3). In turn, B_2^m denotes manageable traffic on link 2. In theory, all this manageable traffic could be shifted from link 2 to link 1 (the orange arrows). The amounts of manageable traffic determine the boundaries for traffic management. Therefore, the optimal solution can be found at a dashed line presented in Fig. 3. The optimization algorithm finds a solution denoted as C', e.g., a minimal cost laying on the line. More sophisticated algorithms are possible.

The optimization algorithm is more complicated if more inter-domain links are available. For N links the cost function is defined over N-dimensional space.

B. The compensation algorithm

Traffic management is performed during the accounting period. First, the optimal cost C' is translated into desired traffic volume at both links B_1' and B_2'. Those values define the so called 'reference vector' $\overrightarrow{B}$ which determines the desired direction of the traffic growth on links over time (Fig. 4).

The whole accounting period is divided into several intervals. During an interval, the traffic volume on links is observed. On the basis of these measurements, a new prediction of the cost which will be achieved at the end of period is calculated. If the predicted cost differs from the optimal solution an intervention is needed to compensate for the undesired trend. At the end of the interval a current traffic vector $\overrightarrow{X_t}$ is calculated, where t denotes a current time. The direction of this vector is compared to the reference vector. The $\overrightarrow{U_t}$ vector represents a difference between measured and expected traffic at time t. This is called a 'compensation vector'. It is found by a compensation algorithm which, in general, takes into account traffic and cost estimation on several links. The compensation vector reflects the amount of traffic that needs to be shifted from one link to the other and is translated to recommendations for overlay applications on the content source selection. The traffic compensation is performed during the interval using the compensation vector determined in the last interval. At the end of the interval, the compensation vector is recalculated.

V. Experimental results

The DTM mechanism has been implemented in a proprietary simulator. We have performed a number of experiments to verify the DTM mechanism and evaluate its performance. We present a simple, but intuitive experiment to picture DTM and its capabilities. The setup of the presented experiment was used in Section IV as an example to describe the optimization algorithm, compensation algorithm. Therefore, we assume an AS with two inter-domain links charged on the basis of traffic volume. Cost functions and the cost map looks exactly the same as in the example (Figures 2 and 3).

The source of a manageable overlay traffic for the experiment was BitTorrent – a file-sharing P2P application. Peers do not use a random list of peers provided by a tracker, but a list ranked by an *Oracle*. The peer ranking is biased by DTM. For instance, if the compensation vector calculated in some interval indicates that the traffic on link 1 should be decreased, the *Oracle* updates the peer ranking rule accordingly to the length of compensation vector. The new ranking is used in the following interval and gives higher preference to the remote peers that can feed local peers via link 2. We assume a success

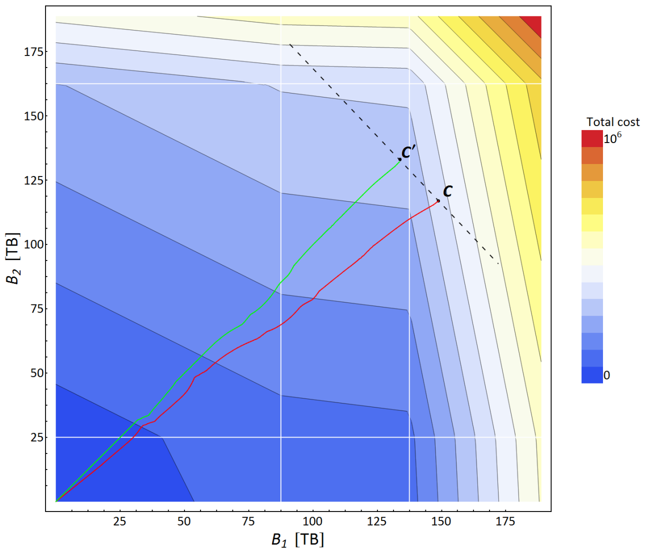

Fig. 5. Comparison of traffic growth on links during the accounting period with (green curve) and without (red) DTM

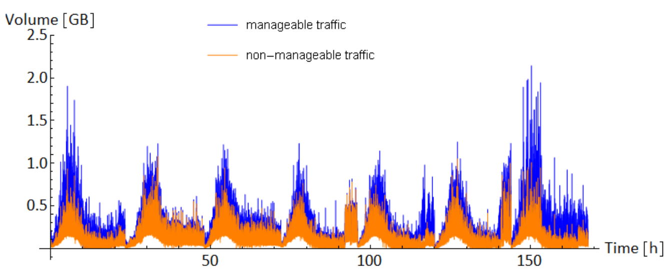

Fig. 6. Traffic pattern in link 2 without DTM mechanism

rate of 30%, since some local peers may be uncooperative or some remote peers may refuse connection with local peers. The algorithm changes peer rating dynamically every five minutes based on current measurements of the traffic. Traffic volume samples are collected every 30 seconds. The goal cost is represented by point C'. We have performed two simulations: with and without DTM. The assumed accounting period was one week.

The red curve in Figure 5 shows the growth of traffic volume without DTM. The final cost of the inter-domain traffic is 34,793 units (point C). The green curve shows the growth of traffic volume when the DTM is used. The total cost is very close to the goal cost (point C') and equals to 28,068 units. The cost is decreased by $\approx$19%. It can be easilly seen from Fig. 5 that the traffic on link 1 was decreased while

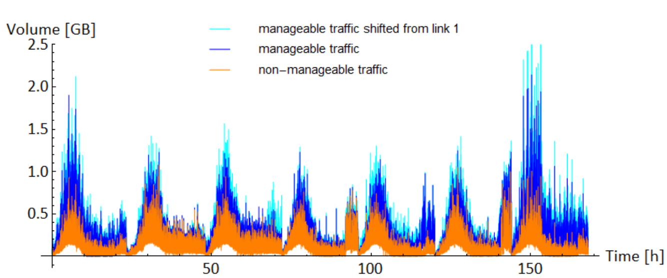

Fig. 7. Traffic pattern in link 2 with DTM mechanism

on link 2 it was increased. The compensation algorithm was able to influence the overlay in such a way that traffic was shifted from link 1 to link 2. This observation is confirmed by observation of traffic patters on these links. In Figures 6 and 7 we show traffic patterns on link 2 without and with DTM, respectively. The orange curve represent non-manageable traffic. The manageable traffic (the blue curve) is shown on top of the non-manageable. It can be noticed that the manageable traffic on link 2 is larger if DTM is applied.

VI. Conclusion

This paper presents a traffic management mechanism which enables management of application layer traffic. The DTM requires cooperation with an application. Only traffic generated by a cooperative application is manageable. DTM is dedicated to the management of overlay traffic in a dynamic way, based on constantly changing traffic and related cost prediction.

At the current stage of development, the DTM mechanism is able to influence an application generating downstream traffic incoming to the autonomous system via several inter-domain links. Cost functions on inter-domain links can be based on traffic volume or the 95th percentile. The DTM monitors traffic on links and reacts to constantly changing traffic patterns. We focus on minimization of cost of transferring an inevitable inter-domain traffic (which must be downloaded via inter-domain link since related content is not available locally).

The paper reports simulation results for a scenario with two inter-domain links, cost functions based on traffic volume, and P2P application. The results show that total transfer cost may be decreased if the DTM is applied even though the overall inter-domain traffic is not decreased.

Future work will be devoted to testing DTM with 95th percentile rule based cost functions. Cloud-based applications will also be involved.

Acknowledgements

This work was supported partially by the SmartenIT project funded by the EU FP7 Program under Contract No. FP7-2012-ICT-317846. It was also supported in part by the Polish National Science Centre under Grant N N517 555239.

References

[1] Cisco, “Cisco global cloud index: Forecast and methodology (2011-2016),” 2012.
[2] G. Dán, T. Hoßfeld, S. Oechsner, P. Chołda, R. Stankiewicz, I. Papafili, and G. D. Stamoulis, “Interaction Patterns between P2P Content Distribution Systems and ISPs,” vol. 42, no. 5, pp. 222–230, May 2011.
[3] I. Papafili, G. D. Stamoulis, R. Stankiewicz, S. Oechsner, K. Pussep, R. Wojcik, J. Domzal, D. Staehle, F. Lehrieder, and B. Stiller, “Assessment of economic management of overlay traffic: Methodology and results,” in *The Future Internet*, ser. LCNS, 2011, vol. 6656, pp. 121–131.
[4] J. Seedorf and E. W. Burger, “Application-layer traffic optimization (ALTO) problem statement,” IETF RFC 5693, Oct. 2009.
[5] Z. Duliński, R. Stankiewicz, P. Wydrych, M. Kantor, and P. Chołda, “Cost-driven peer rating algorithm,” in *IEEE International Conference on Communications (ICC)*, 2011, pp. 1–6.
[6] “SmartentIT project.” [Online]. Available: www.smartenit.org
[7] V. Aggarwal, A. Feldmann, and C. Scheideler, “Can ISPs and P2P users cooperate for improved performance?” *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 3, pp. 29–40, Jul. 2007.