Market-Based Cooperative Resource Allocation for Overlay Networks

Ryota Egashira Ariffin Datuk Yahaya Tatsuya Suda
School of Information and Computer Sciences
University of California, Irvine
Irvine, California, USA
{egashira, ariffin, suda}@ics.uci.edu

Abstract— **In recent years, the notion of service overlay networks has been proposed as a promising solution for providing end-to-end QoS without changing the current Internet architecture. A major issue in deploying service overlay networks is determining how to allocate resources (such as link bandwidth) on a substrate network to overlay networks, while satisfying the end-to-end QoS requirements of applications running on each overlay network. This paper introduces the Market-based Cooperative Resource Allocation (MaCRA) architecture that achieves fair and efficient resource allocation in a decentralized manner. In MaCRA, resources on a substrate network are priced, and each overlay network provider creates an overlay network on a minimum cost basis to meet its application QoS requirements. MaCRA also allows each overlay network provider to trade their current resources with other overlay network providers when resources on a substrate network are not available or expensive. Simulation results demonstrate that MaCRA achieves fairness and efficiency in allocating resources for overlay networks when compared to existing mechanisms.**

Keywords - Service Overlay Network, QoS, Network Resource Allocation, Market Mechanism, Network Resource Pricing, Network Resource Trade

I. INTRODUCTION

The Internet consists of multiple autonomous systems (ASs) that are managed independently by its respective Internet Service Providers (ISPs). End-to-end (E2E) Internet QoS requires the cooperative efforts of multiple ASs on a traffic flow path. Even if an ISP aims to provide E2E QoS for outgoing traffic flow and supports QoS within its own AS, there may be lack of incentives for the other ASs to support a similar level of QoS within their own ASs.

In recent years, researchers have proposed to construct an overlay network that supports QoS in the application layer while preserving the best-effort network layer. An overlay network has been a widely used technique to evaluate and implement various applications as well as new network protocols without any changes to existing infrastructure (e.g., QBone, MBone, Xbone). There is an increasing number of applications with E2E QoS requirements and such applications need a common set of functionalities (e.g., detection of network status, discovery of an overlay topology, fault tolerance, etc) to satisfy E2E QoS requirements on top of an overlay network. Therefore, instead of constructing a dedicated overlay network to support each application's QoS, it has been proposed to deploy a unified overlay network (i.e., service overlay network [1][2]) that works as a substrate to support various applications with E2E QoS requirements and provide common functionalities.

In deploying service overlay network infrastructure, one primary issue is how to allocate resources (i.e., link bandwidth) on a substrate network to each application with E2E QoS requirements. In order to accommodate as much application traffic demand as possible, it is necessary to optimize spatial and temporal allocation of limited resources. Applications may request for resources at random times and such requests need to be handled in a serial manner without any knowledge about future requests. It is thus regarded as an online problem [3] and very challenging.

There have been existing research efforts to address similar issues in several areas including routing over virtual circuits [4], designing a VPN [5] and allocating resources for virtual networks in a diversified Internet [6]. Most of the existing mechanisms rely on a centralized approach in which a server or a group of servers collect information about available resources on a substrate network and make resource allocation decisions for applications. However, a centralized approach has an inherent scalability limitation. Allocating network resources to applications with E2E QoS requirements is a computationally complex problem that can be formulated as a NP-complete problem [5]. Many research efforts have been dedicated to developing heuristic algorithms to obtain approximate solutions of this problem in a reasonable computation time [5]. However, as applications that request for resources on a substrate network increase in number and have additional QoS requirements to be satisfied, the problem becomes even more difficult and intractable. Therefore a decentralized approach needs to be designed and developed.

In developing a decentralized resource allocation approach, fairness among applications needs to be considered. In online resource allocation, there is lack of fairness since if two applications are requesting for the same resources, one application may successfully reserve resources while the other may fail. This occurs since successfully obtaining resources depends on precisely when

This work is supported by the NSF through grants ANI-0083074, ANI-9903427 and ANI-0508506, by DARPA through grant MDA972-99-1-0007, by AFOSR through grant MURI F49620-00-1-0330, and by grants from the UC MICRO and Discovery programs, Hitachi, Hitachi America, Hitachi CRL, Hitachi SDL, DENSO IT Laboratory, NEC, NEC America, NICT, NTT Docomo and Novell.

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

requests for resources are made and what resources are requested. To increase fairness, the chance of obtaining necessary resources should not significantly differ regardless of when an application requests for resources and what resources it requests. Efficiency of resource allocation also needs to be considered. In online resource allocation, optimal resource allocation for the current set of application traffic demands may result in sub-optimal resource allocation of a future set of application traffic demands. It would be desirable to be able to reallocate network resources already assigned to one application to another application to improve overall efficiency of resource allocation.

This paper introduces the Market-based Cooperative Resource Allocation (MaCRA) architecture as a solution to the issues mentioned above. MaCRA aims at satisfying two application QoS requirements: bandwidth guarantee and minimum latency. In MaCRA, an overlay network provider (ONP) deploys an overlay network on a substrate network to support its application traffic. (See Fig.1). A substrate network provides its resources (i.e., link bandwidth) to support an overlay network. A substrate network prices bandwidth of a link based on latency of the link. An ONP selects E2E paths (i.e, a set of links) that constitute an overlay network and reserves link bandwidth on the E2E paths. An ONP selects E2E paths that minimize total cost (i.e., meaning that E2E paths with minimum latency are selected) from several possible candidates. Note that a substrate network doesn't make any decision regarding what E2E paths to reserve. Each ONP is entirely responsible for such decisions.

In MaCRA, each ONP seeks link bandwidth that minimzes total cost in a greedy manner. However, ONPs being greedy without any cooperation does not lead to overall optimal resource allocation, which will be demonstrated in simulation studies later in this paper. In order to further optimize spatial and temporal resource allocation, MaCRA architecture enables trading of link bandwidth between ONPs. A trade allows an ONP to buy link bandwidth from other ONPs when link bandwdith on a substrate network is not available or expensive.

A trade of link bandwidth between ONPs takes place only when the trade benefits both ONPs (i.e., both a seller and a buyer of link bandwidth). A trade is not made if the trade is not beneficial (i.e., fair) for either a buyer or a seller ONP. A seller ONP, upon selling link bandwidth to another ONP, reserves other link bandwidth to keep supporting its application traffic demand. Trading link bandwidth thus enables reallocation of link bandwidth from an already deployed overlay network to another deployed overlay network. It could help improve efficiency in spatial and temporal allocation of link bandwidth.

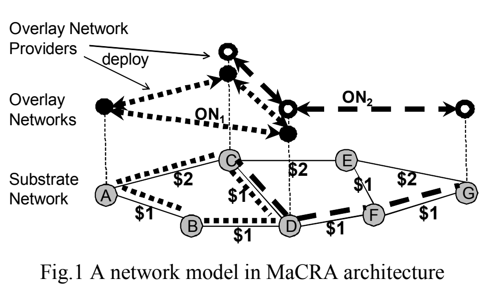

Fig.1 A network model in MaCRA architecture

In MaCRA architecture, each ONP performs computation to determine what link bandwidth to reserve to support its application traffic. The distributed nature of the proposed computing model enhances scalability in computing fair and efficient bandwidth allocation when the number of applications requesting for link bandwidth increases. The design of MaCRA can be extended such that ONPs consider additional QoS requirements (e.g., loss rate, delay jitter) in selecting resources to reserve.

II. Proposed Architecture

This section elaborates on the design of MaCRA architecture, the network model assumed in MaCRA architecture and how a substrate network and ONPs perform resource allocation.

A. Network model

MaCRA assumes a network model that contains a substrate network, overlay network providers (ONPs) and overlay networks as illustrated in Fig.1. The following describes functionalities of each component.

1) Substrate network

A substrate network consists of nodes and links. A node in a substrate network has functionalities to support routing of application traffic. A node keeps a dedicated routing table for traffic of each application (i.e., overlay network). A node is able to differentiate traffic from different applications and forward it to a proper next hop based on the dedicated routing table. A node updates a routing table in response to a request from an ONP. An ONP, upon identifying a set of links and the amount of bandwidth to reserve, requests nodes incident on these links to allocate requested bandwidth to the application and update their routing tables accordingly.

A node is also responsible for maintaining updated state of a whole substrate network. A node periodically exchanges link state with each other using an existing protocol (e.g., OSPF). Link state to be exchanged includes link connectivity, residual (unreserved) bandwidth on a link, maximum available bandwidth on a link and current price of bandwidth on a link.

A link on a substrate network provides bandwidth with guarantee. In service overlay network infrastructure, a link on a substrate network represents a logical (overlay) link that may consist of several underlying links. Bandwidth of a logical link is guaranteed using existing techniques [7] regardless of changes in the status of underlying links. It is, however, out of scope in this paper to discuss how to guarantee link bandwidth. Bandwidth on a link could be shared by multiple overlay networks (applications). In Fig.1, for example, bandwidth on link C-D is shared by two overlay networks, ON_1 and ON_2. Bandwidth on a link is priced based on latency of the link. There are other possible pricing schemes of link bandwidth on a substrate network, and these pricing schemes are compared and discussed in simulation studies later in this paper.

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

2) Overlay network provider

Overlay network providers (ONP) refer to individual users or corporations (e.g., ISPs) that need to satisfy QoS requirements for their application traffic. Examples of applications considered in this paper include video conferences between individuals, VPN between remote branches of companies, value-added routing service (such as RON [8]) and P2P video multicast applications. An ONP is assumed to know the following information regarding its traffic demand: 1) a group of nodes that constitute an overlay network to support application traffic with QoS requirements, 2) an amount of bandwidth required for application traffic and 3) time duration to support its traffic demand (i.e., reservation period of link bandwidth).

3) Overlay network

An overlay network is defined as a set of application traffic demands among nodes. In Fig.1, for example, overlay network ON_2 represents a set of application traffic demands between node C and D and between node D and G. An ONP deploys its overlay network through reservation of link bandwidth and update of a routing table at nodes on a substrate network. In order to deploy ON_2, in Fig.1, an ONP reserves bandwidth on link C-D, D-F and F-G. The ONP also asks node C, D, F and G to update their routing tables and forward its application traffic properly. An ONP decides which link bandwidth to reserve from many potential candidates based on price. An ONP is charged per unit time for using reserved link bandwidth by a substrate network during a reservation period. As the reservation period expires and an ONP still needs link bandwidth, the ONP repeats the reservation process.

B. Resource allocation process

1) Reservation process

Nodes in a substrate network keep updated link state of a substrate network and provide the information to an ONP when requested. An ONP first obtains information of updated state of a substrate network including connectivity, residual (unreserved) bandwidth, maximum available bandwidth and price of bandwidth on each link. An ONP then identifies i) a currently available set of paths with minimum cost and ii) an optimal set of paths with minimum cost assuming that bandwidth on all links is unreserved and maximally available. An ONP then decides whether to reserve link bandwidth from a substrate network or from other ONPs. An ONP chooses to buy link bandwidth from a substrate network when the calculated cost for the currently available paths equals to the calculated cost for the optimal paths. An ONP chooses to buy link bandwidth from other ONPs when the calculated cost for currently available paths is higher than the calculated cost for the optimal paths. This is because bandwidth on certain links of the optimal paths is already reserved by other ONPs and not available from a substrate network. In this case, an ONP starts the trading process with other ONPs that have reserved bandwidth on the links. An ONP reserves link bandwidth for a given reservation period.

2) Trading process

A trading process between ONPs is explained with detail. In the trading process, a buyer refers to an ONP that tries to buy link bandwidth from other ONPs, and a seller refers to an ONP that has reserved link bandwidth and is able to sell the link bandwidth to other ONPs.

A buyer identifies a set of links on which bandwidth needs to be reserved. A buyer then sends requests to sellers of bandwidth on the links to start a trading process. Since each node in a substrate network is designed to keep track of overlay networks reserving bandwidth on its incident link, a buyer accesses the nodes incident to the links of interest and obtains information of sellers. A buyer then informs all sellers about the amount of bandwidth that it needs. Each seller calculates its selling price in a way explained later in this subsection and submits a bid to a buyer. Note that, in MaCRA, a seller is not able to know the selling prices of other sellers for fairness of competition. A buyer selects the cheapest one among all bids received from sellers. However, if even the cheapest bid would cost more than reserving currently available paths from a substrate network, a buyer chooses to buy bandwidth of the currently available paths from a substrate network.

A seller calculates selling price of link bandwidth based on i) cost of alternate paths to support its traffic demand and ii) cost of switching paths.

i) Cost of alternate paths

A seller needs to ensure that its traffic demand can be still supported using an alternate set of paths (a set of paths that doesn't include a link whose bandwidth is being traded) during a remaining reservation time before selling link bandwidth that it is currently using. Once a trade is successfully made, the seller will reserve bandwidth on the alternate paths and support traffic demand for a remaining reservation period. Price of bandwidth on alternate paths would be different from price of bandwidth on the current paths. It is likely that alternate paths would cost more. The seller calculates extra charge that it needs to pay to a substrate network for a remaining reservation time in the case of using alternate paths compared to the case of using current paths. The calculated extra charge is added into a selling price.

ii) Cost of switching paths

Cost of switching paths includes cost of updating routing table stored at nodes on a substrate network (routing table update cost) and cost of service interruption for an ONP (service interruption cost). Once a trade is successfully made, a seller switches from current paths to alternate paths to keep supporting its traffic demand. To this end, routing tables at nodes on a substrate network are updated to forward the seller's application traffic correctly. Physical overhead incurred in updating routing tables at nodes are represented as routing table update cost. Notice that, while routing tables at nodes are being updated, the seller's application traffic may experience packet loss and cause service interruption of application. Some applications (e.g., mission critical applications) may be completely intolerant to service interruption. Other applications may be tolerant to service

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

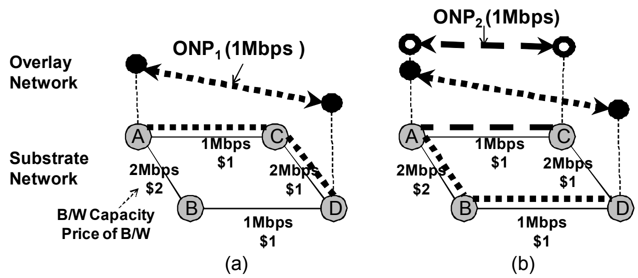

Fig.2 An example of a trade of link bandwidth between ONPs

interruption. The level of tolerance to service interruption is represented as service interruption cost. Service interruption cost for an application that is intolerant may be set extremely high such that the application is less likely to accept the trade unless a buyer is willing to pay the high cost.

Consider an example network illustrated in Fig.2. Notice that each link is annotated with its bandwidth capacity (Mbps) and price of bandwidth ($ per Mbps per hour). Suppose that overlay network provider, ONP_1 has reserved 1Mbps on link A-C and link C-D between 0:00 to 2:00 to deploy its overlay network (See Fig.2 (a)). At 1:00, another overlay network provider ONP_2 tries to deploy an overlay network connecting node A and C. Since a currently available path (A-B-D-C for $4) is more expensive than an optimal path (A-C for $1), ONP_2 requests ONP_1 for a trade of 1Mbps on link A-C. ONP_1 then identifies an alternate path (A-B-D) that doesn't contain link A-C. If ONP_1 would use an alternate path, the total cost during a remaining reservation period would be $3 (A-B for $2 and B-D for $1) times 1 hour (between 1:00 and 2:00). If ONP_1 would keep using a current path, the total cost would be $2 (A-C for $1 and C-D for $1) times 1 hour. Estimated extra charge due to switching to an alternate path (i.e., cost of alternate paths) is $1 ($3 - $2). Notice also that ONP_2 is willing to pay the cost of alternate paths $1 to ONP_1 since using path A-C instead of path A-B-D-C saves more than $1 for ONP_2. In order to switch paths, ONP_1 needs to request node A, B, C and D to update their routing tables, and the cost of switching paths is calculated accordingly. Fig.2 (b) illustrates the case when the trade is successfully made.

III. SIMULATION STUDIES

Simulations have been conducted to investigate the validity and performance of MaCRA architecture in comparison with other mechanisms.

A. Simulation configurations

Substrate network - There is no large-scale service overlay network deployed in the current Internet. In order to investigate MaCRA on a realistic network topology, Level-3 ISP network is chosen as a substrate network as illustrated in Fig.3. There are 23 nodes and each node represents a domain located in a city. There are 34 links and each link represents an OC-48 backbone fiber connecting domains. It is assumed that MaCRA architecture operates on the substrate network, and the same amount of bandwidth on each link (backbone fiber) is available to support overlay networks. Latency of a link is defined to be proportional to physical distance of the link. A pricing for link bandwidth is defined as $p = c_1 * l$ (p: price, c_1: const, l: latency of a link) in simulations.

Arrival and service (reservation) model – A new ONP with application traffic demand arrives following a Poisson process and requests for link bandwidth on a substrate network. Parameters of the arrival and service model are derived from M/M/∞ analysis. Arrival rate of ONPs is varied in order to examine a whole range of traffic load from 10% to 90%. Time period during which an ONP reserves link bandwidth is determined based on Erlang distribution in simulations.

Traffic demand – An ONP knows nodes and links of an overlay network to be deployed. Nodes in an overlay network are selected in a uniform random manner among all 23 nodes on a substrate network. An ONP selects E2E paths to connect each pair of nodes that has an overlay link (i.e, application traffic demand exists between the nodes) to minimize cost. The number of nodes in an overlay network varies from 2 to 5 nodes in simulations. Due to space limitation, the following subsection illustrates results in case of 2 nodes unless stated otherwise. The bandwidth required for traffic demands that each ONP has equals to 10% of maximum available bandwidth on a link.

Cost of switching paths – In simulations, routing table update cost and service interruption cost are assumed to be negligible. In order to investigate the validity and performance of MaCRA, it is assumed that all ONPs are willing to get involved in a trading process in a simulated network.

Other mechanisms to be compared – "Greedy" refers to a version of MaCRA where no trade is allowed and an ONP reserves link bandwidth only from a substrate network. "UBP" refers to Utilization Based Pricing of link bandwidth. The notion of UBP have been used in several research projects [2][9] to achieve load-balancing. Price of link bandwidth is defined by a linear function of bandwidth utilization on a link, $p = c_2 * B_{rsv} / B_{cap}$ (p: price, c_2: const, B_{rsv}: bandwidth already reserved by overlay networks on a link, B_{cap}: bandwidth capacity on a link). In UBP, there is no trade between ONPs, and each ONP selects E2E paths to deploy an overlay network that minimizes the total cost determined by the pricing function above.

B. Simulation results

Bandwidth-weighted latency (BW-latency) of a link is defined as latency of the link multiplied by bandwidth of the

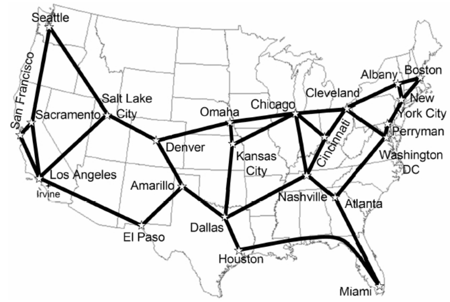

Fig. 3 Level 3 ISP backbone network in U.S.

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

link reserved by a certain overlay network. BW-latency of an overlay network refers to summation of BW-latency of all links whose bandwidth the overlay network reserves. BW-latency of an overlay network indicates the level of latency that application traffic experiences on the overlay network. To compare MaCRA with other mechanisms, a new metric, latency stretch, is introduced. Latency stretch is defined as the ratio of actual BW-latency of an overlay network to optimal BW-latency of an overlay network. Optimal BW-latency is calculated assuming that all of link bandwidth on a substrate network is available and optimal paths can be reserved. Latency stretch indicates how optimal currently reserved paths of an overlay network are.

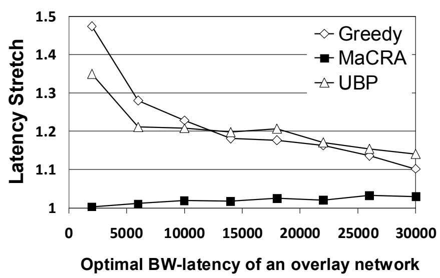

Fig. 4 Latency stretch over optimal BW-latency

Fig.4 shows latency stretch of an overlay network over optimal BW-latency of an overlay network in case of Greedy, MaCRA, and UBP. In Greedy and UBP, an overlay network with small optimal BW-latency experiences high latency stretch, while an overlay network with high optimal BW-latency experiences relatively lower latency stretch. This shows existence of unfairness where an overlay network with the need of reserving short paths (i.e., less number of links) is more likely to obtain paths with much higher latency than an optimal path. In many of networks including the simulated network, difference of latency between the best path (i.e., optimal) and the second (or below) best paths (i.e., alternate) is larger when connecting two closely located nodes than when connecting two distantly located nodes. It is because there could be more possible path options to connect two distantly located nodes than to connect two closely located nodes. MaCRA addresses this unfairness by allowing overlay networks to trade link bandwidth with each other, and it is clearly illustrated in Fig. 4

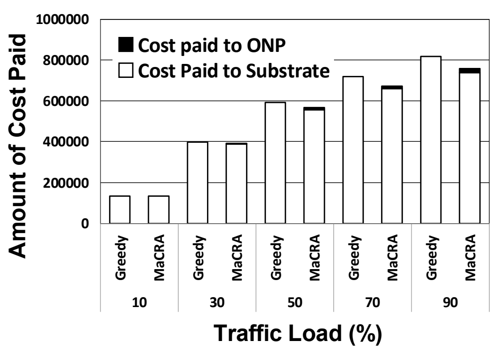

Fig. 5 The amount of cost that ONPs pay over traffic load

Fig.5 shows the amount of cost that all ONPs pay per unit time over traffic load in case of Greedy and MaCRA. In Greedy, ONPs pay cost only to a substrate network for reserved link bandwidth. In MaCRA, ONPs pay cost to a substrate network and to other ONPs if a trade takes place. With low traffic load, the total cost is almost the same between Greedy and MaCRA since all ONPs are able to reserve optimal paths and thus no trade takes place. With the high traffic load, the total amount of cost (cost to a substrate network plus cost to other ONPs) becomes smaller in MaCRA than in Greedy. When the traffic load is as high as 90%, the cost paid to other ONPs in the trade process account for about 3% of the total cost paid by ONPs in MaCRA. Fig.5 demonstrates that the relatively small amount of cost used for performing a trade help reduce the amount of cost paid to a substrate network.

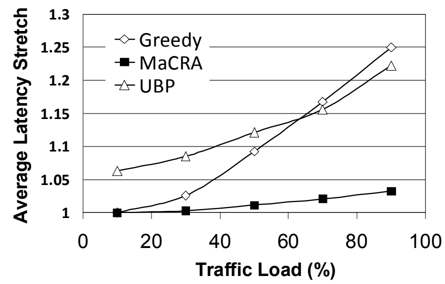

Fig. 6 Average latency stretch over traffic load

Fig.6 show the average latency stretch of an overlay network deployed on a substrate network over traffic load in case of Greedy, MaCRA and UBP. The level of stretch is almost the same between Greedy and MaCRA with small traffic load. With high traffic load, however, average latency stretch surges in Greedy, while it stays low in MaCRA. This result demonstrates that the trade process introduced in MaCRA contributes to reducing average latency stretch of overlay networks. In UBP, average latency stretch is relative high since E2E paths that constitutes an overlay network are selected not to minimize latency but to avoid congestion.

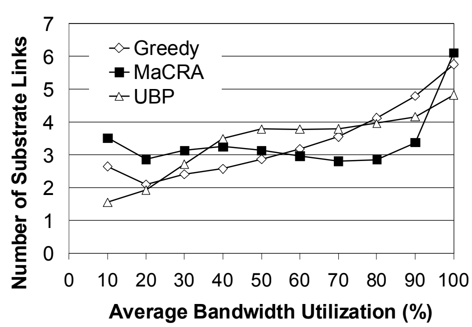

Fig. 7 Number of substarte links with given average bandwidth

Fig.7 shows the number of links with given average bandwidth utilization when the traffic load is high (90%). Average bandwidth utilization is defined as the average ratio of bandwidth reserved by overlay networks on a link to bandwidth capacity on the link. It indicates distribution of average bandwidth utilization on a substrate network. The number of links with high bandwidth utilization being over 60% is smaller in MaCRA than in Greedy. UBP avoids links from reaching high bandwidth utilization by raising price. As

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

a result, the number of links with 100% bandwidth utilization is smaller than Greedy and MaCRA, and the bandwidth utilization is more distributed over links when compared to Greedy and MaCRA.

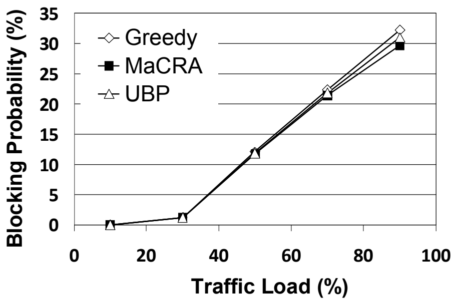

Fig. 8 Blocking probabilty over traffic load

Fig.8 shows the blocking probability over traffic load. The blocking probability is defined as the number of unsuccessful deployment of overlay networks divided by the total number of deployment attempts of overlay networks. The blocking probability is almost the same among three mechanisms with smaller traffic load. With high traffic load, the blocking probability in MaCRA is slightly smaller than Greedy and UBP. This is because trading link bandwidth in MaCRA enables existing overlay networks to reconfigure their reserved paths and release resources for a newly incoming overlay network to fit in.

IV. Related Work

Service overlay network has been actively studied. [1] addresses the problem of bandwidth provisioning from underlying network domains to construct a substrate network. [10][11] focus on optimizing design of topology of a substrate network. MaCRA aims at optimizing allocation of resources on a given substrate network for applications. [2] works on a similar problem with this research and proposes a new source routing mechanism to reserve a high quality path. However, it doesn't support any cooperation between applications reserving a path, and fairness between applications is not considered. [12] optimizes reconfiguration of overlay networks to minimize cost using a centralized approach. It doesn't consider fairness between overlay networks. Cooperation between overlay networks have been discussed in several research projects. [13] applies a bio-inspired mechanism of symbiosis to foster cooperation. MaCRA achieves cooperation of overlay networks for resource allocation and applies a market-based mechanism. [14] proves that selfish routing decisions of end users without any cooperation lead to overall performance degradation (called price of anarchy). MaCRA is intended to address the problem by allowing overlay networks to cooperate through trading resources. [15] outlines distributed algorithmic mechanism designs (DAMD) that provide incentives to each agent in multi-agent systems, such that each agent's decision leads to improved system-level performance. MaCRA belongs to DAMD, and gives incentives for overlay networks to cooperate based on a market mechanism. Our initial idea of MaCRA is presented in [16], and this paper introduces new metrics and evaluates fairness and latency of MaCRA by extensive simulations.

V. Conclusion

This paper introduces the Market-based Cooperative Resource Allocation (MaCRA) architecture. Extensive simulation studies reveal that MaCRA achieves fairness and efficiency in allocating resources for overlay networks and reduces average latency that application traffic running on overlay networks experience especially when traffic load is high. As future work, the authors plan to investigate impact of dynamically changing state of a substrate network (e.g., latency of a link) on the performance of MaCRA. The authors are also designing a distributed lock/mutual exclusion mechanism for resources on a substrate network that allows multiple ONPs to reserve or trade resources concurrently.

References

[1] Z. Duan, Z.-L. Zhang, Y.T. Hou, "Service overlay networks: SLAs, QoS, and bandwidth provisioning," IEEE/ACM Transactions on Networking 11 (6) (2003) pp. 870–883

[2] Z. Li, P. Mohapatra, "QRON: QoS-aware routing in overlay networks," IEEE Journal on Selected Areas in Communications 22 (1) (2004) pp. 29–40.

[3] A. Borodin, R. El-Yaniv, "Online Computation and Competitive Analysis," Cambridge University Press.

[4] Q. Fang, J. Cobb, and E. Leiss, "A pre-selection routing scheme for virtual circuit networks," in Proc. of International Conference on Parallel and Distributed Computing and Systems (PDCS), 2000.

[5] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener, "Provisioning a virtual private network: a network design problem for multicommodity flow," in Proc. of ACM symposium on Theory of computing

[6] Y. Zhu and M. Ammar, "Algorithms for assigning substrate network resources to virtual network components," In Proc. of IEEE INFOCOM, 2006

[7] L. Subramanian, I. Stoica, H. Balakrishnan, R. H. Katz, "OverQoS: An Overlay Based Architecture for Enhancing Internet QoS," in Proc. of USENIX NSDI, 2004, pp. 71-84

[8] D. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris, "Resilient overlay networks," in Proc. of 18th ACM Symp. Operating Systems Principles, 2001, pp. 131–145

[9] A. D. Yahaya, T. Suda, "iREX: Inter-Domain Resource Exchange Architecture," in the Proc. of IEEE INFOCOM, 2006

[10] S. L. Vieira and J¨org Liebeherr, "Topology design for service overlay networks with bandwidth guarantees," in Proc. of IEEE IWQoS, 2004, pp. 211–220

[11] L. Zhou and A. Sen, "Topology Design of Service Overlay Network with A Generalized Cost Model," in Proc. of IEEE GLOBECOM, 2007

[12] J. Fan, M. H. Ammar, "Dynamic Topology Configuration in Service Overlay Networks: A Study of Reconfiguration Policies," in Proc. of IEEE INFOCOM, 2006

[13] N. Wakamiya, M. Murata, "Overlay network symbiosis: evolution and cooperation," in Proc. of the 1st international conference on Bio inspired models of network, information and computing systems

[14] T. Roughgarden, "Selfish Routing," PhD thesis, Cornell University, May 2002

[15] J. Feigenbaum and S. Shenker, "Distributed algorithmic mechanism design: Recent results and future directions," in Proc. of ACM DIALM, 2002.

[16] R. Egashira and T. Suda, "Market-Based Resource Allocation for Service Overlay Networks," in Proc. of IEEE CAMAD, 2009