Interaction of Overlay Networks: Properties and Implications

Wenjie Jiang John C.S. Lui
Department of Computer Science & Engineering
The Chinese University of Hong Kong
{wjjiang, cslui}@cse.cuhk.edu.hk

Dah-Ming Chiu
Information Engineering Department
The Chinese University of Hong Kong
dmchiu@ie.cuhk.edu.hk

ABSTRACT

Although the concept of application layer overlay routing has received much attention lately, there has been little focus on the "*co-existence*" and "*interaction*" of overlays on top of the same physical network. In this paper, we show that when each overlay plays the optimal routing strategy so as to optimize its own performance, there exists an equilibrium point for the overall routing strategy. However, the equilibrium may be *inefficient*: (a) it may not be Pareto optimal, (b) some fairness anomalies of resource allocation may occur. This is worthy of attention since overlays can be easily deployed and overlays may not know the existence of each other, they may continue to operate at a sub-optimal point.

1. INTRODUCTION

In recent years, there has been tremendous interest on the routing and deployment of overlay or peer-to-peer networks[1, 6]. In particular, application layer routing schemes are shown to effectively address the problems of traditional IP routing. Measurements from [1, 6, 7] indicate that in the current Internet, a large percentage of traffic can find better routes by relaying packets with the assistance of overlay nodes. From a theoretic point of view, application layer routing is a form of optimization in which an overlay maximizes its utility based on the available network resources and information.

Although the concept of overlay networks has received much attention lately, there has been little focus on the "*interaction*" of "*co-existence*" of overlay networks. In this work, we consider the scenario when multiple overlays are constructed on top of a common physical network. These overlays have partially overlapping paths and even nodes. Each overlay is "*selfish*" by nature in that it performs overlay routing so as to optimize its own performance without considering the impact on other overlays. We explore this form of interaction and how the interaction can affect the network stability, performance and fairness in resource allocation.

In this work, we derive the fundamental properties of Internet overlay interactions systematically. Firstly, we introduce the concept of *overlay optimal routing* policy. Secondly, we model the interaction of overlays as a non-cooperative strategic game and show that even when multiple overlays each striving for its own optimality, there always exists a *Nash equilibrium*, under very mild assumptions about the delay functions of the overlay's links and general network topology. Finally, we report a number of important properties of overlay optimization. Namely,

- **Sub-optimality**: The equilibrium point is not Patero (or social) optimal, which can cause "tragedy of the commons", meaning that the performance of *all* overlays can be seriously degraded, despite their individual routing optimization.
- **Fairness paradox**: Another more interesting and important result is on the issue of fairness in resource allocation. Namely, at the equilibrium, it is possible for some overlays to obtain a higher percentage of the common resource (*e.g.*, link bandwidth) as compared to other overlays and cause these overlays to experience a significant performance degradation.

2. OVERLAY OPTIMAL ROUTING: MODELING

In this section, we first propose the "*overlay optimal routing*" policy. Normally there are multiple source-sink pairs in an overlay network. Under the optimal routing policy, every source node decides to split traffic among *all* its available overlay paths, coordinated by the overlay network to minimize the average delay for the *whole* overlay. In this model, we assume that there is no underlying traffic in the physical network. We show that even under a less dynamic environment (which justifies in the long run), the interaction between overlays can result in some undesirable properties.

Consider a *physical network* with a set $\mathcal{J}$ of *resources*, which denotes a set of physical links. For each link $j \in \mathcal{J}$, let C_j represent its finite capacity (unit is bps). Let a *route* r be a non-empty subset of $\mathcal{J}$, and denote $\mathcal{R}$ as the set of all possible routes of the physical network. Let $|\mathcal{J}| = m$ and $|\mathcal{R}| = q$, we use A to represent an $m \times q$ matrix with $A_{jr} = 1$ if $j \in r$, and $A_{jr} = 0$ otherwise. Thus, the matrix A defines a $0-1$ link-route indicator matrix. Let $d_j(l_j)$ denote the delay function for the physical link $j \in \mathcal{J}$, where l_j is the *aggregate* rate of traffic that traverses link j. Let $L = (l_1, l_2, \ldots, l_m)^{\mathcal{T}}$ denote a traffic rate vector for all physical links, and we define a delay function as $\mathcal{D}(L) = (d_1(l_1), \ldots, d_j(l_j), \ldots, d_m(l_m))^{\mathcal{T}}$. In this work, we only assume that the delay function is continuous, non-decreasing, and convex. Note that this is a reasonable assumption since this applies to a link with a fixed propagation delay, or a link whose delay is represented by general queueing delay models.

An overlay network is a connected sub-graph of the underlying physical network which consists of a set of logical nodes and logical links. A logical path is interpreted as a set of logical links, each of which may consist of one or more physical links. The logical topology of the overlay network depends heavily on how this overlay is organized. With proper translation, we can map every logical path to a set of corresponding physical links. Thus, the routing matrix for overlay s can be similarly defined as $A^{(s)}$, which is a partial matrix of A. Within an overlay, there can be "*multiple*" source-sink pairs and each source-sink is associated with a traffic flow f, which has a constant traffic demand of x_f (units is bps). Suppose there are a set $\mathcal{N}$ of overlays, and for each overlay $s \in \mathcal{N}$, there is a finite set $\mathcal{F}_s$ of source-sink pairs. For each flow $f \in \mathcal{F}_s$, there is a set $\mathcal{R}_f$ of distinct paths that can be used by the flow f to deliver information from its source to its sink. In our model, each overlay has the ability to control the routing of its traffic *within* its overlay network. Therefore, source nodes of an overlay network may choose to split and assign their traffic onto different paths so that the weighted average delay of the *whole* overlay network can be minimized. Equivalently, the overlay needs to decide, for all its flows, how to assign traffic to every possible path $r \in \mathcal{R}_f$ so as to optimize its desired performance. Thus, each flow f in the overlay s has a routing decision vector $y^{(s,f)} = \left(y_1^{(s,f)}, y_2^{(s,f)}, \ldots, y_{|\mathcal{R}_f|}^{(s,f)}\right)^{\mathcal{T}}$, where $y_k^{(s,f)}$ is the amount of traffic along k-th path for flow f in overlay s, and $\sum_{k=1}^{|\mathcal{R}_f|} y_k^{(s,f)} = x_f$ such that the traffic demand for flow f is satisfied. For the compactness of presentation, we rewrite the routing decision for overlay s as a concatenation of the flow vectors of all its source-sink pairs: $y^{(s)} = \left(y^{(s,f_1)}, y^{(s,f_2)}, \ldots, y^{(s,f_{|\mathcal{F}_s|})}\right)$.

In the overlay optimal routing, one overlay's goal is to minimize the average weighted delay of traffic within this overlay, which can be interpreted as the sum of weighted end-to-end delays on all possible paths, with the weight being the rate of overlay traffic traversing on that corresponding path. With these notations, the overlay optimal routing for overlay s can be formulated as the following convex optimization problem [1] **OVERLAY**$^{(s)}$:

$$
\begin{aligned}
\textbf{Minimize} \quad & \text{delay}^{(s)} = y^{(s)\mathcal{T}}[A^{(s)\mathcal{T}}\mathcal{D}(\textstyle\sum_i A^{(i)}y^{(i)})] \\
\text{s. t.} \quad & \forall f \in \mathcal{F}_s,\ \textstyle\sum_{k=1}^{|\mathcal{R}_f|} y_k^{(s,f)} = x_f, \\
& Ay \leq C, \quad y^{(s)} \geq 0. \qquad (1)
\end{aligned}
$$

Since the routing decision optimizer $y^{(s)}$ of overlay s depends on the current routing decisions $y^{(-s)}$ of other overlays, there will be inevitable interaction of routing behaviors between these co-existing overlays. Therefore, the interaction can be understood as an iterative process, wherein each overlay periodically calculates its optimal routing strategy based on other overlays' routing decisions (assuming fixed during the calculation), which may be performed at different time scalings. To analyze the interplay between overlays, we model the interaction between multiple co-existing overlays as a non-cooperative Nash routing game G_{overlay}. First, the set of players consists of all overlays. Second, the strategy that overlay s can take is a feasible routing vector $y^{(s)} \in \Gamma^{(s)}$, which is defined as the feasible region of the optimization problem (1). Lastly, each overlay prefers a lower average delay. Accordingly, we have the following definition for the *Nash Equilibrium* in G_{overlay}:

Definition 1. *A feasible strategy profile* $y^* \in \Gamma_1 \times \ldots \times \Gamma_n$, $y^* = (y^{*(1)}, \ldots, y^{*(s)}, \ldots, y^{*(n)})^{\mathcal{T}}$ *is called a Nash equilibrium if for every player* $s \in \mathcal{N}$, $\text{delay}^{(s)}(y^{*(1)}, \ldots, y^{*(s)}, \ldots, y^{*(n)})$ *is less than or equal to* $\text{delay}^{(s)}(y^{*(1)}, \ldots, y^{'(s)}, \ldots, y^{*(n)})$ *for any other feasible strategy profile* $y'^{(s)}$.

Theorem 1. *For the overlay optimal routing game defined above, there always exists a Nash Equilibrium Point (NEP) if the delay function* delay$^{(s)}$ *is continuous, non-decreasing and convex.* [2]

3. IMPLICATIONS OF INTERACTION

In this section, we discuss some intrinsic problems of overlays interaction. These problems include sub-optimality in performance and certain fairness anomaly in resource allocation. It is important to point out that these problems are not unique to overlay optimal routing policy, but rather, common to all forms of application layer routing that have interaction among overlays. Worse yet, because overlay may not realize the existence of other overlays, these problems will persist due to the convergence to the equilibrium point.

[1]For the detailed derivation of this optimization model, see [2].
[2]Please refer to [2] for a detailed proof.

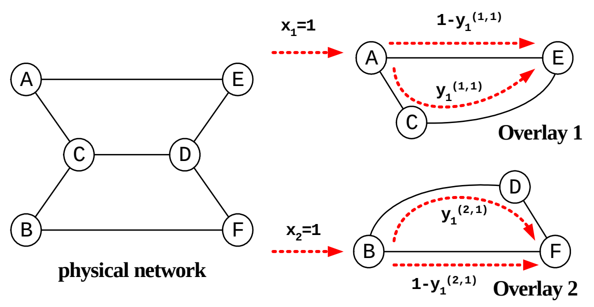

Figure 1: A simple network with two overlays to illustrate potential problems

Let us use an example to illustrate these issues. A physical network consisting of six nodes is depicted in Figure 1. There are two overlays in the network: overlay 1 consists of node A, C, E while overlay 2 consists of node B, D, F. For overlay 1, all logical links map to the corresponding physical links except for the logical link between node C and node E, which corresponds to the physical links $C - D - E$. For overlay 2, all logical links map to the corresponding physical links except for the logical link between node B and node D, which corresponds to the physical links $B - C - D$. Thus, the physical link $C - D$ is the common link which is *shared* by these two overlays.

3.1 Sub-optimality of Nash Equilibrium

Assume that both overlays have one source-sink pair and one unit of traffic demand: $x_1 = x_2 = 1.0$. We define the following delay functions for various physical links in the physical network: $d_{A,E}(y) = a + y$; $d_{C,D}(y) = by^{\alpha}$; $d_{B,F}(y) = c + y$, while other links have zero delay. Here, y represents the aggregate traffic traversing a link, and a, b, c, α are some non-negative parameters of the delay functions.

Let us consider the routing decisions of these two overlays. For overlay 1, it routes $y_1^{(1,1)}$ *fractional* unit of traffic through the logical path A-C-E and $(1-y_1^{(1,1)})$ (because $x_1 = 1$) fractional unit of traffic via the logical path A-E. On the other hand, overlay 2 routes $y_1^{(2,1)}$ fractional unit of traffic through the logical path B-D-F and $(1 - y_1^{(2,1)})$ fractional unit of traffic via B-F. To find out the Nash equilibrium point, we write down the Karush-Kuhn-Tucker (KKT) conditions for overlay 1:

$$a+2\left[1-y_1^{(1,1)}\right] = b\left[y_1^{(1,1)} + y_1^{(2,1)}\right]^{\alpha} + y_1^{(1,1)} \cdot b\alpha\left[y_1^{(1,1)} + y_1^{(2,1)}\right]^{\alpha-1} \qquad (2)$$

while the KKT conditions for overlay 2 is

$$c+2\left[1-y_1^{(2,1)}\right] = b\left[y_1^{(1,1)} + y_1^{(2,1)}\right]^{\alpha} + y_1^{(2,1)} \cdot b\alpha\left[y_1^{(1,1)} + y_1^{(2,1)}\right]^{\alpha-1}. \qquad (3)$$

where $y_1^{(1,1)}, y_1^{(2,1)} \in [0, 1]$.

One can easily show that in the overlay optimal routing game described above, the Nash equilibrium point is *not* Pareto optimal. A Pareto optimal point is defined as a strategy profile for all overlays such that no overlay can use another routing strategy that can decrease its own weighted average delay *without* increasing other overlays' weighted average delay. Namely, the equilibrium point is not Pareto optimal since there exists another routing strategy by which *all* overlays can achieve a *better* performance than at the Nash equilibrium.

To show the sub-optimality of the Nash equilibrium in the example network depicted in Figure 1, we consider the KKT conditions specified by Equations (2) and (3). Assume we have the following parameters for the delay functions: $\alpha = 1, a = 1, b = 1$ and $c = 2.5$, one can simply verify that the Nash equilibrium in this example is $\{y_1^{(1,1)} = 0.5, y_1^{(2,1)} = 1\}$, that is, overlay 1 uses both paths while overlay 2 uses a single path, which consists of the shared link. The weighted average delay for overlay 1 and 2 is both 1.5. However, if we consider another routing strategy profile of $\{y_1^{(1,1)} = 0.4, y_1^{(2,1)} = 0.9\}$, one can find that the weighted average delay for overlay 1 and 2 are 1.48 and 1.43 respectively, which are *lower* than the delay achieved at the Nash equilibrium.

3.2 Fairness Paradox

Another more severe problem is the notion of fairness in resource allocation. We use the same network in Figure 1 to illustrate the problem. Note that these two overlays are symmetric, each having two paths: a shared path and a private path. As in the previous example, although overlay 2 is "*worse off*" by having a private path (link $B-F$) with higher delays than that of overlay 1's private path (link $A-E$), it is able to achieve the same average delay as overlay 1 in the Nash equilibrium. This is because overlay 2 is able to fully take advantage of the lower delay of the shared path, whereas it only makes sense for overlay 1 to send part of its traffic over the shared link due to its superior private path. In fact, one can find delay functions such that the situation is *arbitrarily worse*.

To illustrate, note that the delay function for the shared link $C-D$ is $d_{C,D}(y) = by^{\alpha}$. One can ask for what values of a and c, which are the parameters of the delay functions for the private link of overlay 1 and 2, so that the Nash equilibrium solution remains at $\{y_1^{(1,1)} = 0.5, y_1^{(2,1)} = 1\}$? The values for a and c do exist, in particular, when $b = 1$ and $a < c$, we have:

$$a = \left(\frac{3}{2}\right)^{\alpha} + \frac{\alpha}{2}\left(\frac{3}{2}\right)^{\alpha-1} - 1 \quad ; \quad c = \left(\frac{3}{2}\right)^{\alpha} + \alpha\left(\frac{3}{2}\right)^{\alpha-1}.$$

In Section 3.1, we showed when $a < c$ and $\alpha = 1$, $\text{delay}_1 = \text{delay}_2$. When $\alpha > 1$:

$$\text{delay}_1 = \left(\frac{3}{2}\right)^{\alpha} + \frac{\alpha}{4}\left(\frac{3}{2}\right)^{\alpha-1} - \frac{1}{4}; \qquad \text{delay}_2 = \left(\frac{3}{2}\right)^{\alpha},$$

and observe now that delay_1 becomes *greater* than delay_2. This implies that the overlay 2 is able to achieve *better* performance despite starting with a *worse* private link in an otherwise symmetric situation with overlay 1. Furthermore, as we increase α, this unfairness can even be *unbounded*, that is:

$$\left.\frac{\text{delay}_1}{\text{delay}_2}\right|_{\alpha\to\infty} = \infty$$

as depicted in Figure 2(a). This type of anomaly also exists in other operating range, for example, when $a = 2, c = 4, \alpha = 1$ and we vary the values of b, one can observe that overlay 1 will have a worse performance compared to overlay 2 (though the unfairness is bounded in this case). In summary, we illustrate that there exist delay functions for the links such that although overlay 1 seemingly has better paths than overlay 2, it is destined to lose the routing game to overlay 2 by an arbitrary margin: a rather "*paradoxical*" situation.

4. RELATED WORK AND CONCLUSION

Efficiency loss of selfish routing in a non-cooperative environment has been studied by Roughgarden in [5]. A probabilistic routing protocol [8] was proposed to implement selfish routing as well

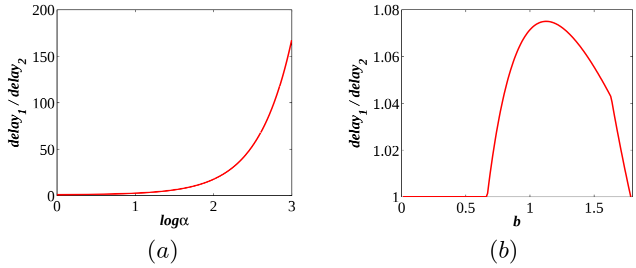

Figure 2: $delay_1/delay_2$ **ratio v.s. (a)** $\log(\alpha)$**: unfairness becomes unbounded; (b) parameter** b **with** $a = 2, c = 4$ **and** $\alpha = 1$**: bounded unfairness.**

as optimal routing in an overlay network. A seminal work on the interaction of an overlay and the underlying ISP was presented in [3], in which authors use a two-player non-cooperative game model to study the interaction between one overlay and the underlying ISP. The contribution of our work is that we use a non-cooperative game framework to study the interaction between *multiple* overlays, in a static underlay environment. All the players (*overlays*) are in an equal stand of the game, while in [3] the overlay and the underlay may have different levels of knowledge about the underlying network. We show that under this fully competitive environment, unregulated behaviors will converge to an inefficient equilibrium, and finally results in a series of undesirable properties. Interaction of overlay and underlay was also discussed in [4].

In this paper, we explore the interactions between multiple co-existing overlays. We consider the situation wherein overlays can determine their routing so as to optimize their individual performance measures. We model the interaction of these co-existing overlays as a non-cooperative strategic routing game and analyze the system stability. We further demonstrate several inherent properties of this form of individual routing optimization, namely the performance measure is non Pareto-optimal, and the fairness paradox in terms of performance and allocation of common resources. We believe our result will shed light into the cause of these problems and bring awareness and stimulate further research in solving these problems.

Acknowledgement: The work of John C.S. Lui was supported in part by the Hong Kong RGC Grant.

5. REFERENCES

[1] D. Anderson, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resillient Overlay Networks. In *ACM SIGCOMM*, Oct. 2001.

[2] W. Jiang, D.-M. Chiu, and J. C. S. Lui. Interaction of Overlay Networks: Properties and Implications. Tech. Report, CUHK, 2005.

[3] Y. Liu, H. Zhang, , W. Gong, and D. Towsley. On the Interaction Between Overlay Routing and Underlay Routing. In *Proceeding of IEEE INFOCOM'05*, Miami, FL, USA, Mar. 2005.

[4] R.Keralapura, N. Taft, C.N.Chuah, and G.Iannaconne. Can ISPs Take the Heat from Overlay Networks? In *ACM HotNets*, Nov. 2004.

[5] T. Roughgarden and E. Tardos. How Bad is Selfish Routing? *Journal of ACM*, 49(2):236–259, 2002.

[6] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan. Detour: a Case for Informed Internet Routing and Transport. *IEEE Micro*, 19(1):50–59, Jan. 1999.

[7] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson. The end-to-end Effects of Internet Path Selection. In *Proceedings of ACM SIGCOMM '99*, pages 289–299, Cambridge, MA, Aug. 1999.

[8] H. Xie, L. Qiu, Y. R. Yang, and Y. Zhang. On self adaptive routing in dynamic environments. In *Proceedings of International Conference on Network Protocols*, Berlin, Germany, Oct. 2004.