Service management for multi-domain Active Networks

Alvin Tan, Walter Eaves, Chris Todd

Department of Electronic & Electrical Engineering,
University College London
E-mail: {atan, weaves, c.todd}@ee.ucl.ac.uk

Abstract: . The fundamental issue that we address is the allocation of limited network communication and computational resources. Our goal is to exploit the economic situation of the stakeholders so that they can apply a range of simple local policies that should give rise to more stable global situations. We model the inter-domain environment such that Axelrod's notion of norms in social setting can be mapped on to our problem space. We further use the Paris Metro Pricing approach as we foresee the need for an explicit charging mechanism to prevent a 'tragedy of the commons'. We finally describe our software implementation of an inter-domain manager (IDM), which has the objective of keeping state on an administrative domain basis as opposed to per network element.

1. Introduction

Norms exist in a given social setting to the extent that individuals usually act in a certain manner and are often punished when seen not acting this way [1]. We view it as an emergent pattern of interaction that is perceived to be 'correct' by the community of equals that comprise this group. We categorise the Internet 'stakeholders' into network service providers (NSPs) who provide transit, long-distance, raw capacity, backbone connectivity; Internet service providers (ISPs) who provide points of presence, *i.e.*, sell access to the Internet; and end-users. A more specific role played by ISPs, who are also providing 'open' interfaces (*i.e.*, those provided by active routers), is the Active Network provider (ANP). Our goal is to exploit the economic situation of the stakeholders so that they can apply a range of simple local policies that should give rise to more stable global solutions.

The advent of Active Networking (AN) technology necessitates node-level, per-flow computational resource allocations, in addition to prevailing resource reservation approaches for communication bandwidth. The node operating system is responsible for allocating and scheduling the node's resources (*e.g.*, link bandwidth and CPU cycles). Each execution environment (EE) implements a virtual machine that interprets active packets that arrive at the node. This is an abstraction that is used only for the purpose of partitioning the resources of the active node such that different communities of users can stay isolated from each other.

2. Adaptation to a network context

This general setting has multiple end-users transmitting packets through an upstream bottleneck (see Figure 1), competing for bandwidth within the upstream provider's domain and CPU time if AN capabilities are required. The abstraction for this analysis is sufficiently general for consideration of inter-ISP[1] flows.

If a real-time, operating system (RTOS) were to run on the network node where the bottleneck is formed, then we may allocate periods of guaranteed processing time on a per-flow basis. However, this would mean periods when processing time is over-allocated, thus giving rise to inefficient usage of valuable computational resource time. A more common scheduling mechanism is that of a Unix-like operating system. Each flow is allocated a fixed time slice. Problem arises when an incoming active packet flow is too large. The allocated time could be insufficient to complete the task, thus forcing the running application to an abrupt end. This may cause the dropping and resending of the application stream.

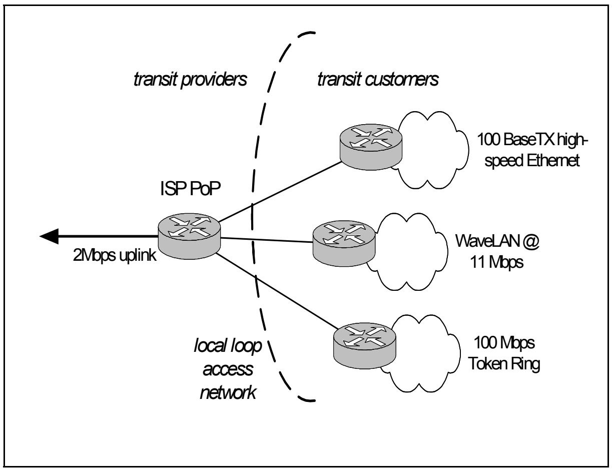

Figure 1: Customer flows coming in at higher speed than ISP's uplink

The bottleneck is a first-in-first-out (FIFO) queue. In this competitive setting, the ISPs are egoistic in that their utility function reflects only their own performance, and not that of the network as a whole. If

[1] Recall that we introduced the Active Network provider (ANP) as role that can be played by any ISP that has active nodes within its network; hence we use the acronyms interchangeably depending on the context.

S_i is the set of strategies s_i for end-user i, let S_{-i} be the product set of all the other end-users' strategy sets. Using a game theoretic approach, the best reply correspondence can viewed of the following relation, R_i: $S_{-i} \rightarrow S_i$, where R_i seeks to maximise i's utility. We then ask whether these end-users will reach an equilibrium point, where dominant strategies emerge. By intuition, for two end-users, Nash equilibrium [4] occurs when one end-users best reply to another will also yield the recipient's best reply. Using simulation, we will show that these strategies, which vary in terms of vengefulness and boldness, will self-regulate towards relatively low levels of boldness if metanorms are adopted at bottlenecks.

2.1. Metanorms

Axelrod's idea of metanorms [1] attempts to establish a norm that one must punish those who do not punish a defection. In practice, logs at the access router for each sub-network will know which downstream sub-network defects (*i.e.*, using up a lot of bandwidth or taking a huge slice of processing time). Over time, a sub-network (*e.g.*, transit ISPs) gets 'frustrated' and starts enforcing norms within its downstream users. This practice will propagate all the way down to per-user granularity of metanorm enforcement. The access router, via the inter-domain manager (IDM), sends a log report to the other IDMs of the downstream sub-network, which replies with the choice of whether to punish or not.

An opportunity to defect is accompanied by a known chance of being observed (S). Boldness (B) and vengefulness (V) define a player's[2] strategy. If B > S, the player will defect, hence obtains a payoff of 3 credits. Every other player in the group gets 'hurt' by 1 credit. Should the defector be spotted, it shall be heavily punished at a cost of 9 credits, while the punisher, in turn, incurs a cost of 2 credits for its enforcement action. After receiving the list of punishment decisions from the sub-networks, the upstream IDM further sends out a final list of non-punishers in order to enable metanorm enforcement. The sub-networks can either choose to punish the non-punishers and incur a further cost of 2 credits or do nothing about it. As with the basic norms game, this depends on a sub-network's vengefulness and the non-punishers who get punished will suffer cost of 9 credits.

If end-users were simply Web browsing, there would be no noticeable deterioration in the maximum speed of the line, due to burstiness. However, if end-users are transferring large amounts of data such as MP3 files through peer-to-peer file sharing, or streaming video content, then congestion can quickly become obvious as the connection becomes saturated. The latter is our context of defection.

2.2. Paris Metro Pricing (PMP)

The fundamental issue that we are addressing is the allocation of limited resource[3], *i.e.*, having multiple sub-networks that attempt to maximise the usage of their upstream provider's uplink as well as computational resource; and any solution would need an explicit charging mechanism to prevent 'tragedy of the commons'. In Odlyzko's Paris Metro Pricing (PMP), providers do not provide formal quality of service (QoS) guarantees [6] and packets are still handled on a best effort basis. The network is logically partitioned into separated channels. One of the unused bits in IPv4 is needed to separate the flows into premium channel and best-effort channel. Physically, no traffic engineering is performed. Differentiation in congestion levels is provided by the price discrepancy between the two channels. The simple rationale is that when premium channel gets congested, end-users will realise that it is not worth paying for the quality that they are getting, so they drop to the best-effort channel. This self-regulating behaviour is expected to restore the differential in QoS.

Thus, summarising §2.1 and §2.2, traffic from a sub-network may stay in the premium channel: if it pays more (hence increases the credits); or of it does not defect, *i.e.*, self-imposed discipline on sending traffic through provider's uplink. With regards to metanorms game, when a sub-network has been punished too much, its credits will drop. The provider must allow the sub-networks to replenish credits by paying more. As a mechanism to avoid dropping to the best-effort channel in the middle of a transmission, a derivative market can be established in at the queue. The downstream subnetwork can exercise call options when credits drop below a certain level. We adapt Semret's [8] approach to enable an end-user to pay a reservation fee, which gives the right to buy extra credits at any time in the future at its bid price for the premium channel to stay within it.

3. Inter-domain manager

The inter-domain manager (IDM) is the software component that represents a domain (or sub-network) when communicating control signals with an upstream provider. Nichols *et al* [5] observed that multilateral agreements rarely work. Instead, end-to-end services need to be constructed out of purely bilateral agreements. The IDM thus

---

[2] We use the term player to denote both a sub-network as a single entity of domain, or end-user.

[3] In fact, due to the bursty nature of Internet traffic, when an end-user connects to the Internet, her link is shared with other Internet end-users at almost every step of the way (except the modem dialup). Even for leased lines, end-users contend within the corporate LAN itself.

keeps state on an administrative domain basis as opposed to per network element. We note that the IDM is implemented using a server-based approach where the control entity is physically separated from the router. This has the advantage of removing the message processing and storage burden from routers.

Before two IDMs negotiate, they must have had a prior agreement to enable each other to send request for participation in our resource reservation mechanism (otherwise, it is simply best-effort) and subsequently to 'haggle' for the initial best price for the premium channel. This agreement could be done at the SLA definition phase. There must be at least one router in a domain that understands the protocol so that it can divert request (REQ) messages to the IDM.

In a multi-domain environment, the reservation path depends as much on an ISP's policy as it does on resource availability. Thus, a domain en-route might have the necessary bandwidth and node processing power to accommodate an additional EE thread, but could choose not to allocate them to the requestor due to its own reasons. As such, it will be in the ISP's discretion to issue a reject (REJ) message back to the original sender. Otherwise, if the ISP agrees to let future sub-network flows use the premium channel, then an acknowledgement (ACK) message is replied.

4. Simulation results

4.1. Basic norms game

In our simulation using Java, we adapted the basic norms game to reflect the bottleneck at upstream ISP's uplink. Boldness in our context means having sub-networks push high traffic load through the uplink. The strategies of the sub-networks for the initial population are chosen at random. For each generation (indicated by different shades of plot in Figure 2), the scores are calculated using the payoff explained in §2.1. The sub-network with the most successful boldness-vengefulness strategies, *i.e.*, one standard deviation (σ) above the average (μ) score for that population, will be retained as well as having its strategy replicated; while those between μ and (μ+ σ) will simply be retained. The other sub-networks are eliminated, while new ones are regenerated randomly to maintain the total population at 20.

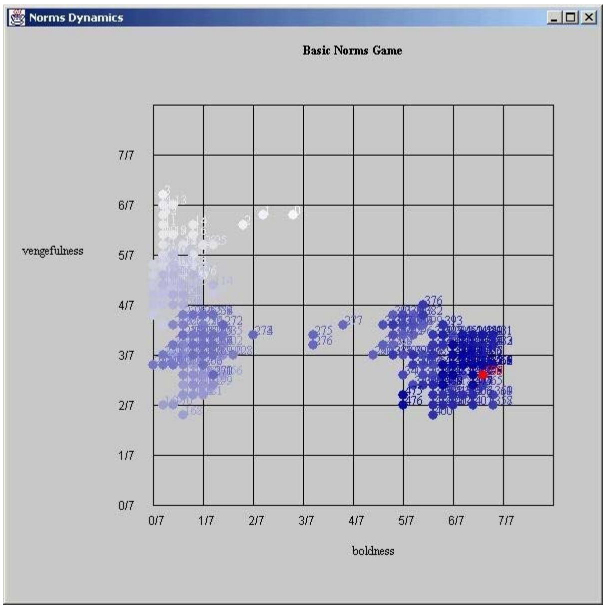

Figure 2: Basic norms game. Lighter shaded plots indicate earlier generation.

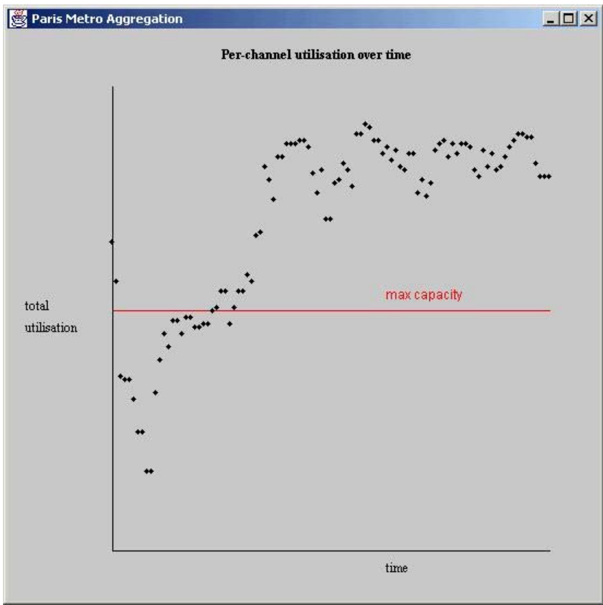

Figure 3: Temporal adaptation of channel utilisation with no incentive to punish 'defection'.

Consistent with Axelrod's results, the first thing to happen is the fall in boldness level. Once the boldness level fell, the vengefulness level dropped as well. This enabled an evolution towards bolder sub-networks, thus giving rise to channel congestion (see Figure 3). Maximum capacity is arbitrarily set to equal to 75% of the max boldness (analogous to the capacity of the premium channel of the uplink) of 20 sub-networks. A high level of boldness indicates that a large amount of packets will be dropped.

4.2. Metanorms game

In the basic norms game, the sub-networks are not motivated to punish a non-punisher (*i.e.*, a defection). The metanorms approach is expected to provide more incentive for being vengeful. A drop in the level of boldness (see Figure 4) justifies this incentive. As with the simulation in §4.1, we repeat the scoring and evolution for 500 generations. The result shows a distinctive drop of boldness level that stays at that low level.

In Figure 5, where we use boldness to represent bandwidth usage per sub-network for the premium channel, the drop in total channel utilisation is consistent with the low level of boldness in the metanorms game.

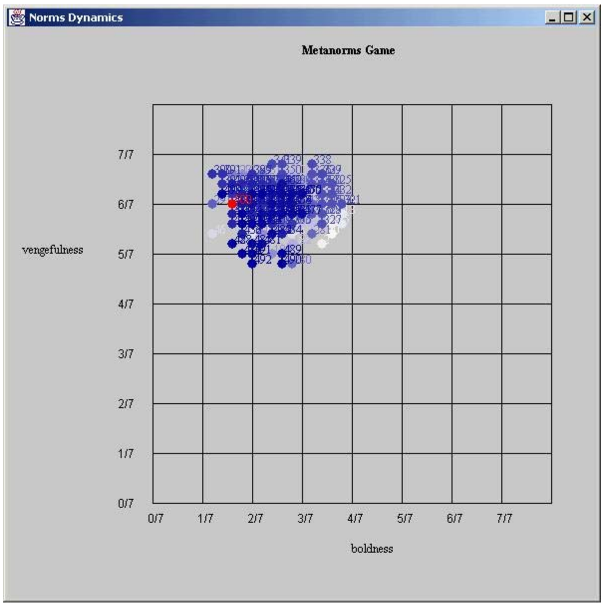

Figure 4: Metanorms game. Boldness level drops and remains low.

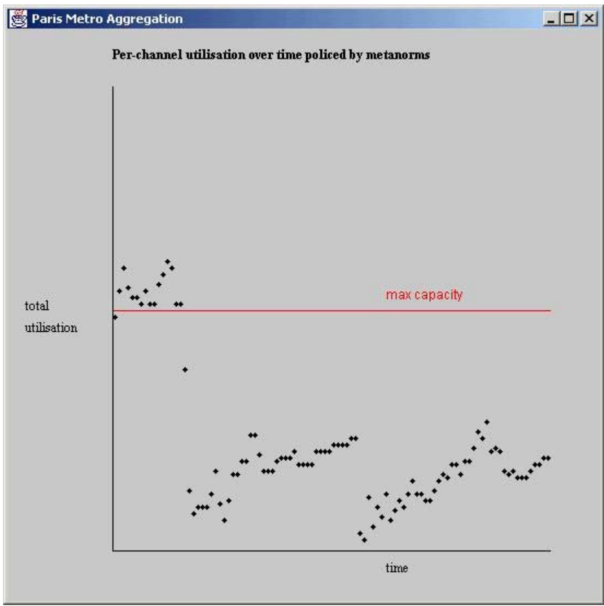

Figure 5: Temporal adaptation of channel utilisation with incentive to punish 'defection'.

5. Conclusion and discussion

We have presented a model for multi-ISP interaction at a network bottleneck in view of allocation of limited network communication and computational resources. We adapted Axelrod's metanorms approach to show that channel utilisation can be used to place users into groups, which can then be allocated to the premium or best-effort classes. In a related work by Gibbens and Key [3], which is based on similar tenets of game theory, the end-user's strategy for resource optimisation is investigated. The key similarity in our work is their use of congestion signals [7] that indicate a cooperative approach, where user behaviour is constrained, and policed. We incorporate the concept of a pricing mechanism, *i.e.*, the PMP.

On a separate but relevant context, basic ADSL services in the typically have a contention ratio of 50:1, with the more expensive or business packages having 20:1. The ADSL line from the customer's property terminates at an ADSL card at the local exchange, and the DSLAM unit combines all the ADSL data onto a single pipe, which carries all Internet traffic to and from the exchange. Contention therefore occurs within the DSLAM unit, with all the ADSL users in an exchange competing for the pipe; hence the possibility for using our same approach for contention management.

6. Bibliography

[1] Axelrod, Robert. An evolutionary approach to norms. In *American Political Science Review* 80, no. 4, December 1986.

[2] Fagin, R., J.Y. Halpern, Y. Moses, and M.Y. Vardi. Reasoning about knowledge. *MIT Press*, 1995.

[3] Gibbens, R. J., and P. B. Key. The use of games to assess user strategies for differential quality of service in the Internet. In *Workshop on Internet Service Quality Economics (ISQE)*, December 1999.

[4] Nash, John. Equilibrium points in n-person games. In *Proceedings of NAS*, 1950.

[5] Nichols, K., V. Jacobson, and L. Zhang. A two-bit differentiated services architecture for the Internet. In *IETF RFC 2638*, July 1999.

[6] Odlyzko, Andrew. Paris Metro pricing for the Internet. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, 1999.

[7] Ramakrishnan, K., S. Floyd, and D. Black. The addition of explicit congestion notification (ECN) to IP. In *IETF RFC 3168*, September 2001.

[8] Semret, N., and A. Lazar. Spot and derivative markets in admission control. In *16th International Teletraffic Congress*, Edinburgh, June 1999.