Towards a reasoning for interaction between inter-domain management entities

Alvin Tan, Walter Eaves, Alex Galis, Chris Todd
Department of Electronic & Electrical Engineering,
University College London
E-mail: {atan, weaves, a.galis, c.todd}@ee.ucl.ac.uk

Abstract: . Internet Service Providers mutually interact for connectivity's sake, but the fact remains that two peering agents are inevitably self-interested. Contention occurs when the entity with excess resource intends to offer it to another, which is the case between two peering ISPs. This scenario is brought on by the active and programmable networks paradigm where network nodes offer dedicated execution environments (EEs). We discuss this interaction using a formal approach. We will also describe the problem space for defining a control protocol for inter-domain network management. Our motivation is to ensure that response to a specific service deployment is consistent and predictable. We want to control the contention for network resources such that a service element is provided with a superior level of network resource while the other service element does not obtain an unfair allocation of resources (to some definition of 'fairness').

1 Introduction

Underneath the veneer of the service provider's[1] competitive retail and wholesale environment, every service provider must 'strategically' interoperate with neighbouring (whether upstream, downstream or peering) networks in order to provide comprehensive connectivity and end-to-end service. The Internet is an example of a multi-agent system [2]. In this context, an agent is synonymous to a service provider. The service providers are able to act in an environment (*i.e.*, the Internet) where different service providers would have different 'spheres of influence', in the sense that they will have control and management responsibilities over different parts of the environment. We will refer to these service providers as Internet Service Providers (ISPs[2]) henceforth.

A domain is defined as a sphere of influence. An Internet domain refers to a set of routers under a single technical administration, using interior gateway protocols to route packets within the domain and using an exterior gateway protocol to route packets to other domains. When a domain routes transit traffic, resources are being consumed. Therefore, some domains might be willing to route some types of traffic but not others. We discuss this interaction using a formal approach using first-order logic [3].

2 ISP-customer problem space and policy issues

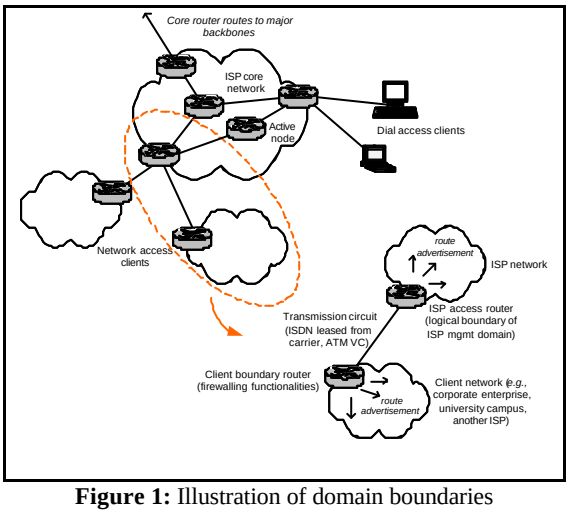

Figure 1: Illustration of domain boundaries

The ISP offers its customer a default route, and the ISP has the responsibility for announcing the customer's route to all other customers of the ISP and to all external connections. The policies that the customer network wishes to represent to the Internet through the ISP should be exactly aligned to the service provider's default case policies. The ISP's firewall rules should be aligned because the gateways of a domain may act in their own behalf or on behalf of other ISPs. As such ISPs and gateways must trust one another to different extents [4].

Due to the dependence between ISPs, end-to-end QoS is difficult to be improved beyond the quality provided by their peers. The ability to manage end-to-end services calls for traffic conditioning and admission control at network edges and resource provisioning at network nodes (*e.g.*, PHB). Should routing domains trust each other since algorithms might not be implemented properly and bugs might propagate? Bigger ISPs tend to make it difficult for the smaller ISPs to peer with them because of the concern over the smaller ISPs' ability to properly administer a safe BGP peering connection.

In addition, through multi-homing (*i.e.*, using two or more upstream ISPs) allows the ISP to switch its traffic from one to the other in the event of routing or connectivity failures on any single upstream ISP. This

[1] In this paper, we refer to organisations that operate networks as 'network operators'. We divide network operators into two categories: service providers and customers. A service provider is a network operator who operates a network to provide Internet services to different organisations, *i.e.*, its customers [1]. For example, JANET acts as a service provider to universities but it also buys (hence acts as a customer) international connectivity from another service provider.

[2] Currently, the scope of the Internet business can be subdivided into three major areas: ISPs that provide access services; ISPs that provide transit or backbone connectivity; and ISPs that provide value-added services

presents an approach that allows the ISPs to continually engineer its traffic flows between upstream providers in order to minimise the total costs of the upstream service. Each ISP would prefer to manage its own network resources and enforce its own internal traffic engineering policies. There should be an inter-domain reservation system that uses these delivery commitments to establish a reservation path through multiple domains. There is a need for an inter-domain management systems protocol between ISPs that includes resource advertising to complement route advertising as offered by the Border Gateway Protocol (BGP).

3 Policies in a multi-user system

There ought to be at least two different levels of granularity for these policies, *i.e.*, the need for a semantic translation from policies to firewall rules, BGP Loc-RIB, and/or IPsec filter rules is important for effectuating the intended actions [6]. The same concept applies to allocation policies. At node level, the resource 'currency' would include processor cycles, buffer spaces and link capacity. At network level, resource allocation decisions should consider parameters that would contribute to the quality of end-to-end connectivity.

Consider a set of customers vying to reserve resources on an active node, owned by an ISP, in order to further provision service to their end-users. The following formalised policies represent our first iteration of inter-domain management protocols.

Definition 1: The set $S = \{s_1, s_2, s_3, \ldots, s_n\}$ represents the general set of users who intend to access systems across administrative domains $D = \{d_1, d_2, d_3, \ldots, d_n\}$ that lie across the Internet. There also exists a set of privilege users who owns the nodes. We identify them as ISPs ? $= \{f_1, f_2, f_3, \ldots, f_n\}$, and it comes naturally that ISPs perform administrative functions a(d) on their domains. It also follows that there is exactly a single ISP per administrative domain. Formally, $\forall d \exists !f(a(d))$.

Definition 2: Within a single administrative domain, we note a subset of the general group of customers, who have satisfied a certain admission or access control policy to enable the use of its node resources; and are, thus, known as authorised entities denoted by $A = \{a_1, a_2, a_3, \ldots, a_n\}$, where $A \subseteq S$, and $S \backslash ?$ effectively gives $\{x \mid x$ represents all the contending customers$\}$.

3.1 Admission policies

No policy: *any entity may access system without discrimination*. Formally,

$$\forall x(x \in S \rightarrow x \in A) \tag{3.1.1}$$

Tickets: *any entity may access system if it has a requisite credential*. The credential is awarded by the owner of the system (root) based on a quality the entity possesses (*e.g.*, paid subscription, barter trades on network resources).

Definition 3: $C(p, q)$ represents a credential given by p to q. Formally,

$$\forall x(x \in ? \ \vee \ C(?, x) \rightarrow x \in A) \tag{3.1.2}$$

Sponsor/guarantor: *any entity may access system if at least another single entity within the system can guarantee its credibility*. Obviously, the entity within the system knows the potential new entrant, and can vouch for the safety of the latter's operations within the system. Formally,

Definition 4: $F(p, q)$ represents q as p's alliance, e.g., two organisations with prior bi-lateral agreements.

$$\forall x(\exists y(F(x,y) \ \wedge \ y \in A \ \wedge \ y?x) \rightarrow x \in A) \tag{3.1.3}$$

Majority: *any entity may access system if a certain percentage r of the existing entities within the system at that particular instance allow it.* These entities may or may not know the potential new entrant but make decisions (to allow or deny) based on their satisfaction of current resource usage (*i.e.*, refer to allocation policies below to reason about the semantics of negotiation protocols).

Definition 5: The satisfaction function s(p) returns a Boolean result based upon the satisfaction level of entity p. Formally,

$$\forall x(\exists y(\sum_{n=1}^{i} y > r * n\{A\} \ \wedge \ s(y) \ \wedge \ y \in A) \rightarrow x \in A), \text{ where } 0<r<1 \tag{3.1.4}$$

It may be fairly argued here that an entity that gains initial access to the system obtains *de facto* control, and thereby the benefit it can secure for itself. As such, apart from having a formal notation to identify the order of access, we also introduce additional restriction on these voting rights:

Limited voting rights: *any entity that has already gained access to the system may vote to allow or deny a new entrant if it has not already done so for w number of times.* Note that this is not an access control policy but rather a constraint enforced upon an entity's voting rights.

Definition 6: $V(p, q)$ denotes the number of times p has voted on the decision on q's access permission. If q is left empty, i.e., $V(p)$, the default notation denotes the number of times p has voted in total; and $?$ denotes the right to vote in the next decision making process. Formally,

$$\forall x\,(x \in A \wedge V(p) < w \rightarrow x \in ?), \text{ where } w \in \boldsymbol{Z}^{+} \tag{3.1.5}$$

The motivation here is to restrain reckless denial of access and, at the same time, ensure that current entities within the system can maintain the right to negotiate and carefully decide if their interests and satisfaction would be compromised in view of the impending new entrant to the system.

3.2 Allocation policies

In resolving resource contention, the ISP must recognise what is akin to the 'theory of unlimited territorial integrity' that forbids a country to alter the natural conditions of its own territory to the disadvantage of a neighbouring country [7]. On the other hand, according to the Harmon doctrine[3], which advocates the 'theory of absolute territorial sovereignty', where a country has absolute sovereignty over the area of any river basin in its territory. Evidently, one can foresee these doctrines in conflict, just as how the ISPs has to ensure that its resource allocation policies are foolproof and will not compromise a customer's interest in riposte to another customer's benefit.

Definition 7: Physical resources ? can be divided between bandwidth $B=\{b \mid b$ represents units of bandwidth$\}$, processor cycles $P=\{p \mid p$ represents units of CPU cycles$\}$, and memory $M=\{m \mid m$ represents units of memory space$\}$. Formally, $B \subseteq ?$, $P \subseteq ?$, $M \subseteq ?$, and $B \wedge P \wedge M = \emptyset$.

The constraint[4] for resources is represented by subscript notation, *viz.*, $b_{max} \in B$ indicates the total bandwidth that can be provided per node, $p_{max} \in P$ indicates the total processor cycles that can be achieved by the node, and $m_{max} \in M$ indicates the total disk space per node. The temporal limit for the duration of resource usage is represented by the function $?_{max}$.

Definition 8: Resources are further categorised into total resources that are still available $?_{avail}$ and the total resource that have been allocated $?_{alloc}$, where $?_{avail} \wedge ?_{alloc} = \emptyset$, and $?_{max} = \{?_{alloc}, ?_{avail}\}$ such that $(?_{alloc} \rightarrow 0) \rightarrow (?_{avail} \rightarrow ?_{max})$

Basic limits: For the lower boundary, *no entity can buy less[5] then than x amount of memory, y amount of CPU cycles and z amount of bandwidth.* For temporal limit on each type of resource reservation, *all entities must reserve a resource within a range of time limits.* For the upper boundary, *no entity can be allocated more than the total amount of available memory, CPU cycles and bandwidth.*

Definition 9: The parameters for the *allocate* function are defined as ?(p, q, u, t) which describes a network operator $q \in ?$ allocating an amount of resource $u \in \{? \mid ?$ *represents the available resource*$\}$ to entity $p \in A$ for a duration of $t \in \{? \mid ?$ *represents the available duration for resource reservation*$\}$. Formally,

$$\forall x \in A\ \forall y \in ?\ (?(x, y, ?, ?) \rightarrow ((?_{min} \leq ? \leq ?_{avail}) \wedge (?_{min} \leq ? \leq ?_{max}))) \tag{3.2.1}$$

Minimum allocation: *Entities will always be allocated the least possible quantity of available resources.* Formally,

$$\forall x \in A\ \forall y \in ?\ (?(x, y, ?, ?) \rightarrow (\neg\exists\, ?'(?' \leq ?)), \tag{3.2.2}$$

where $?'$ *= alternative amount of resource to be allocated*

History: *No entity may request more than a certain percentage of its prior highest reservation parameter.* Formally,

$$\forall x \in A\ \forall y \in ?\ (?(x, y, ?, ?) \rightarrow (? \leq rh(?))) \tag{3.2.3}$$

where $h(?)$ *= prior amount requested for resource* $?$*, and* $0 < r < 1$

---

[3] First authoritatively stated by Judson Harmon, an American Attorney-General who made the declaration concerning the Rio Grande.

[4] Representation of constraints without the '*max*' subscript indicates instantaneous values. Intuitively, adding the '*min*' subscript indicates minimum values for resource and period of subscription.

[5] The minimum limit for resource requestors is imposed from a business aspect, whereby the revenue from resources provisioned should, at least, justify the cost (*e.g.*, maintenance and management costs) of provisioning in order to satisfy the break-event point.

Pareto-optimal[6]: *No entity can reserve so much resource such that the next request coming from any other entity is denied of taking equally as much*. In other words, the remaining resource available should be at least equal or more than what is to be reserved. Intuitively,

$$\forall x \in A \; \forall y \in ? \;\; (?(x, y, ?, ?) \rightarrow (?_{avail} \geq 2?)) \qquad (3.2.4)$$

However, cynics of the Pareto Optimal Principal argue that you cannot make anybody better off without making someone else worse off, *viz*., if b_{avail}=30Mbps and service provider 1 (ISP_1) is allocated 15Mbps of bandwidth (in line with the intuitive rule), the next provider ISP_2 that comes along with a request for 15 Mbps is denied from taking the remaining 15Mbps. The rule would forbid such a request because it would mean that the subsequent requestor ISP_3 is not be able to obtain 15Mbps of bandwidth if ISP_2 gets the remaining 15Mbps.

Thus we observe here that the same rule that gives the green light to ISP_1 (thinking that ISP_2 would be unaffected) later denies ISP_2 in consideration of ISP_3. The interim conclusion achieved is that, 'unless we know in advance the total number of potential customers to be provisioned per node, we cannot effectuate the Pareto Optimal Principal as an allocation policy'.

Ideal aspiration: *an entity is allowed to reserve the maximal available resource if no other entities access the system during that period of time*. Thus, in the absence of other entities, the lone entity would experience an *'ideal aspiration'* level.

$$\forall x \in A \; \forall y \in ? \;\; ((?(x, y, ?, ?) \wedge \bar{x} = 0) \rightarrow (? = ?_{max})) \qquad (3.2.5)$$

However, consistent to other resource sharing predicaments, it is impossible to guarantee every entity its *'ideal aspiration'* level as more entities gain access to a system.

Anti greed: *No entity may reserve resources that it cannot use* (*e.g*., does not have much traffic to justify a huge bandwidth reservation). Admittedly, this is the least important in terms of priority since the rationale in a commercial sense should be, *'as long as you can pay, join the club'*!

$$\forall x \in A \; \forall y \in ? \;\; ((?(x, y, ?, ?) \wedge s(x) = 0) \rightarrow (? = 0)) \qquad (3.2.6)$$

4 Conclusion and future work

We described the complexities inherent in multi-domain management. We discussed the complexity of accessing resources in other administrative domains and this complexity impedes provisioning of end-to-end services. We have also presented the first iteration of formal description for admission and allocation policies. It is expected that logical representations can help derive logical consequences of the policy rules and therefore test them before they are put in force. Formal method is chosen over other representational schemes because it has a sound and rigorous theoretical framework.

As part of future work, we hope to represent intentional notations (*e.g*., trust and belief), in which the standard substitution rules and truth functional property of first-order logic do not apply. Modal logic is the alternative formalism approach for possible world semantics, while sentential and first order logic (predominantly set theory in this paper) are used to categorise the network domain constituents at a reasonable level of abstraction.

5 References

[1] Bates, T., Gerich, E., Joncheray, L., Jouanigot, J-M., Karrenberg, D., Terpstra, M., Yu, J., **'Representation of IP routing policies in a routing registry'**, *RFC 1786*, March 1995 (http://www.ietf.org/rfc/rfc1786.txt)

[2] Flores-Mendez, R. A., **'Towards a standardisation of multi-agent system frameworks'**, *ACM Crossroads*, Issue 5.4, 1999

[3] Enderton, H. B., **'A mathematical introduction to logic'**, 2nd Edition, *Hardcourt/Academic Press*, 2001

[4] Abadi, M., Burrows, M., Lampson, B., Plotkin, G., **'A calculus for access control for distributed systems'**, *SRC Research Report 70*, Palo Alto, California, August 1991 (ftp://gatekeeper.research.compaq.com/pub/DEC/SRC/research-reports/SRC-070.pdf)

[5] Huston, G., **'Next steps for the IP QoS architecture'**, *RFC 2990*, November 2000 (http://www.ietf.org/rfc/rfc2990.txt)

[6] Patz, G., Condell, M., Krishnan, R., Sanchez, L., **'Multidimensional security policy management for dynamic coalitions,'** *DARPA Information Survivability Conference and Exposition 2001 (DISCEX II)*, Jun 2001

[7] Kilgour, D. M., Dinar, A., '**Are stable agreements for sharing international river waters now possible?**', *Policy Research Working Paper No. 1474*, 1995

[6] Named after the work of Italian economist and sociologist Vilfredo Pareto (1848-1923), the *Principal of Pareto Optimality* is an evaluative principle that says, 'The community becomes better off if one individual becomes better off and none worse off'.