

The Direction of Value Flow in Connectionless Networks

Bob Briscoe<rbriscoe@jungle.bt.co.uk>

BT Research, B54/74, BT Labs, Martlesham Heath, Ipswich, IP5 3RE, England

21 Sep 1999

Abstract

This paper argues that all network providers in a connectionless multi-service network should offer each class of their service to each neighbour for each direction at a single price. This is called ‘split-edge pricing’. If sets of customers wish to reappportion their networking charges between themselves, this should be tackled end-to-end. Edge reappportionment should not be muddled with networking charges, as is the case in the telephony market. Avoiding the telephony approach is shown to offer full reappportionment flexibility, but avoids the otherwise inevitable network complexity, particularly for multicast. ‘Split-edge pricing’ is recursive, applying as much to relationships between providers as to edge-customers. Various scenarios are discussed, showing the advantage of the approach. These include phone to Internet gateways and even inter-domain multicast conferences with heterogeneous QoS. The business model analysis suggests a new, purely financial role of end-to-end intermediary in the Internet industry.

Keywords: Charging, pricing, clearing, end-to-end, multi-service, multicast, connectionless, business models.

1 Introduction

Traditionally, data communications has been sold so cheaply that charging for it on usage basis has not seemed feasible or sensible. While flat-rate subscription or connect-time charging prevails, the question of reappportioning the value of a particular communication between its ends rarely surfaces. With the possibility of variable quality of service (QoS) approaching, the need for some form of usage-charging for high QoS has arisen. This has led to new thinking on cheaper usage-charging systems for packet networks [17, 3], which in turn has brought the issue of reappportionment of charges between the end customers back into the limelight [6]. However, to tackle this issue, the traditional model

of a unit of communication value is inadequate because it grew in the context of connection-oriented telephone systems. Worse, it has often been misapplied to the connectionless Internet.

The connection-oriented mind-set also leads to confusion over blame and liability for each unit of communication. This paper explicitly clarifies these fundamental issues. Our minimalist model for a connectionless network business has boundaries that match the service access points above and below the network layer of the OSI stack [22]. This business can be modelled as buying in lower and higher layer services (e.g. links, virtual connections or naming services). It still applies to ISPs that use their own links and services — the cost just becomes internalised.

This paper proposes a simple charging model that can be applied between any pair of multi-service connectionless networks for each class of service and for send and receive separately. It works whether the pair are both providers or even if one is an edge customer. The model’s simplicity ensures charging will always be straightforward at every border in the Internet, whether for unicast or multicast flows. No matter how many networks are connected together, any one network is only dependent on prices from its direct neighbours. Therefore, the model is intrinsically scalable. We call the model ‘split edge pricing’. From the end customers’ points of view, this means that any flow through the Internet is sold on entry and on exit. As a consequence, the model *appears* to restrict all end customers to each have to pay the price of their local network provider. This appears to restrict any customers who would rather reappportion the costs differently between themselves (termed clearing).

Invariably, network providers offer their services at a set price regardless of the value each customer derives from each transmission. This is a natural consequence of a competitive market, often called a ‘buyer’s market’. If usage-based charging is in operation, no one bothers with any communication of less value than this market price. However, transmissions naturally have at least two ends. (In fact, we consider two-ended flows as just a specific

case of multipoint flows.) Often a transmission just never happens because one of the ends derives less value than their local price. Often in such cases, the total value derived from the transmission by all ends would have been greater than the total charges levied by all providers on all end customers. Therefore, it is in any network provider's interest to matchmake customers who derive surplus value with those who would otherwise be in deficit. That is, clearing plays an important role in encouraging network demand.

Matchmaking in the traditional telephony market is well understood. Various ways are available for end customers to share the cost of a call besides the normal 'originator pays'. Examples are 'calls free to the originator', 'local charge only' etc. Telephony interconnect arrangements ensure that wherever payment enters the system, it ends up being cleared between the providers who bore the cost of each call. However, the interconnect pricing scheme that drives clearing blurs the distinction between clearing of edge payments and the market price of interconnect. This paper argues that these over-complicated clearing arrangements are the result of evolution from a fully connected matrix of single country providers and are flawed for the Internet. Instead we propose 'split-edge pricing' as a more flexible replacement. The apparent problem of no flexibility to clear between the ends is solved simply. Clearing can be achieved end-to-end, directly between customers or their edge providers, bypassing the core network businesses. If instead, clearing follows the same path as the data flow, we show that core network complexity becomes inevitable, particularly for multicast, but also for unicast. Incidentally, end-to-end clearing was never possible on the PSTN because there was no convenient way to form independently routed end-to-end data connections simultaneously to call progress. Clearly, this *is* possible on the Internet.

Clearing requirements will differ on a per session basis, therefore the model where clearing takes place end-to-end involves per-session accounting without involving the network providers along the data path (other than those at the edge). Thus, it seems natural for clearing to re-use existing e-commerce concepts and mechanisms. This results in a scenario where the traditional telephone bill becomes an anachronism for the Internet. Instead, edge-provider charges can be settled by any other third party across the ends of a communication, leaving the 'bill' as just the balance of those charges that are directly retailed between the edge-provider and its edge-customer. E-commerce-based clearing allows part of local customer A's usage to be 'wholesaled' to remote customer B or to third party C. While part of customer B's usage can be wholesaled

to customer A and so on.

The cost of the act of clearing is significant; therefore it is important that the default apportionment in the core model matches the most common case. We establish that the common case is where both senders and receivers pay, at all ends of each transmission. This incidentally causes the perceived need for the network to report how many receivers are subscribed to a multicast to evaporate. We then consider the unpleasant fact that, on the Internet, a receiver can never protect itself from being sent to. We suggest a rather novel business model that is still optimised for the common case, but simultaneously has no receiver liability. For completeness, we examine the specific problems with the traditional clearing model used in telephony. The paper draws to a close by working through some example scenarios to suggest how the models would work in practice. Finally, limitations and further work are listed before conclusions are drawn.

2 Related Work

Some authors state that they believe the business model of the current fixed access rate Internet is 'sender takes all' [11, 23]. This phrase is used to imply the sender's ISP receives all the revenue. This is completely erroneous. ISP's rates relate to access bandwidth, regardless of the direction in which it is used. Thus, 'sender and receiver each take half' is more appropriate (approximately). This is a similar position to the half-circuit charging common for data links, but applied end-to-end. MacKie-Mason *et al* asserts that the blame for a transmission is impossible to determine at the network level [15], an argument that can descend into sophistry. However, later, using precise definitions of the terms, we argue that the sender is always to blame for a transmission in a connectionless network.

Clark analyses the apportionment of charges between senders and receivers [6], and proposes an engineering solution, which he admits would introduce considerable complexity to the Internet if implemented. Shenker *et al* describes edge pricing [18], a business model that appears regularly in communication networks and which forms much of the background to this work.

3 The value of place

The value of communication concerns the incremental value of having information in a certain place (or places) by a certain time, instead of or as well as the original places. Usually, the more

the information is worth, the more value is placed on having it in the right places in a timely manner. There is no value to the customer at all while information is in transit. It is delivery that is important. Strictly one also has to take account of the mitigating cost of storage in both places (or only in the second place if the sender deletes after sending). In summary the added value of transmission is the marginal change in value caused by associating a new location at the delivery time with the intrinsic value of the information to the customer.

However, because the data communications market is fairly competitive, charges for communicating information tend instead to follow the ‘cost plus margin’ rule. This is particularly so because it is very difficult for providers to predict what value their customers put on moving any one piece of information.

Any payment to an edge-network provider has the two aspects — ‘who pays’ and ‘who is paid’. ‘Who is paid’ can only be each local provider collecting its local price. With competitive ‘cost plus’ pricing there is no scope for any provider to break out of that. But, because communications naturally involves at least two parties, in order to cover the total costs of all the providers involved, ‘who pays’ can be on a different apportionment. The edge customers *do* know the value to them of having the information at a certain place in time. Thus, although apportionment is difficult for network providers, it is very relevant to edge-customers. Clearly, the network providers can stimulate more use of their networks by making arrangements for customers to efficiently apportion costs between themselves.

4 End-to-end pricing

If a price is higher than the perceived value for any customer, she is free to get the remote party (or anyone else) to make up the difference through some higher level arrangement. On the other hand, if the value to her is higher than her local price, she is also free to offer to cover some of the costs of the remote end(s). However, our minimalist provider doesn’t have to be concerned with matchmaking multiple customers to get round local discrepancies between price and customer value. This is an issue that can be dealt with end-to-end, not locally. We are not saying ISPs shouldn’t offer end-to-end pricing — it is clearly in their interest to matchmake between customers with surplus value and those with deficit. All we are saying is that, if they do, end-to-end pricing should be considered as a separate role (Fig 1). Such a role could be a separate business — it could gain on some combinations and lose on

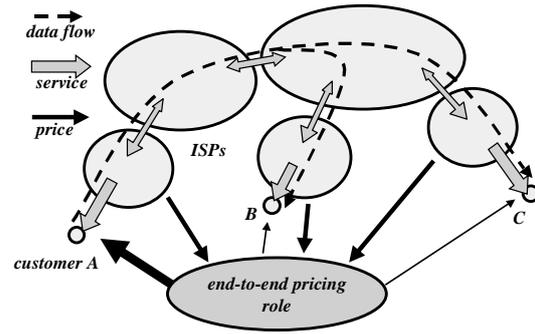


Figure 1: ‘End-to-end pricing’ role

others, possibly making a profit overall. In this case it would be a retail service that used the networking services as wholesalers. It is also possible that edge customers could effectively take on this role themselves. Fig 1 shows three end customers using a data path through multiple connected ISPs. The relative value of the service flows and prices for one direction of one class of service is represented by the thickness of the arrows. Note that the size of the proportions of prices represents a choice by the end-system that is willing to pay more than its local price. In fact, the end-to-end pricing role may advertise identical prices to each customer. However, these could be modified by an offer by A to cover a proportion (possibly all of it). Pricing between providers is omitted for clarity (but see later).

Telephony firms have traditionally offered end-to-end pricing because they are selling an application. The role of network provider has always been muddled with selling the end-to-end application. This is already putting considerable strains on the International Accounting Rate System (IARS) [13] with potentially $s(n-1)^2$ prices having to be negotiated (where n is the number of edge providers and s is the number of global schemes for sharing the proportions of the price between the ends, e.g. local rate only, free to sender). In practice, end providers are grouped together to reduce the number of prices presented to customers. The PSTN uses addressing conventions (e.g. +800 for free to sender), but this limits commercial flexibility to the few schemes that are widely recognised. Clark proposed an Internet-based solution to allow flexibility [6]. However, catering for various combinations of sender and receiver payments through the core of the network needs packet format changes and router involvement. Further, wholesale prices between providers would have to be negotiated for every possible scheme for sharing charges between ends as well as for every possible grouping of end points beyond that boundary. Worse still, inter-

provider accounting would then require traffic flows to be isolated then further sub-classified by how much each end was paying on a per-flow basis.

The ‘ n^2 problem’ would still exist for our end-to-end pricing solution but this is fairly easy to contain by grouping. An example scenario is given at the end of the paper. Importantly though, end-to-end pricing gets rid of all the inter-provider problems described above. There is no longer a need to identify end-to-end flows at inter-provider boundaries. Thus inter-provider charging could be based on bulk measures like average queue lengths, number of routing advertisements etc. Also, most importantly, end-to-end pricing can be introduced without changing the Internet at all, and it allows future flexibility. To summarise so far, we should ensure any discrepancy in the willingness to pay across end customers is normalised end-to-end first, so that edge ISPs always receive payment at their local price.

5 Common case value apportionment

Although we have delegated the problem of apportioning sender and receiver payments to a higher layer, it is still important to cater for the common case at the network charging level so that the higher layer functions are unnecessary in most cases. We propose that all edge providers should charge their local customers for both sending and receiving as the default case. Our primary justification for this stance is that the large majority of communication occurs between consenting parties. In this section we also argue that other possible default scenarios (e.g. ‘only senders pay’) would be unstable anyway and collapse back to our proposed model, which appears stable. Further, delegating charge reapportionment to a higher layer eliminates all but local pricing, so that we can extend charging for sending and receiving recursively to apply at the boundary between any pair of providers. This greatly simplifies inter-provider metering — essential for Internet scalability as QoS and multicast are introduced. Thus our edge pricing model applies to any ‘edge’ — whether at the edge of the Internet or just the edge of a backbone. The stability analysis intrinsically applies equally to this more general case. Incidentally, *allowing* for prices for each direction to be different hardly needs justification. It allows for asymmetric costs (e.g. access technology like xDSL or satellite) and for asymmetric demand (e.g. some ISPs might host more big senders, while others might host the mass of receivers). If these factors aren’t asymmetric, the two prices can simply be set to be the same.

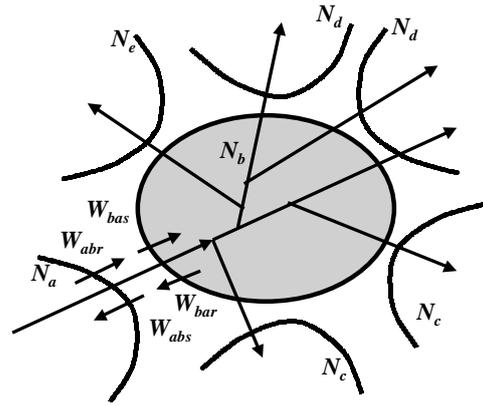


Figure 2: Split-edge pricing

Fig 2 shows a generic scenario with multiple networks, N , all connected to the network of interest, N_b . Each connected network has a status relative to N_b based on whether it provides more or less connectivity to other hosts at that class of service. Although the diagram gives the impression that N_b is a backbone network, any one of the neighbouring networks could be a simple link to an edge customer’s single host. The model is designed to be general enough for N_b to be an edge customer, an edge network, a backbone network or some hybrid. Those networks with the same suffix are of similar status relative to N_b . For instance, those labelled N_c may be edge customers, N_d may be equally large backbones and N_e a peer network.

In fact, this is a simplification. To be more specific we propose that a provider should offer each class of service in each direction at a separate price. Thus, Fig 2 shows the situation for one of possibly many classes of service. Class of service is defined as a unique combination of the service mode (unicast, multicast) and quality (latency, instantaneous bandwidth, reliability, jitter). Quality specifications within one class may leave one parameter to be specified by the customer while others remain fixed, thus generalising both RSVP and diff-serv [21, 1]. Appendix A justifies treating each class of service independently. Appendix B gives an introduction to the model for each class of service, but allows heterogeneous QoS per leg of the multicast. The full model is given in an earlier version of this paper [2]. However, all this detail would obscure the summary of the analysis we attempt to give here, which we now continue.

A packet of a particular class of service is shown being multicast from N_a into N_b and onward into the other networks. Because multicast is a general case of unicast this allows us to model both topologies. We will also be able to treat the topology as aggregation¹ by reversing the direction of trans-

¹Examples of packets that are forwarded until aggrega-

mission. The term packet is used, but the arrows could represent flows of similar class packets for a certain time. Fig 2 highlights the pricing between networks N_a and N_b . W_{bas} and W_{bar} denote the per direction weightings applied to the 'nominal charge' that N_b applies to N_a (for more detail on exactly what the nominal charge means, see Appendix B). W_{abs} and W_{abr} likewise weight the charge N_a applies to N_b . Each weighted price is for transmission between the edge in question and the remote edge of the Internet, not just the remote edge of that provider. For full generality, there have to be four price weightings like this for every class of service at every inter-network interface, but the weights would take different values unless the neighbours were of the same status. The relationship between any two parties across the edge of their networks is split into prices for each class of service and each of these is further split into two prices for each direction, each of which are again split into 'half' prices that each party offers the other. Hence, we call this model 'split-edge pricing'.

Thus the payment for traffic in any one direction across each interface depends on the difference between the two weighted prices offered by the networks either side. In other words, no assumptions are made about who is provider and who is customer; this purely depends on the sign of the difference between the charges at any one time. Clearly, edge customers (N_c , say) have no provider status in the networking market. So, for all j , $W_{cjs} = 0$ and $W_{cjr} = 0$.

In Appendix B we analyse policies like 'only senders pay' or 'only receivers pay' using the model (by simply setting all receiving weights to zero or all sending weights to zero). Stability of a policy is determined by assessing whether one network would gain from a maverick policy. The results are summarised here.

'Only senders pay' or 'only receivers pay' are only stable policies if all providers agree to adopt the same policy, and none break ranks. As soon as one goes maverick, customers who are primarily receivers and those who are primarily senders migrate to different providers. Income appears to remain stable, but the source of the income switches from retail customers to interconnect causing the inter-provider link to become a bottleneck. Thus costs increase without any increase in revenue.

'Only senders pay' is also unstable where multicast

tion (reverse multicast or 'concast' — converging cast) are:

- RSVP receiver initiated reservation (RESV) messages;
- pragmatic general multicast (PGM) [19] negative acknowledgement (NACK) messages or the 'lay breadcrumb' messages [7] suggested in their place.

is concerned. ('Only receivers pay' is all but meaningless for multicast.) To support an 'only multicast senders pay' policy, all domains have to trust each other to faithfully report receiver community size. In Appendix B we show that it is simple for a domain to lie about its local receiver community size to increase its profits. Proposed mechanisms such as 'EXPRESS count management protocol' [10] suffer from this flaw. Solving this problem is unlikely to be successful without breaking the scalability benefits of receiver initiated IP multicast that ensure upstream nodes are unaware of downstream join and leave activity.

In contrast, 'both senders and receivers pay' is stable in both unicast and multicast cases. It also doesn't lead to inefficient network utilisation unlike the above cases. It is also possible to cater for different balances of predominant senders and receivers by weighting the sending price differently to the receiving price. For instance if there are a few big predominant senders but many small predominant receivers, the economy of scale in managing a large customer can be reflected in a lower sender weighting. Similarly, the inefficiencies of multicasts to small receiver communities compared to multiple unicasts can be discouraged by slightly weighting multicast sender pricing. The aggregation case is similar, with 'both senders and receivers pay' stable while the two other policies go unstable for the same reasons as for multicast, but swapped round.

If end customers want a different apportionment of charges, we have made the case for this being arranged end-to-end. The remainder of the paper concentrates on issues surrounding clearing. Also, at the end, various worked examples are given to illustrate how it would be achieved in practice. However, first, we introduce one further relevant issue — that of how a receiver can control its costs, if it can't stop itself being sent to.

6 Blame, liability and control

We have shown that all ends paying is the common case and a stable one so should be the default. We can share the cost differently at higher level if end user value is shared differently from this default (and it is worth bothering given the cost of another financial transfer). However, we must remember that a sender can decide not to send but a receiver can not avoid being sent to (in the current Internet). We must be careful here to define the context of the question of blame. We are only concerned with blame for sending into or receiving from the service access point of the network layer. Clearly, if someone operates a Web service, they don't normally decide whether to send replies on

a request by request basis. But this doesn't mean they have been forced to send *at the network level*. They have chosen to put the service on a well-known port with public access. They can stop certain people requesting them to send by securing the Web server or interposing a firewall. But, whenever they send it is because they have arranged it to be so.

Ultimate sender blame presents a problem. In cases where the sender derives surplus value from a communication and the receiver derives less value than their provider charges, receivers are vulnerable to being exploited (e.g. adverts). Such cases are much rarer than it first appears, mainly because of confusions that can be cleared by considering the following factors:

- The value of the information isn't relevant when considering the networking service — only the value of *moving* the information — getting it to a useful place
- Often the value of moving information is transitory — getting it to a useful place to discover that moving it wasn't useful
- Often the value of moving lots of information is to get a small part of it to a useful place, but it isn't possible to know which part before moving it
- The cost of transmitting information is often far less than the cost of the effort of targeting which information should be transmitted
- Information in one direction often controls the flow of information in the other

Nonetheless, genuine cases remain where the receiver is being persistently forced to pay for transmission that is valuable to the sender but not to the receiver. The only solution to this seemingly intractable dilemma is for it to be *customary* for all ends to pay, but the ultimate liability should remain with the sender. Any receiver could then dispute the customary apportionment (end-to-end) with no risk of denial (unless the sender had proof of a receiver request). A similar but opposite situation used to prevail with the UK postal service. It was customary for the sender to pay for the stamp, but if it was missing or insufficient the receiver was liable for the payment, because the Royal Mail had an obligation to deliver every letter.

7 End-to-end clearing

We have discussed how prices can be apportioned between the ends of a communication. We now discuss how payment will follow the same path. We

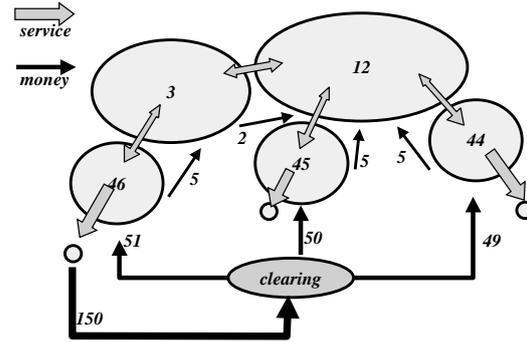


Figure 3: 'End-to-end clearing' model

can assume electronic commerce will make it possible for anyone to pay anyone else's ISP on the Internet, even if a clearinghouse is needed. We shall call this the 'end-to-end clearing' model (Fig 3). These arrangements will typically be made through higher level protocols. The act of making a financial transfer has similar order costs to the cost of transmission of a couple of small e-mails. In addition there is the cost to the ISP of provision of processing resources for authentication. Therefore, arranging a different apportionment of charges between ends is more likely for long-lived sessions, such as Internet telephony or conferences on the mbone, than short connections, such as are typical on the Web. However, a collection of related short connections may be combined into one longer-lived session for these purposes.

In the 'end-to-end clearing' model, the clearinghouse role deals with the end-to-end 'half-circuit' sharing (including the straightforward price differences between the two ends) leaving inter-provider accounting to be purely about wholesaling. The figure shows one end paying and follows example proportions of this money as they are distributed among the providers. The clearing role may be involved in taking payments for higher level service from each customer (e.g. conference fees or pay-TV charges). In such cases it knows the number of participants and can charge the customer on the left (who is paying for everyone) accordingly. In this example each leg is charged at fifty units and no profit is made. Note that inter-provider money flows match the flow of networking service in the opposite direction across the same interface. This is the deliberate aim of the model, to ensure that bulk measurements at these boundaries can drive interconnect charges based solely on local conditions. It also allows each pair of providers to choose their own basis for metering independently of other arrangements at other interfaces.

There is nothing to stop providers or customers

assuming the clearinghouse role, but the accounting information model needs to be based on a third party clearing system to allow for the most general case. To clarify, the paying customer may make payment:

- either to a dedicated clearing house
- or direct to the ISP at the remote end (the remote customer need only notify her ISP's payment interface to the payer)
- or even direct to the remote customer so that she can pay their own ISP

In all cases, the *role* of clearing must be separate even if there is no separate enterprise to achieve the function. Note that the last case is special — the clearing role is null, but it still appears in the information model. In other words, the charges for all ends should never be lumped together while accounting. If, instead, end-to-end half-circuit sharing were achieved through the provider chain, end-to-end clearing information would have to be identified separately from that needed for wholesale accounting. If clearing information were not identified separately, the types of model that could be built on the infrastructure would be restricted.

8 Iterative clearing

We have presented what we believe to be an optimum business model, but other models need to be considered. In particular, we will now consider a model similar to the public 'phone service, which has one or two implicit features that need to be separated out for full understanding. We will consider payment in the model first, rather than pricing, as it will then be easier to understand the pricing issues.

In this model, ISPs don't expect payment for all sent and received traffic to be made to *all* edge providers (Fig 4). Instead a customer might pay their own provider on behalf of both (all) ends as in the normal case for telephony and as proposed by Clark for the Internet [6]. This alternative business model allows customers to decide into which end(s) payment enters the system, on a per flow basis. We shall call this model the 'iterative' model for reasons that will become clear as we go. The financial flows between providers in this model depend on at which ends payment is entering the system on a per flow (or per packet) basis. For some flows, there may even be proportional sharing of costs between the ends. For business model flexibility an accounting system would need a 'payee percentage' field — the percentage of the total cost to be paid by the

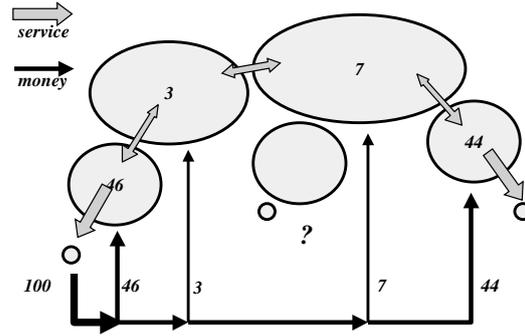


Figure 4: 'Iterative clearing' model

customer at the end being accounted for. Usually it would be 100% or 0% in the typical cases of 'paid completely to local provider' or 'completely to remote'. The balance would be the remote end's payment. Note, though, that the perceived purpose of this model is the transaction efficiency when the local payee gets 100%.

If Fig 4 is compared with the end-to-end clearing model in Fig 3, both models end up with two of the edge ISPs paid the same amounts on a half-circuit basis (we will explain where the third leg of the multicast has gone later). The difference is merely in the route the payment takes from payer to payee. With 'iterative' clearing the payment follows the data path. Along the way, providers take their cut with two types of money sharing being mixed together:

- wholesale cut
- half-circuit sharing

However, the amount deducted from the flow at each boundary doesn't match the level of service crossing that boundary. This can lead to complexity in the network, as there is pressure to design the network itself to reveal the apportionment of costs. This was why Clark was concerned about how much complexity would be added to the Internet to cater for arbitrary combinations of sender and receiver payments. This is also why international and interconnect on the PSTN have limited flexibility to arbitrarily apportion charges between the ends. Even 'free to sender' calls are blocked between a lot of countries because they don't yet have prices set or the interconnect accounting in place. Specifically, there are five points stacked up against the 'iterative clearing' model:

- As already pointed out, a 'payee percentage' field would have to drive inter-provider accounting, whether it was in accounting messages or packets. Otherwise the revenue of an

edge ISP and its upstream providers would depend on a factor completely outside their control — to which end its customers chose to make payment. The ‘payee percentage’ field would therefore have to be trusted by upstream providers. To help prevent the field being tampered with, it would need to be signed by the remote ISP. How signed fields can be aggregated without losing the signature integrity would be a matter for further research.

- Still further complication might be introduced for some future applications if the share of payment between the parties wasn’t fixed but depended on characteristics of the flow or other parameters only understood at a higher level — higher than the provider would normally be interested in.
- Worse still, the payment should ideally be split taking into account the current prices of all the edge providers who will eventually be paid. The only alternative (used in the international accounting rate system (IARS) for telephony) is for ISPs to agree compromise prices between themselves that average out price inconsistencies. This is what has been causing all the tensions in IARS as some countries liberalise earlier than others causing huge variation in prices around the world, between which no happy compromise can be found. This is difficult even for a system where every end to end path only passes through two international carriers at maximum, each pair setting compromise prices with each other. With eight ISPs on many end to end Internet paths, five typical [16] and considerable peer interconnection, it is likely that it will take longer to negotiate prices than the time available, thus leading to distortions to providers’ supply and demand signals.
- Finally, because of the much longer provider chains typically found on the Internet, unacceptable delays will be introduced before the revenue arrives in the correct place. Any delay in clearing hugely increases the cost of the payment system, as extra trust mechanisms have to be invoked while the payment remains unconfirmed. These trust mechanisms have to be applied to the edge customers, not just the providers, therefore hugely increasing the total cost of the system.
- If multicast is to be catered for by iterative clearing (e.g. conferences), each provider needs to know how many ends they are serving locally, both to inform the person paying and check settlement. In contrast, with the end-to-end clearing model, if senders and receivers

all pay, no-one needs to count ends nor trust others to count ends for them. If clearing is desired by the ends, only the ends need to know how many ends there are to pay for — no-one needs to calculate how many ends are attached to each provider. Thus, for instance, if there is a charge to join a conference, this can cover the cost of paying each participant’s communications charges as well as the content (each participant would have to declare their ISP when they join). The more who pay the host to join, the more there is to cover charges. Then the host can send bulk payments to the relevant ISPs, either directly or through a single clearer (see worked example below). Multicast has been omitted from Fig 4 simply because there is no generic way to handle multipoint networking with iterative clearing.

The only advantage of the ‘iterative’ model is that it appears to reduce (by one) the number of transactions to achieve the desired apportionment. Also all the inter-provider transactions can be fairly lightweight because they can be batched up. For example, consider the case where both the parties in an Internet ‘phone conversation are being paid for by the caller. It appears less complex for the caller to pay everyone’s payments to her own ISP, then let the ISP transfer the correct amount to its upstream provider as part of a bulk transaction. However, on the other side of the bargain is a considerably more complicated network, compromise pricing, increased credit time lags and less flexibility in inventing new ways to apportion charges, particularly for multicast.

9 Example scenarios

9.1 Finding an end-to-end price

Let us assume some way has been invented for an ISP’s edge customer, C_a , to announce her intention to cover some part of the transmission costs of parties communicating with her, C_b , C_c etc. Some suggestions are given in [6]. Kausar suggests modifications to SDP [8] to achieve this for longer sessions [14]. A price needs to be set and settlement made between C_a and each party. If this is achieved, end-to-end, between the parties involved there are no further engineering implications — the pairs of parties clearly trust each other enough to enter into a financial arrangement and are willing to accept the cost of the transaction. However, there will be many occasions where the parties have no trust relationship. In these cases the problem reduces to, C_b , C_c etc finding suitable

intermediaries. First they must know C_a 's ISP. C_a may have already given this information in the session description protocol. Alternatively a directory of ISPs could be operated by the Internet address allocation registry (IANA) or any private concern, in which one could look up the network address of C_a and be given the payment interface of the associated ISP. This directory might be operated by one organisation monolithically (feasible for the current 74,000 ISPs) or it could be hierarchical like DNS. They may then choose to check whether their own ISP has a direct relationship with C_a 's ISP. Alternatively, they could go straight to a different directory we postulate would be necessary. This directory would accept lists of ISPs and return a list of organisations that would act as an intermediary between them all. The mechanism would be identical to a Web search engine — an invert index of ISP-intermediary pairs that would accept queries of logically ANDed ISPs. There may be some intermediaries who will deal with any combination of ISPs through their own network of secondary clearing arrangements.

Whatever, the resulting intermediary could then be contacted to find the price being offered for the particular combination of ISPs. The same organisation would naturally take the payments and clear them between providers. The intermediary would also have to find out the prices being charged by the relevant edge-providers, which would represent its back-end costs. We assume the providers would be using tariff dissemination protocols such as in Rizzo *et al* [17], Carle *et al* [5] or Yemini *et al* [20] that could be listened to by the clearer as easily as by the edge-customers, particularly if transported over multicast.

9.2 Accounting and clearing with 'sender liable but local payment customary'

We concluded earlier that only the sender should be ultimately liable for usage charges. However, we suggested that the customary position should be to expect every customer to pay for both reception and sending. We will now describe how this necessary but rather novel business model would work, assuming also end-to-end clearing. At this point we have to make a clear separation between accounting and liability for payment. We propose that each edge customer and her edge network provider should first reconcile their usage records for sent and received traffic, whoever is expected to pay for that usage. That is, we require the edge ISP to be willing to sign the relevant accounts if asked. Whoever ends up paying, the price applied to each

account will be that set by the edge-provider supplying the service. We will also describe the more demanding case of post-payment.

The following discussion is easiest to understand if we consider the accounts for one flow at a time. We assume that accounting is operating in near-real-time on a per-session basis, thus such an approach makes sense as a session consists of a set of flows known to at least a subset of the participants. Again, for ease of understanding, let us consider one incoming flow to one customer before we consider her outgoing flow. Note that two or more edge-charges are raised per flow depending on the number of ends, but we are focusing on one customer at one end (termed 'local'). For brevity, we will use feminine pronouns for the local customer and masculine for the remote customer(s).

If the local customer doesn't wish to pay her provider's charges for reception, she will present the account of her received flow to the clearer (discovered using the procedure in the previous section). She will identify the remote sender who she expects will pay. The clearer looks up the sender's ISP and debits its account, referring to the sender's address. When the sender's ISP settles its account with the clearer, it will deduct the amount from its account with the sender. The price used is the clearer's, which may or may not be identical to the local provider's price. If the sender is a large organisation, it might have an account directly with the clearer. The clearer also credits its account with the local ISP at the local provider's price, referring to the local customer. When this account settles, it will clear the local customer's debt with her local provider.

The sender cannot dispute the charge unless he has evidence that the receiver previously agreed to accept payment liability for traffic of this type from him. Purely for the clearer's information, the local customer may identify which records she believes the remote party expects to pay and which she is simply disputing. Thus despite the customary case being each end paying its own charges, the onus is on the sender to ensure it can prove this. In most cases, there is a relationship between the parties in a communication which will allow the sender to safely assume the customary case so that no receiver is expected to bounce its customary charges back to the sender. Where such trust isn't present, the sender might wish to ensure receivers have confirmed they are accepting the flow on the customary terms (where they pay their own end's charges). However, some ISP's may offer sender debt collection as a service, bundling this liability with their service.

We now move on to consider outgoing flows sent

by the local customer. If she has evidence that the remote receiver has agreed to cover her local provider's charges for her sent traffic, she will present her account for the sent flow to the clearer. The mechanics of looking up accounts and debiting and crediting are as before. The receiver can only dispute this charge if he can prove the evidence saying he agreed to pay is invalid.

The local customer may optionally notify her local provider that she expects settlement for each flow to come from the clearer not herself. If the clearer never pays the local provider, the local customer remains liable to her own provider for the charges. If she has also already paid the clearer, she can ask the clearer for evidence that it has settled for the usage in question with her local provider. If it can provide the evidence signed by the provider, her liability to her provider is cleared. If it can't provide such evidence, it must return her payment so she will never have to pay twice.

This all seems very complicated, given it is per flow. However, in practice, settlement will be done in batch between each pair of parties as providers and clearers all have long term relationships with their customers. It is only the accounting that is per-flow. Indeed, even the accounting may be done in batches, but each batch will contain per-flow granularity of usage records. Also, it is clear that such reapportionment will only be cost-effective for flows where the value of the communications quality is higher than the cost of reapportionment. Examples would be long-lived flows, or collections of flows between the same ends where premium quality transmission service is used. Also, it should be recalled, that the clearing intermediary is merely a role. This role may be taken by one of the edge-ISPs or by one of the edge-customers, which removes the need for half the messaging.

Essentially, receipts are being traded much in the same way as employees claim travel expenses from their employer. This results in a scenario where the traditional telephone bill becomes an anachronism. Such a bill represents two commercial processes wrapped into one, which we propose should be separate. The first stage is reconciliation of all local usage records, whoever will pay. The second stage is agreement over who pays for which usage record.

9.3 Inter-domain multicast with heterogeneous QoS

Illustrating the power of the principles set down so far, we can take an example like multicast with heterogeneous QoS per receiver and show that charging for it with correct apportionment will 'just hap-

pen', even inter-domain. However, it will only be efficient by using the 'split-edge pricing' model. For illustration, let us have two provider networks with an edge customer of N_a sending into a multicast address where the tree crosses into N_b reaching receivers who are customers of N_b . N_b will have given a price for multicast reception and N_a one for sending to an address in the multicast range. The tree may spread to other receivers on other networks too. The receivers in each domain will note when they join the multicast and each start being charged for the traffic they receive at their local price. The sender will be charged at her provider's price. At the domain boundary between N_a and N_b , N_b will be charged N_a 's price for sending to a multicast while N_b will charge its price to N_a for receiving from a multicast. This usage may either be measured exactly at the inter-provider border, calculated statistically by combining customer usage data with multicast routing tables or simply covered within bulk measurements at the border. The receivers of a particular multicast group may happen to be located so that the multicast tree fans out immediately at its entrance to the network. With heterogeneous QoS per receiver (e.g. RSVP), any message to set up the QoS must emanate from the receiver and can therefore be charged for locally. Again this can be treated identically at the inter-provider boundary.

We believe edge pricing allows enough flexibility to charge differentially for broad ranges of route lengths because it allows different charges for different administrative domains. Even if a single domain spanned the globe, if desired it could be divided into internal pricing domains to achieve the same effect.

We now go on to discuss how reapportionment of charges between the ends might work in this multicast case. It would be unlikely that anyone would volunteer to pay for all receivers however many there were, unless they had a prior arrangement with each receiver. For instance, if a conference organiser were offering to pay everyone's communications expenses, part of the charge for each participant to join the conference would most likely cover these costs. Thus, any number of participants could join but the host would still have all payments covered from income. The host might just allow enough in the conference charge to cover most prices of most providers and probably make a little profit to cover the risk. Other models might be possible. The host may only agree to pay each participant's charges up to a ceiling. Alternatively, the host may ask for receipts authenticated by ISPs and pay the exact charge of each ISP. Thus, the end-to-end pricing model allows full flexibility for reapportioning charges in both the multicast and

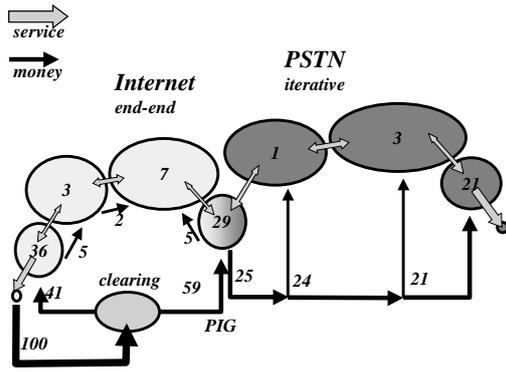


Figure 5: Clearing across a PIG

unicast cases.

9.4 Phone to Internet gateway

Others are taking the approach of allowing the telephony charging model to determine that for Internet telephony. Because of the complexity implications this approach has, we suggest the Internet should take advantage of the opportunity for a fresh start. The Internet should reject the complexity of iterative PSTN charging before it becomes endemic. Instead, phone to Internet gateways (PIGs) should be treated exactly as end-systems are treated above. Any apportionment of sender and receiver payments should be dealt with from one end of the Internet to the PIG. That is end-to-end across the Internet's patch, rather than end-to-end across both the Internet and the PSTN (Fig 5). This also allows for multicast models on the Internet side (multipoint on the PSTN side remains as complicated as today). Clearing between end-parties *must never* become muddled in with network provider pricing. If this were to happen, the Internet would be for ever saddled with interworking with a legacy, even when the legacy had virtually withered and died. It is always better to make the legacy interwork with the new model than the other way round.

Note that the end customer will see no difference if they rely on their edge network provider for all Internet telephony charging. This is purely an internal re-arrangement between ISPs. However, the customer could make these arrangements herself, if she desired.

10 Limitations and Further Work

The models described in this paper only become critical for any Internet communications that are usage-charged (e.g. QoS requests). Such cases are rare at present, however the author is engaged in related work, which suggests that the subscription and connect time charging models for Internet communications are only viable as long as capacity utilisation is low. However, even ultra-lightweight usage-charging [3] is yet to be proven cost-effective, therefore the context that this paper relies on is not at all certain without considerable further work.

Also this paper argues the case for 'sender and receiver both pay', but it is not proven. The stability of inter-related ISP policies requires 'war-game' simulation as a step towards a proof. Further work is also required to exercise scenarios based on these models, through simulation and prototyping in order to fully work through the performance and security issues.

11 Conclusions

We have defined a generalised pricing model for any number of interconnected multi-service networks. Each network offers each of its neighbours each class of service at a separate price for each direction of transmission. We call this 'split-edge pricing'. This model scales naturally to any size inter-network because all prices only depend on direct neighbours.

We have shown that the common case for apportioning value between the ends of a connectionless communication network is catered for if all users are charged for both sending and receiving. We have also shown that this is the most stable and efficient case, particularly for multicast and aggregation. It should therefore be used as the default in the 'split edge-pricing' model.

We have suggested that a new business model would be useful and more efficient to cater for the cases where there is a large discrepancy from this default in terms of value apportionment — large enough for it to be worth the cost of making a balancing transaction. This new model requires a new role in communications markets — an intermediary for end-to-end pricing and clearing. This new role could be conducted by existing ISPs or customers themselves, but there appears to be considerable added value, making this a viable business in its own right. It appears that this role is a threat to existing ISPs business. The role is one suitable for a purely financial processor using common

e-commerce mechanisms with relatively low costs and the ability to take a share of the surplus value available on top of charge reappportionments. This role relegates edge ISPs into wholesalers for a potentially large class of Internet applications. The intermediary would become the retail face of the Internet in many cases.

Further, we suggest a subtle twist to the recommendation that customers should pay for both sending and receiving. We suggest this should be customary, but that ultimate liability for sending should lie with the sender. Disputes could then quickly be resolved through the end-to-end clearing role. This stems from the unavoidable fact that receivers can't avoid being sent to.

Acknowledgements

Richard Gandon (BT International Carrier Services), Martin Tatham, Mike Rizzo, Jérôme Tassel, Steve Rudkin, Chris Russell (BT Research), Jon Crowcroft (UCL).

References

- [1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An architecture for differentiated services. Request for comments 2475, Internet Engineering Task Force, URL: [rfc2475.txt](http://www.ietf.org/rfc/rfc2475.txt), December 1998.
- [2] Bob Briscoe. The direction of value flow in multi-service connectionless networks. In *Proc. International Conference on Telecommunications and E-Commerce (ICTEC'99)*, URL: <http://www.btexact.com/projects/mware/>, October 1999.
- [3] Bob Briscoe, Mike Rizzo, Jérôme Tassel, and Konstantinos Damianakis. Lightweight, end to end, usage-based charging for packet networks. In *Proc. IEEE Openarch 2000*, pages 77–87, URL: <http://www.btexact.com/projects/mware/>, March 2000.
- [4] Frances Cairncross. *The Death of Distance; How the Communications Revolution Will Change Our Lives*. Harvard Business School Press, October 1997.
- [5] Georg Carle, Felix Hartanto, Michael Smirnow, and Tanja Zseby. Generic charging and accounting for value-added IP services. In *Proc. IEEE IFIP International Workshop on QoS (IWQoS'99) pricing workshop*, June 1999.
- [6] David D. Clark. Combining sender and receiver payments in the Internet. In G. Rosston and D. Watterman, editors, *Interconnection and the Internet*. Lawrence Erlbaum Associates, Mahwah, NJ, URL: <http://diffserv.lcs.mit.edu/>, October 1996.
- [7] Ross Finlayson. Breadcrumbs. Minutes of Third Reliable multicast research group meeting, Orlando, FL, and following on the mailing list URL: <http://www.east.isi.edu/rm/>, February 1998.
- [8] Mark Handley and Van Jacobsen. SDP: Session description protocol. Request for comments 2327, Internet Engineering Task Force, URL: [rfc2327.txt](http://www.ietf.org/rfc/rfc2327.txt), March 1998.
- [9] Shai Herzog, Scott Shenker, and Deborah Estrin. Sharing the cost of multicast trees: An axiomatic analysis. *Proc. ACM SIGCOMM'95, Computer Communication Review*, 25(4), August 1995.
- [10] Hugh W. Holbrook and David R. Cheriton. IP multicast channels: EXPRESS support for large-scale single-source applications. *Proc. ACM SIGCOMM'99, Computer Communication Review*, 29(4), September 1999.
- [11] Accounting rates and how they work. Technical report, ITU/Teleography Inc, URL: <http://www.itu.ch/intset/whatare/howwork.html>, 1996. Extract of chapter 3, Box 3.1 from [12].
- [12] The direction of traffic. Technical report, ITU/Teleography Inc, Geneva, URL: <http://www.itu.ch/ti/publications/traffic/direct.htm>, 1996.
- [13] Reforming the international accounting rate system. Web site, ITU, URL: <http://www.itu.ch/intset/>, at latest 1998.
- [14] Nadia Kausar, Bob Briscoe, and Jon Crowcroft. A charging model for sessions on the Internet. In *Proc. IEEE ISCC'99, Egypt*, volume 1629, pages 0246–, URL: <http://www.rennes.enst-bretagne.fr/~afifi/iscc99.html>, July 1999. Springer LNCS.
- [15] Jeffrey K. MacKie-Mason and Hal Varian. Some economics of the Internet. In *Proc. Tenth Michigan Public Utility Conference at Western Michigan University*, URL: <http://www.sims.berkeley.edu/~hal/people/hal/papers.html>, March 1993.
- [16] Sean McCreary and kc claffy. How far does the average packet travel on the Internet? Technical report, CAIDA, URL: <http://www.caida.org/Learn/ASPL/>, May 1998.
- [17] Mike Rizzo, Bob Briscoe, Jérôme Tassel, and Kostas Damianakis. A dynamic pricing framework to support a scalable, usage-based charging model for packet-switched networks. In *Proc. Int'l Wkshp on Active Networks (IWAN'99)*, volume 1653, URL: <http://www.btexact.com/projects/mware/>, February 1999. Springer LNCS.
- [18] Scott Shenker, David Clark, Deborah Estrin, and Shai Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM SIGCOMM Computer Communication Review*, 26(2), April 1996.
- [19] Tony Speakman, Dino Farinacci, Steven Lin, Alex Tweedly, Nidhi Bhaskar, Richard Edmonstone, Kelly Morse Johnson, Rajitha Sumanasekera, Lorenzo Vicisano, Jon Crowcroft, Jim Gemmell, Dan Leshchiner, Michael Luby, Todd Montgomery, and Luigi Rizzo. PGM reliable transport protocol specification. Internet draft, Internet Engineering Task Force, URL: [draft-speakman-pgm-spec-05.txt](http://www.ietf.org/draft-speakman-pgm-spec-05.txt), November 2000. (Work in progress) (expires 24 May 2001).
- [20] Y. Yemini, A. Dailianas, and D. Florissi. MarketNet: A market-based architecture for survivable large-scale information systems. In *Proc. 4th ISSAT International Conference on Reliability and Quality in Design, Seattle, WA*, URL: <http://www.cs.columbia.edu/dcc/marketnet/>, August 1998.
- [21] Lixia Zhang, Stephen Deering, Deborah Estrin, Scott Shenker, and Daniel Zappala. RSVP: A new resource ReSerVation protocol. *IEEE Network*, September 1993.

- [22] H. Zimmerman. OSI reference model — the ISO model of architecture for open systems interconnection. *IEEE/ACM Transactions on Networking*, 28(4):425–432, April 1980.
- [23] Chris Zull. Interconnection issues in the multimedia environment. In *Proc. Interconnection Asia '97, IIR Conferences, Singapore*, URL: http://www.cutlerco.com.au/core/content/speeches/Interconnection%20Iss%ues/Interconnection_Issues.html, April 1997.

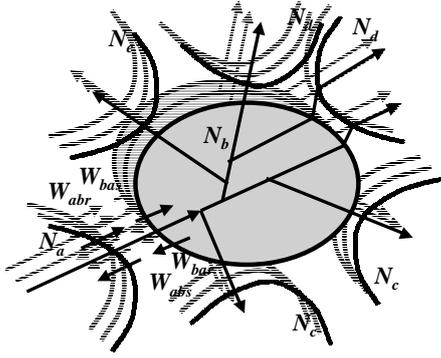


Figure 6: Split-edge pricing per class of service

A Independence of logical classes of service

Here we argue that each class of service can be treated independently of other logical classes of service. Fig 6 attempts to show the split-edge prices for different classes of service between N_a and N_b by layering the diagram in the third dimension ‘out of the page’. Each class of service ‘layer’ is a logical independent inter-network in its own right as it has its own share of resources and its own inter-network prices. Relating this to currently proposed technology, for integrated services [21], class of service is defined as either best effort, controlled load, or guaranteed service with the particular flowspecs reserved being dealt with as heterogeneous QoS within a class (Appendix B). For differentiated service (DS) [1], each DS code-point represents a class of service.

However, because each network is managed autonomously, there may be disjoint mappings between classes of service in neighbouring networks (as allowed in diffserv). Such a case is shown between N_b and the right-most N_d , which uses one class of service where everyone else uses two. At such a boundary, W_{bds} and W_{abr} for the merged class appear from N_b ’s point of view to each be a pair of prices for each of its classes of service that happen to be identical. On the other hand, N_b might offer two different W_{bdr} prices and two W_{bds} prices to N_d for each of N_b ’s two classes. N_d would just see each pair as a single price that varied depending on the relative proportions of traffic coming from or going to each class. Therefore, even with disjoint class mappings, we need not concern ourselves with more than one class of service at a time.

B Split-edge pricing model

There follows a description of the ‘split-edge pricing’ model. From this, one can analyse the surplus (or deficit) income that any party on a network can expect for any flow topology (unicast or multicast). Also pricing can be different for traffic in each direction so reverse multicasting (aggregation) is catered for in the model). There is no room here for analysis based on the model, but this is presented in [2] and conclusions are described here backed by natural language rather than mathematical argument.

We make the assumption that differences in length of routes or shapes of routing trees through any party’s network will not produce cost differentials that are significant enough to be worth measuring and charging for. I.e, we assume the ‘death of distance’ [4], or a ‘black box’ network routing assumption where a network’s routing is hidden within its interfaces. This means we trust a network or inter-network of federated domains to find the cheapest route without end-system intervention or, put another way, that routing protocols approximate to a competitive market. If the route isn’t truly the cheapest, we assume the discrepancy is so minor that the end-customer isn’t concerned. Border routing policy distorts this considerably, but one of our long term motivations is to make most border-routing policy redundant, simplifying inter-provider interfaces using usage-charging instead. This also implies that internal network design inefficiency should be absorbed by providers in their overall pricing. For instance in current multicast routing, tree stability always takes priority over tree efficiency. Providers are free not to make this decision if they can design better multicast routing algorithms, but there is no need to expose internal cost differentials in external pricing (when they are purely for provider convenience in the first place). This is deliberately unlike any of the scenarios analysed in Herzog *et al* — the classic work on sharing the cost of multicast [9]. Herzog *et al* has a scenario where all receivers share the cost equally, but not where receivers are divided into sets based on provider domains and only charged equally per domain (edge pricing). Nor does it consider heterogeneous QoS. If some degree of distance-based pricing *is* required, we assume edge-address-based pricing mechanisms can be used without having to concern the end-systems with the route between addresses. But we believe even this is unlikely.

Fig 7 shows a generic topology that will be used to illuminate the analysis. The general model and terminology have already been introduced in Section 5. We explained the scenario of a packet or flow being multicasted between inter-connected networks of various statuses relative to the one of interest

and we explained the four price weights at any interface. However, we glossed over the details of the model (e.g. not saying what the weights were applied to). We will now correct those omissions.

The figure shows the classes of service, Q , set for each branch of the tree. Q may be confined to discrete levels or allowed to take any from a bounded continuous range of values, depending on the QoS mechanism. It is assumed that, wherever a packet is duplicated for multicasting, the multiple copies might each have different classes of service. Note that it has not been assumed that the branches *all* have different classes of service. This depends on the (independent) requirements of the ultimate ends of each branch. There need be no correlation between neighbouring network status and the value of Q for that branch. RSVP is an example of such a heterogeneous QoS scheme (diffserv has the potential to become heterogeneous, e.g. by setting the class of each branch based on the DS-bytes in the multicast routing packets or in the IGMP join packets from end systems or even using RSVP). The packet or flow being modelled could be data or signalling.

V represents some measure of the volume or size of the service consumed. It might be the amount of data in the packet or in a flow of similar packets for a certain time. It might be the time for which a reservation of a certain size is held. We simplify the model by requiring V to be the same for all branches of the tree. This is justified because a branch leaving the tree, a packet loss or a network filter can be dealt with as an alteration to the topology rather than allowing V to be heterogeneous. We define the 'nominal charge' function that N_b levies for the packet or flow as $C_b(V, Q)$. This is a *nominal* charge because next we will describe how it is weighted to determine the *actual* charge for each different type of neighbour.

W_{bas} & W_{bar} denote the per direction weightings applied to the charge that N_b applies to N_a as shown at the highlighted interface between these two in Fig 7. The third digit of the suffix denotes the direction of traffic that the weighting applies to; s being the weight for traffic *sent* into the provider, r for traffic *received* from the provider setting the charge. However N_a is also offering service to N_b . So similarly N_a weights its charge C_a with weights W_{abr} and W_{abs} . The first digit of the suffix of W denotes the network provider setting the charge being weighted. The second digit denotes the type of neighbour network provider to which this weighting applies.

Often, we can consider scenarios where many of these price weightings are set to be either equal to each other or to be zero, but formulae derived for

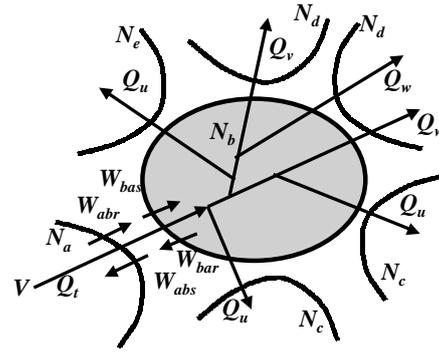


Figure 7: Split-edge pricing with heterogeneous QoS

this model allow fine granularity of price weighting for any scenario we might dream up in future work. We have initially experimented with extreme policies like 'sender pays all' but the formulae allow consideration of more subtle scenarios where prices might be slightly unequal in the different directions, perhaps because of asymmetric access technology like xDSL. The analysis in [2] takes this model and produces formulae for the surplus of each party. Then various scenarios such as 'sender pays all' or 'senders and receivers pay equally' are exercised to determine whether some policies are more advantageous to providers than others.

It is concluded that for *intra*-provider packets, the revenue is unaffected by whether senders, receivers or both are charged, but which customer(s) contributes to it, obviously depends on the direction of the unicast.

However, for *inter*-provider packets where peer providers charge each other as much as any edge customer, if both providers adopt a 'receiver pays all' or a 'sender pays all' policy, the revenue ends up always moving to the edge provider furthest from the customer paying. As long as each provider has a similar customer mix in terms of senders and receivers the net effect is that all providers make similar gains (and incidentally they could mutually agree not to charge each other with little change in their income — peering). Charging for only one direction has the benefit of halving the cost of measuring and accounting for both directions.

However, if all provider policies start as 'sender pays all', N_a might notice it has more receivers than senders. Let us assume N_a switches to charging receivers only, while N_b continues to only charge senders. A packet in the direction from N_b to N_a would result in each provider charging their edge customer for Q_t , but neither charging each other. A packet in the opposite direction would result in

neither provider charging their end customers at all, but both nominally charging each other. This would, as one might expect, cause a migration of all predominant receivers to other providers like N_b and all big senders to N_a . This leaves all providers with very little edge revenue, but all just nominally charging each other similar amounts. Clearly this will not be tolerated for long as the edge-customers are exploiting the providers, all the traffic is being inefficiently and expensively funnelled across the inter-provider links and there is consequent pressure for all providers to reverse their policy. Changing the wholesale pricing policies will make no difference. If the retail policies remain imbalanced there will be little revenue entering the system. Thus neither all providers charging only for sending nor all charging only for receiving can be stable unless there is some industry-wide agreement as to which one to standardise around. If there is such tacit agreement, charging for sending seems preferable as it avoids the problem of charging for unsolicited receipt.

What this teaches us is that any extreme policy where either sending or receiving are offered at low or zero price encourages instability simply because local pricing doesn't match true local costs. Therefore, when end customers arrange themselves to their best advantage, providers suffer unless they collectively organise themselves, which is unlikely. On the other hand, all providers charging for both sending and receiving is also stable, and without artificial standardisation. If any provider breaks ranks and charges, say, only senders, receivers will quickly migrate towards it and senders away to the rest of the industry. Thus the most stable arrangement is for send and receive pricing to approximately match costs.

When multicast is analysed, this only strengthens the case for all ends being liable for their use of the network, whether sending or receiving. Here we consider single-source trees for brevity, but the argument is even stronger for shared trees. Today, multicast senders are being charged exorbitant rates such that only those with income from (or advertising to) large numbers of receivers become customers. In the longer term, 'sender pays all' will only be tenable if it is at a rate that reflects the number of receivers. However, if this problem is considered, the trust required to achieve a solution appears to put up a theoretical barrier to 'sender pays all'. The cost of an inter-domain multicast is spread throughout the tree. The cost to each domain is approximately proportional to the number of branches within that domain. The proximity of the tree's fan-out to its ingress into the domain determines the actual cost per branch. This in turn depends on how meshed the network is, which is

usually a characteristic of the design of a whole domain (and if not, the internal design is under the domain's control). Thus, a 'leaf-domain' (one with only edge receivers and no downstream domains on the tree) can work out its cost for a certain number of receivers in its domain and report it to the upstream domain. This report is effectively a bill presented to the upstream domain, requesting a share of the income from the sender as it trickles down through the domains in the same direction as the tree. This upstream domain can collect similar 'bills' from other downstream networks and report the sum onwards upstream, eventually reaching the sender. Thus the head-end provider charges the sender the costs of the whole tree and has to pass on much of this income to meet the downstream 'bills'. Clearly, any domain can over-report its bill and make a profit. This is because the sender cannot determine the topology directly from its receivers without destroying the scalability benefits of multicast, which deliberately hides receiver activity.

It seems far simpler to apply the same argument as for unicast above and charge each receiver and the sender. The charges need not be identical for each, but the sender charge doesn't need to reflect the number of receivers. Where the tree crosses a domain boundary, each domain simply charges the other as one sender or one receiver. This ensures charging is completely distributed. There is not even a need to correlate together the receivers for one tree. All receivers are just usage-charged as they occur, with no need for co-ordination or totalling. As before, if the sender (or one receiver, or even a third party) wants to offer to pay for all receivers, this can be arranged end-to-end, rather than through the networks.