Dynamics of large-scale optical ATM networks

Martin Biddiscombe – UCL, Supervisor: Paul Kirkby – Nortel

A recent study has shown that even using today's most advanced opto-electronic technology, a future global-scale multimedia Internet network with unconstrained bandwidth and community of interest will be excessively expensive. This project is studying the conjecture that an all-optically switched global scaled core network could be built and would solve this problem. This would need buffering and timing of synchronous cells and 'super-cells' of ATM traffic at the periphery. Progress on testing this conjecture has concentrated on routing in super-symmetric core networks and dynamic pricing as a method of implementing edge buffering of ATM cells.

1 Introduction

If in the future we are going to be able to access a wide variety of broadband and narrowband services using a simple – unified access method, it is likely that these services will be provided using a common transport method – ATM. As the number and variety of services increase, the bandwidth requirements will also increase. It has been demonstrated (West, 1996) that if Internet-style connectivity (i.e. no restriction on range or bandwidth of connections) is allowed, the costs of switching in the core network will be prohibitively expensive. In other words, even with development of faster switches and larger buffers, current methods of ATM switching cannot be scaled up to a global level. This fact alone suggests that ultimately the core network switches will be fully optical.

It is therefore necessary to eliminate electronic components from the core of the network. Switching and processing within the core network must be optical. The network operator must be given the opportunity to utilise optical switching and routing at the VC/VP level without the need to process the ATM cells electronically.

An immediate consequence of this elimination of optical to electronic conversion at switches and multiplexers is that since the signal remains optical, it is not possible to buffer it as is possible with an electronic signal. It is therefore essential that ways are developed to eliminate those features of ATM traffic that necessitate buffering if significant numbers of cells are not going to be lost.

This paper considers two approaches which, if used in conjunction, may well provide the solution to loss-less ATM without buffers. These approaches are:

- Building super-symmetric networks with uniform inter-node delay so that network-wide synchronous switching can be employed;
- Utilising dynamic pricing of bandwidth so that buffering is "pushed out" to the periphery of the network, eliminating the requirement for buffering in the core network.

2 Super-symmetric networks

A three dimensional solid may be considered to be regular if its faces are all identical and its edges are all the same length. It is straightforward to demonstrate that there are only five such solids – the Platonic solids – which are:

Solid	Vertices	Edges per vertex	Faces	Face shape
Tetrahedron	4	3	4	Triangle
Cube	8	3	6	Square
Octahedron	6	4	8	Triangle
Dodecahedron	20	3	12	Pentagon
Icosahedron	12	5	20	Triangle

A polyhedron with any other number of vertices or faces than 4, 6, 8, 12 or 20 cannot be made up of identical faces. For instance the "Buckyball", which has 60 vertices, consists of pentagons and hexagons.

However, in four dimensional space it is possible to construct symmetric polyhedra much more easily. These super-symmetric objects can be projected into three dimensional space. Certain properties are not retained under projection, such as equality of lengths, areas or enclosed volumes (as happens when three dimensional objects are projected into two dimensions). However vertices connected by an edge in the four dimensional (real) object remain connected in the projection, and are connected by the projection of the edge that connects the vertices in the real object. Similarly, projections of edges which meet at a vertex in the real object meet at the projection of that vertex.

It is therefore possible to construct networks where the nodes are the vertices of a super-symmetric four dimensional object, and the links consist of the edges of that object. It does not matter that the three dimensional projection of the object does not have equal length edges – the links in the network can be “padded out” as necessary with fibre so that the time taken for a signal to travel between any pair of connected nodes remains the same as it would in the four dimensional network.

2.1 Hypercube

One of the simplest objects that could be considered is the four dimensional hyper-cube. This has sixteen vertices, each of which is connected to four others. A hyper-cube is constructed from eight cubes in a similar way to that in which a cube is constructed from six squares. A three dimensional projection of the hyper-cube will also have sixteen vertices, each adjacent to four edges. A suitable object is the 16-node torus. This is a torus with four rings running in each “direction”. Actually, we can consider a torus made up of any number of rings. Any such torus is equivalent to a four dimensional symmetric object – obtained by stretching and bending the edges of the torus to produce an object with straight, equal length edges. However, the simplest is the 16 node torus. The following discussion could be applied to a torus with any square number of nodes. The diagrams below are representation of a four dimensional hyper-cube and a 16-node torus (links bold):

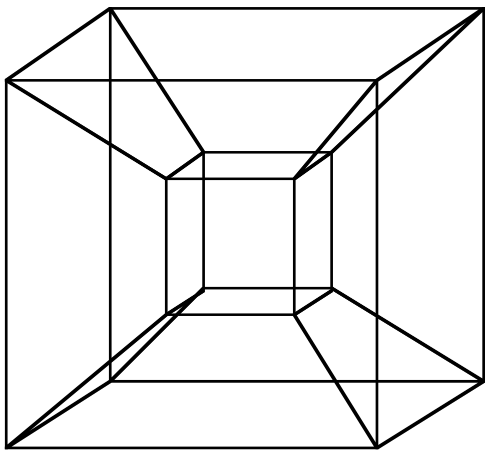

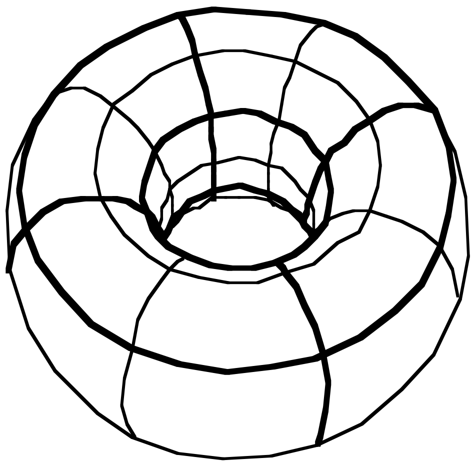

2.2 Routing and traffic considerations

The highly meshed nature of the torus based network means that there are many routes between a given pair of points. We shall be primarily concerned with *shortest* routes between nodes. We shall derive expressions for calculating the number of routes between nodes and the traffic flow on those routes under a variety of conditions.

The torus was introduced as consisting of interlocking *rings*. Indeed, that is how it is constructed. The first question to consider therefore is whether or not the traffic can flow in both directions between adjacent nodes – is the torus constructed from one set of overlapping rings or two sets, superimposed? We can derive results in either case without difficulty. There is no significant difference between these cases. Care must be taken in determining the shortest route in the “2-way” case – since it is possible for up to four distinct sets of shortest routes to exist due to the symmetry of the network.

The initial step in understanding the traffic flow around a torus network is to realise that a torus can be “unwrapped”, resulting in a grid network. Firstly the torus is cut vertically – resulting in a cylinder. Then the cylinder is cut lengthways – resulting in a square grid. This is represented diagramatically below.

Each ring has been cut, and so it is important to realise that a signal reaching the edge of the grid “loops back” to the other side. On the diagram shown, the opposite edges of the grid represent the same links in the network.

Having opened out the torus into a grid, it is now little more than a counting exercise to determine the number of routes between nodes. If the nodes are numbered using Cartesian co-ordinates, then the number of shortest routes from (0,0) to (*x*,*y*) is given by: $\binom{x+y}{x} = \frac{(x+y)!}{x!y!}$. This is readily extended to give the number of routes between arbitrary nodes (*u*,*v*) and (*x*,*y*) as: $\binom{(x-u)+(y-v)}{(x-u)} = \frac{(x+y-(u+v))!}{(x-u)!(y-v)!}$.

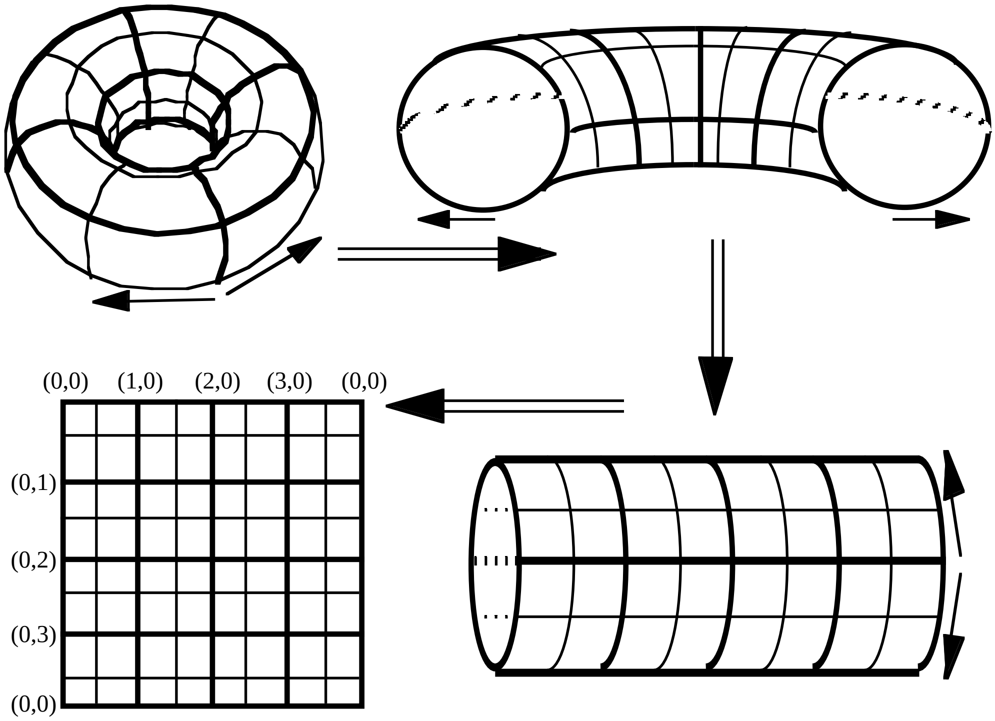

A torus can be unwrapped, producing a square grid

Care must be taken if the originating and terminating nodes straddle the “edge” of the grid. In this situation it is necessary to consider the terminating node as being at (*x*+4,*y*), (*x*,*y*+4) or (*x*+4,*y*+4) as appropriate (assuming the size of the grid is 4x4). These expressions can be used in conjunction with the assumption of a uniformly global community of interest, and that traffic is distributed evenly across available routes, to derive values for traffic flows under various conditions.

3 Dynamic pricing

The symmetry of the network can be used to eliminate cell-scale queuing at switches, since at any given time the network can be considered to be globally employing a form of time division multiplexing. The configuration can be changed (relatively slowly) as a result of time-varying user requirements. It remains to determine a method whereby buffers intended to accommodate temporary bursts of cells can be eliminated. These are required when contention occurs for a limited resource over a period of time. In essence, for the period of the burst, the demand for bandwidth exceeds the supply.

Traditional methods for coping with bursts of data rely on large buffers at the bottlenecks. If the number of incoming cells exceed even the buffer capacity they are lost. There is little or no scope for discrimination between cells in terms of priority of access to the resource. This does not reflect the “real world” situation, where cells will be valued differently by their owners according to the application they are supporting and the level of importance that their owner attaches to the successful execution of the application. For instance, cells carrying real-time data (such as for a telephone call) can bear substantial loss before quality is significantly degraded. However, if a cell is delayed it rapidly becomes worthless. On the other hand, during an operation such as off-line file transfer, substantial cell delay variation is unimportant but minimal cell loss is imperative. It is not difficult to envisage services which require any combination of cell loss and delay tolerance levels.

These variations in performance requirements can be expressed as a *value* (as a function of time), which represents the importance to the user of the cells being delivered within that period of time. The value of cells for a telephone call, for example, will drop off rapidly as time passes. The initial value of the cells need not be uniform across all telephone users though – for instance, within a company, a senior manager's cells may be considered more valuable than those of a temporary junior employee.

Bandwidth may then be priced as a function of the spare capacity of the system – as the total demand increases, the price increases. Only those users who consider their cells sufficiently valuable are allowed access to the network. If an existing user's cells are not sufficiently valuable, then the user must either increase their willingness to pay, or else lose access. Therefore as demand increases, either bandwidth is freed to meet the demand or revenue increases. With the correct pricing structure, consistently high demand provides sufficient increased revenue to finance the necessary increase in capacity (Mackie-Mason & Varian, 1993 and 1994).

3.1 Implementation

There are several ways in which such a dynamic pricing mechanism could be introduced. Two potential implementations are outlined below.

3.1.1 Auctioned access

Each user decides the absolute maximum that they are willing to pay for each service, in terms of (say) £/minute. Their application submits a bid to the network operator consisting of a) how much bandwidth is required and b) how much the user is willing to pay for it. Starting with the user who places the greatest value on their data, the network operator calculates the bandwidth that would be required if successive users were admitted. When no more users could be accommodated, the network operator instructs those applications above the cut-off point to send their cells. The users are not charged at the rate they indicated they would be willing to pay, but at the rate indicated by the user immediately above the cut-off point. Technically this is very similar to a Vickrey auction.

3.1.2 Advertised prices

The network operator broadcasts the current price of bandwidth. A user's application compares the cost of transmitting with the user's willingness to pay. If the cost is less than or equal to that value then the data is sent. The network operator imposes rules on how rapidly cell transmission rates can be increased – an application cannot be allowed to send a sudden burst of cells (since if several sources sharing a resource did this, the resource would overload without warning and cells would be lost). Instead, such bursts are buffered at the periphery of the network and the output cell-rate is increased gradually. This allows time for the network operator to react to the changing demand, varying the price accordingly.

4 Conclusions

If the concepts described above are valid, then network symmetry will enable cells to be aggregated into 'super-cells' with a common destination and common routing. The switches can then operate at a lower rate than required if individual cells were to be switched – the network will reconfigure slowly as user demands change. The network capacity is not exceeded, due to the edge-buffering, and as the network utilisation increases, the dynamic pricing mechanism ensures that the value of data being transmitted is maximised.

5 References

Mackie-Mason & Varian (1993). Pricing the Internet – prepared for Public Access to the Internet, JFK School of Government, May 26-27 1993.

Mackie-Mason & Varian (1994). Pricing Congestible Network Resources, University of Michigan Department of Economics, July 1994.

West (1996). Modelling of High Capacity Telecommunication Networks, UCL Department of Electrical and Electronic Engineering, EE/CS BEng – Final Report, March 1996.