8-15-1997

Tradeoff Decisions in the Design of a Backbone Network Using Visualization

Kemal Altinkemer
Purdue University, kemal@mgmt.purdue.edu

Indranil Bose
Purdue University, bosei@vm.cc.purdue.edu

Alok Chaturvedi
Purdue University, alok@mgmt.purdue.edu

Tradeoff Decisions in the Design of a Backbone Network Using Visualization

Kemal Altinkemer
Indranil Bose
Alok Chaturvedi
kemal@mgmt.purdue.edu bosei@vm.cc.purdue.edu alok@mgmt.purdue.edu
Krannert Graduate School of Management
Purdue University, West Lafayette, IN 47907
Phone: (765)-746-1633 Fax: (765)-494-9658

ABSTRACT :

Visualization provides a useful tool for analyzing large, complex data sets. In the design of backbone computer networks rough cut design decisions can gain from a visual analysis of the generated solution with respect to design parameters. An important step in network design is to decide on the tradeoff between average delay incurred in routing the messages over the network and the total cost of network design. In this research we show how two dimensional and three dimensional surface and glyph representations can be used for understanding the cost-delay tradeoffs involved in the network design problem. The major contribution of this research is that it provides a flexible, visual front-end to the abstract design process and allows the users to decide on the efficient frontier where they would prefer to operate.

INTRODUCTION :

Most problems related to the design of computer networks are combinatorially explosive and involve decisions about a large number of parameters such as average message length, average message delay, utilization of links, cost of delay, capacities of links, routing protocols, etc. Since these problems cannot be solved to optimality, efficient heuristics are usually suggested to solve these computationally intractable problems using the well established methodology of 'divide and conquer'. The solutions produced by these heuristics are not the 'best' in general and computational tests need to be performed to test the robustness of the heuristics to system parameters. This is often referred to in the optimization literature as sensitivity analysis. Traditional sensitivity analysis may not recognize the 'hidden' relationships that exist between the generated solution and the problem parameters. With two dimensional and three dimensional visualizations these relationships sometimes become very obvious. The major premise of this paper is to discover such relationships that may exist between heuristic solutions generated using a Lagrangean relaxation based procedure and the parameters involved in design of a computer network.

Optimization problems in general are visually oriented and in certain situations a visual representation of the problem enables a human expert to identify 'patterns' and come closer to the best possible solution. In the area of design of computer networks the use of visualization based techniques have attracted the attention of researchers recently. Examples of these can be found in Becker et al. (1994), Dean et al. (1995), Jack et al. (1992), Liu et al. (1992) and Martin (1992).

PROBLEM STATEMENT :

Among the many parameters involved in the design of backbone networks, message length (bits), average message delay (msec) and unit delay cost ($) are some of the more significant ones. The unit delay cost represents the opportunity cost of the user. The delay cost arises due to the waiting time of the computer users in front of their terminals. Unit delay cost represents the amount of money that the users of the network are willing to pay for delivery of their messages.

The total cost of network design consists of three components - the fixed cost, the variable cost and the queuing cost. The fixed cost can be broken into two components - a fixed setup cost and a term proportional to the length of the link. The variable cost is proportional to the message rate and represents the cost per message unit charged by a common carrier. The queuing cost is the cost associated with the average delay of the messages being transmitted over the network.

Gavish and Altinkemer (1990) studied the problem of assigning capacities to the links and deciding on routes for messages for each origin-destination communication pair in a backbone computer network. The topology of the network was given and the end to end traffic requirements were known. The optimization procedure sought to find a least cost design for the network. An interesting aspect of this research was the study of the tradeoff between the different costs namely the fixed cost, the variable cost and the queuing cost. Another important issue was determination of the tradeoff between total cost of network design and average message delay.

The purpose of the current research is to study the tradeoff between the different cost components, average message delay, total network operating costs and various combinations of message length, link capacities and unit delay cost using visualization techniques.

DATA GENERATION :

Gavish and Altinkemer (1990) used a Lagrangean relaxation based method for determining the least cost network design in terms of routes of messages and capacities of links. They tested their algorithm on four different networks - namely ARPA, OCT, USA and RING. In this paper ARPA is taken as the example network for the study of the cost-delay tradeoff.

The Lagrangean relaxation based solution technique is used for generating the data points. For the link capacities 7 choices are allowed. The choice of capacities and the corresponding setup cost, distance cost and variable cost are the same as in Gavish and Altinkemer (1990). A route generation procedure is used as an integral part of the Lagrangean procedure that generates routes as needed during the solution process. The experiments are run for 5 different message lengths - 200, 300, 400, 500, 600 bits respectively. The unit delay cost is varied as 100, 500, 1000, 2000 and 3000 (in dollars). These lead to 25 possible combinations. However, these are for a particular value of the fixed cost multiplier. The fixed cost multiplier is then varied in 3 steps - 0.5, 1.5 and 3.0. For the above experiments, the variable cost multiplier is kept fixed at 1.0. For each choice of fixed cost multiplier, the solution procedure generates the total cost of network design, components of the total cost (namely fixed cost, variable cost and queuing cost) and the average message delay. The solution also gives the gap between the Lagrangean value and the feasible solution. As a second part to the procedure, the same set of experiments are repeated, keeping the fixed cost multiplier fixed at 1.0 and varying the variable cost multiplier in 3 steps - 0.5, 1.5 and 3.0. The main idea behind changing the fixed and the variable cost multipliers is to find out how the total cost and the average message delay changes with change in the fixed cost or the variable cost. The visuals based on this data are obtained using IBM's visualization software - Data Explorer.

VISUAL TECHNIQUES :

Several visual techniques are available for data representation. Some of the techniques that are available in IBMís Data Explorer are color map, surface plot, color glyph, isosurface, contour line, volume rendering, sequencer slice, streamline and streakline.

For the research in question three visual techniques are chosen for creating the visuals. The first technique used is a color map. A color map represents a relationship between a range of data values and a set of colors. In this case, spectral colors are used in the color map where ëredí signifies a ëhighí value of the variable and ëblueí signifies a ëlowí value of the variable. Color maps are useful if the user is interested in the values of the fourth variable on the hull of the surface formed by three variables. The second technique used is a surface plot. A surface plot is a surface that connects all points in a three dimensional space. Surface plot is useful when it is possible to conceive of a surface that joins all data points. The third technique used is a color glyph. Glyph representations are made by copying a generic object, usually a sphere, and positioning each copy appropriately and coloring them according to the data associated with that sample point.

The visuals generated using the results of the numerical experiments are displayed in either the ëoff-diagonalí or the ëfrontí view. For any visual that particular view is chosen which revealed the most information. Data Explorer provides 14 different views and also provides capabilities for rotation, zooming, panning, roaming, navigation and animation of the generated visuals.

EXAMPLE :

The following visual represents the variation of total network design cost and average message delay with the variation in the components of the total cost (i.e. fixed cost, variable cost and queuing cost) for different choices of average message length and unit delay cost. The fixed cost and the variable cost multipliers are equal to 1. The average message length and the average delay cost is varied in five steps as detailed before, giving rise to 25 data points. The total cost is represented by a surface plot and the average message delay is represented at each data point by a glyph. Both the surface plot and the glyphs are colored using a color map.

Some conclusions that can be drawn from this example visual are :

- Total cost is higher for higher fixed cost
- Highest total cost does not imply lowest average message delay and vice versa
- Lower queuing cost leads to higher average message delay and vice versa

Using this visual, the designers would be able to see what investment in fixed, variable and queuing cost leads to what average message delay for a chosen value of average message length and average delay cost. This implies that for a particular application that the user wants to send they will be able to choose different combination of the cost components. For example, for real-time video with stringent delay requirements the user will choose a point on the total cost surface where the glyph color indicates the lowest value of the average message delay. With that choice, they will be able to determine the total network design cost and how much needs to be spent as fixed, variable and queuing cost to support a video application over the network.

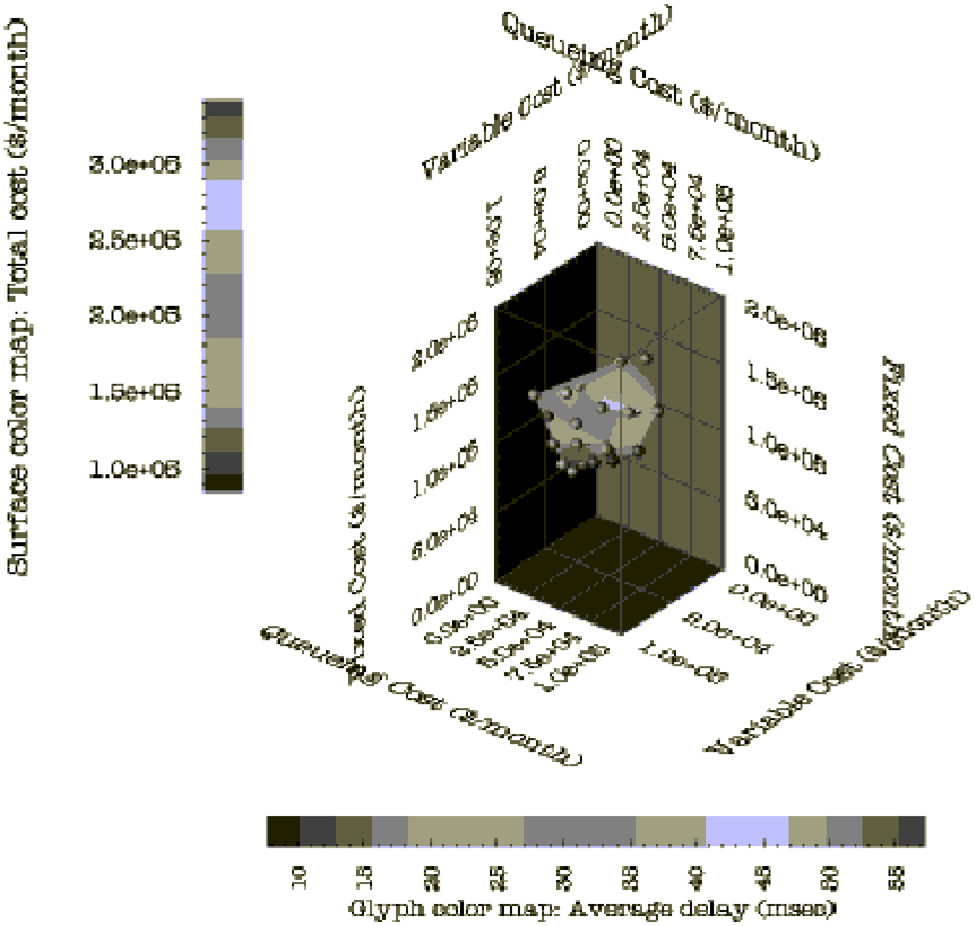

SUMMARY OF RESULTS :

A set of 24 visuals were created showing the tradeoff between the various network design parameters. The following observations can be made from all the visuals that were constructed :

- The total cost increases with increase in the message length irrespective of the multiplier for fixed cost and variable cost.
- Highest total cost does not imply lowest average message delay and vice versa. However, for high total cost the average message delay is usually low.
- Higher total cost does not imply lower average message delay. The effect of message length seems to be dominant.
- The total cost is highest for medium range delay.
- The change in the total cost is high for a small change in message delay at high delay cost.
- The total cost is higher for higher fixed cost.
- The range of average message delay (for the same message length combinations) increases as the multiplier for the fixed cost or variable cost increases.
- Higher delay cost leads to higher total cost for all combinations of fixed and variable cost multipliers.
- Higher delay cost leads to lower average message delay and vice versa.
- Lower queuing cost leads to higher delay and vice versa.

Some observations are intuitive and some are not. They can be used by designers to decide how to allocate investment for specific delay guarantees depending upon the type of application that they would like to support over the network.

CONCLUSION :

The main emphasis of this research has been to understand the relationships that exist between the generated solutions and the problem parameters in a network design problem using visualization. Taking the example of the ARPA network it is shown that several meaningful relationships exist between total cost, average message delay, message length and delay cost. In future we would like to repeat the experiments for sparse and dense network topologies and with variable demands between origin-destination pairs.

REFERENCES :

Will be provided on request