Extended Abstract

Dissertation Title: Oblivious Buy-at-Bulk Network Design Algorithms
Author: Srivathsan Srinivasagopalan (LSU-ID: 89-761-4646)
Major Professors: Prof. Costas Busch and Prof. S.S. Iyengar
Department: Computer Science

1 The Big Picture

1.1 Scenario - 1

Let us suppose that you are an independent contractor who deploys/installs telephone lines. A local telecom company wants to expand its reach by offering telecommunication capabilities to a new suburb and releases a tender. You want to win this contract. To address this tender, you have to answer some questions to yourself: How would you go about buying expensive copper cables (in dollars/unit length) and installing them? Would you connect all the households in the suburb to the local exchange by direct lines? What would be its cost? Is this approach an *optimal* approach with respect to the cost incurred? If so, is this approach scalable? Is there a more optimal way to connect the neighborhood to the local exchange? If so, how will you find out?

The above is an example of a *Network Design* problem. It is a very difficult problem to solve in a reasonable amount of time and currently there is no known solution to problems of such kind (known as "NP-Complete" class of problems in Computer Science jargon). However, approximate solutions exist. Researchers, from Computer Science and Operations Research fields, have been working on many variants of such network design problems for over 40 years. These problems are often called "discrete cost network optimization" in Operations Research.

Since there is only one local telephone exchange and all telephone calls will have to be routed to that exchange, we call this a *Single-Sink Network Design* problem. The term "sink" refers to the destination point where every call gets routed to.

1.2 Question

Following the above example, we want a solution where no household is left out and where no subset of households are connected by a disproportionately long cable. In other words, we want to find a network where *all* households in the suburb are connected to the local exchange in the most *compact* manner. The distance from *any* household to the local exchange must be as short as possible. Can we find such a network?

1.3 Scenario - 2

Let us consider that the same telecommunication company now wants to provide high-speed internet and cable-TV (besides regular phone lines) to the same neighborhood. You want to win the tender and get the contract!

You have to consider economies of scale because in such networks, bandwidth is purchased in some discrete units. The economies of scale exhibits the property where the cost per bandwidth decreases as the number of units purchased increases. This property is why network capacity is bought/sold in "wholesale", or why vendors provide "volume discount". In order to arrive at a good solution, you have to think about installing "fat" cables/optical fibers from the local exchange up to a certain distance and then gradually taper-off by smaller and smaller capacity cables/fibers until you reach individual households. This is the same as you would imagine about laying water-pipes from a water-tank to a neighborhood - start with a big, fat pipe and gradually shrink the pipes' diameter until you hit each household. Big fat pipes are more expensive (per unit length) than the smaller ones. Therefore, the more the number of households that connect to a big fat pipe, the better for both parties. The customers pay less to get what they want and at the same time, your initial investment in laying those cables/fibers can easily be recovered when there are plenty of customers at the end of the cable.

1.4 Modified Question

How will you lay out the expensive high-speed optical fibers (for Internet) and TV cables? The demands (in the form of movie/music/file downloads and TV channel viewing) from various households vary in a dynamic fashion. There is no known "fixed" demand from any of the households. The network you install must accommodate for *any* demand scenario at any *time*. Hence, the previously mentioned concept of laying "fat" pipes initially and gradually "tapering" will not work here. Each segment of the cable and the network as a whole must have the capacity to accommodate the demands by aggregating them as required. Furthermore, when data from different sources gets routed/aggregated through a common cable, there is a cost associated with it. The more data gets aggregated along its path, the greater is the "fusion" cost.

This is what is called *Single-Sink Buy-at-Bulk* network design problem (SSBB) where you try to find a set of *minimum-cost* paths from all the households to the local exchange. Here, "minimum-cost" refers to the combination of both the cost/length of the cable plus the fusion cost. Can you find such a set of paths that will satisfy any demand scenario?

1.5 Scenario - 3

The same telecommunications company now wants to enable the customers to send/receive data among themselves in an efficient, robust and reliable manner. This means that the data must be routed using a shortest-path from the source to the destination. This means that, most often, the data should not be routed through the local exchange for many reasons: (i) the distance the data traveled would be far greater than if the data was routed via different households and avoiding the local exchange. The extra distance adds to additional cost and delay (ii) If all data is routed through a single point (the local exchange), then, that point will not only become a bottleneck, but also prove to be a single-point-of-failure. That is, if the local exchange breaks-down (or compromised by a malicious user or virus), then, telecommunication/Internet/TV services would be down for the entire neighborhood.

Typically, the demand flows are in discrete units and are unsplittable (indivisible), i.e., the flow follows a single path from the demand node to its destination.

1.6 Modified Question

To address the above problem, you, as a contractor, must figure out a way to find out a set of *minimum-cost* paths from *any* set of sources to *any* set of destinations in that neighborhood. This problem is called *Multi-Sink Buy-at-Bulk* network design problem (MSBB).

2 Solution

My dissertation is about the study of *oblivious* buy-at-bulk network design problem (SSBB and MSBB) with the following constraints: an unknown set of demands and an unknown fusion cost function (hence the term "oblivious"). In other words, I describe a novel approach for developing an oblivious spanning tree (or set of paths without any cycles) in the sense that it is independent of the number and location of demand sources and cost function at the edges. Several publications (conferences and journals) resulted from this research study. The list of related publications can be seen in my vita.

Many practically relevant instances of network design problems are very hard and thus are likely intractable. The research work presented in this dissertation focuses on approximation algorithms as one possible way of circumventing this impasse. Approximation algorithms have been used widely and for a long time. They are known to be very efficient (i.e., they run to completion in a reasonable time) and provide solutions to instances of many different optimization problems whose objective values are close to those of their respective optimum solutions. More specifically, the problems discussed in this work are minimization problems.

I have provided tight solutions for three types of graphs - doubling-dimension, planar and minor-free graphs. Many real-world scenarios exhibit the properties of one of these three graph types. For example, the Internet exhibits the properties of doubling-dimension graphs, oil/gas and transportation/railroad networks exhibits planar (or minor-free) graph properties.

3 Why should we care?

An abstraction of this problem can be found in many applications that can be formulated as networks on a plane. Network design lies at the core of them. An optimal solution to this fundamental problem will (immediately) positively impact in the design of *all* the related practical applications.

To the best of my knowledge, the solutions presented in my dissertation and the publications are the first of its kind to propose a single spanning-tree (for SSBB) and set of minimum-cost paths (for MSBB) problems. In the following paragraphs, I have provided a few applications (among dozens) where my approximation algorithms can be used. I hope this will give a sense of why my research is important.

3.1 Significance of the Research

VLSI Power Circuitry: The exponential scaling of feature sizes in semiconductor technologies has side-effects on layout optimization, related to effects such as interconnect delay, noise and crosstalk, signal integrity, and power dissipation, that invalidate the assumptions that form the basis of previous design methodologies and tools.

In a microprocessor, there are several components that need power. To minimize power usage and heat generation, microprocessors work by activating only those components that need to work while others are inactive. So, at any instant of time, only a subset of components must be powered by a *single* power circuitry. Also, this single circuit that connects all the components must be of near-optimal length for all demand scenarios. The smaller the length of the wires, the lower the IR-Drop (power dissipation).

Wireless Sensor Networks: Distributed Wireless Sensor Networks collect and send information to one or more sinks via multiple hops in the network. During this multi-hop relay of information, it gets aggregated with other information at the fusion points (nodes). Typically, sensor network applications may care only about aggregate information (e.g., average temperature, humidity etc). An important aspect in such networks is the dynamism in the set of sources that needs to send data. At various instances of time, different set of sources might have data to send to the sink(s). Since wireless sensor nodes are energy constrained, they are incapable of computing an optimal tree for every instance. In such cases, one needs to build a single tree (or a set of paths) to route data to the sink(s) .

Oil/Gas Pipelines: There is a cost in laying oil/gas pipes to connect various stations/cities. Naturally, the larger the capacity of a pipe and the greater the number of consumers using the pipe, the cheaper would be price to pay for using the pipe (by the consumers). Hence, to build an optimal pipeline, buy-at-bulk network design principles comes into play.

Cloud Computing/Data-Center Networks: Cloud Computing is quickly being adopted by various industries and customers alike. A key factor in the performance of cloud-computing is the network efficiency of the associated data-centers (DC). Data centers are located geographically apart to serve customers in all regions. The inter-DC network bandwidth poses a high-risk in performance if the network is not properly designed. This problem boils down to designing an efficient network for a variety of scenarios.

A few others: Transportation & Logistics, smart power grids, publish- subscribe systems, facility location (in urban planning), quality-of-service in mobile applications, mapping applications etc.

4 Conclusion

Network design problem lies at the core of many practical applications. My dissertation focuses on providing good approximate solutions to the fundamental problem of network design (SSBB and MSBB). The solutions immediately impacts the related practical applications. Preliminary simulations corroborate the analytical results. Many top-tier conference and journal publications have resulted out of this research work.

Related Publications

[1] Srivathsan Srinivasagopalan, Costas Busch, S.S. Iyengar, "Oblivious Buy-at-Bulk in Planar Graphs.", In the proceedings of theWorkshop on Algorithms and Computation (WALCOM 2011), in Springer-Verlag Lecture Notes Computer Science, LNCS Vol. 6552, pp.33-44, 2011.

[2] Srivathsan Srinivasagopalan, Costas Busch, S.S. Iyengar, "Oblivious Buy-at-Bulk in Planar Graphs.", Submitted to SIAM Journal on Discrete Mathematics (SIDMA).

[3] Srivathsan Srinivasagopalan, Costas Busch, S.S. Iyengar, "An Oblivious Spanning Tree for Single-Sink Buy-at-Bulk in Low Doubling- Dimension Graphs", IEEE Transactions on Computers.

[4] Srivathsan Srinivasagopalan, Costas Busch, S.S. Iyengar, "Universal Steiner Trees for Data Aggregation in Low Doubling Metrics", Best Presentation Award at the Ph.D. Forum in the International Conference on Distributed Computing and Networking, ICDCN-2010, Kolkata, India. Jan 03-06, 2010.

[5] Srivathsan Srinivasagopalan, Costas Busch, S.S. Iyengar, "Universal Data Aggregation Trees for Sensor Networks in Low Doubling Metrics", brief announcement at International Workshop on Algorithmic Aspects of Wireless Sensor Networks, ALGOSENSORS'09, Rhodes, Greece, July 10-11, 2009. LNCS, Springer Verlag, Vol 5804/2009, pp 151-152.

[**Complete list of publications, talks, awards:** http://csc.lsu.edu/~ssrini1]