

Doctoral Written Exam in Networking, Fall 2007

December 18, 2007

Answer all parts of all questions. There are three multi-part questions, each of equal weight. Turn in your answers by Friday, December 21, at the same time you picked up your exam, to Jennifer Streubel at MCS-140G.

Brevity is the soul of wit. Be brief and to the point. This is a depth exam. So, your answers should show research maturity and depth. Do not pad your answers and do not write for the sake of writing.

You are free to consult any notes, books and papers during the examination, but make sure to include your references. You must develop your own solutions, which might use existing ideas and techniques, as long as you cite them and clearly explain how these existing ideas/techniques fit within your own solution.

You are not allowed to consult with any person during the examination.

If you have any doubt as to the interpretation of a question, make a reasonable assumption and explain your interpretation in your answer. No explanations will be given during the exam.

1. Assume you work for a major backbone ISP like Sprint or AT&T. You have complete access to their infrastructure. You can turn on any features in routers that you like, you can instrument any links any way you like, and you have complete access to all the data you can collect.
 - (a) Your boss comes to you with a simple question: What fraction of all the traffic on the Internet flows through our network?
 - i. Describe the challenges you are faced in answering this question.
 - ii. Describe what steps you would take to answer this question. If you need to make assumptions, do so, but state them clearly. Naturally, the fewer assumptions you make, the better.
 - iii. Discuss as quantitatively as possible the level of confidence that you would have in your results.
 - iv. Describe what the practical use might be to answering this question.
 - (b) Your boss comes to you with a new problem. She thinks there may be traffic in your network that is generated by botnets. Your boss asks you: “How much traffic is due to botnets? Which hosts are infected? Can you identify specific points in time when botnets are active?”
 - i. Describe the challenges you are faced in answering these questions.
 - ii. Describe what steps you would take to answer these questions. If you need to make assumptions, do so, but state them clearly. Naturally, the fewer assumptions you make, the better.
 - iii. Discuss as quantitatively as possible the level of confidence that you would have in your results.
 - iv. Describe what the practical use might be to answering these questions.

2. Assume you have satisfied your old boss, and have now been promoted within your ISP to a job where you must make strategic corporate decisions.
- (a) Your first task is to set prices for transit traffic through your ISP. Suppose you have obtained competitive information that reveals the demand curves of all other ISPs in the network for transit traffic through your ISP, e.g. functions $f_i(b)$ that specify the utility ISP i receives when transiting b bits of bandwidth through your network. Assume that it costs your ISP a function $c(b)$ to supply b bits of bandwidth in aggregate. Your job is to maximize profit.
- i. Formulate this problem as an optimization problem and devise an algorithm that chooses the profit-maximizing price per bit to charge. Also, specify its running time. You may make reasonable simplifying assumptions, e.g. convexity, but if you do, specify why they are reasonable and necessary to carry out the analysis.
 - ii. Your boss, the CEO, wonders why the other ISPs should all be charged the same amount. He suggests two-tiered pricing: Price A will be offered to one set of customers and Price B will be offered to the rest. Assuming customers cannot negotiate the price and act in their own self-interest, devise an algorithm that sets the two prices and partitions the set of customers into Price A and Price B groups so as to maximize total profit. Can this be done in polynomial time? Under what conditions (if ever) does this increase overall profit?
 - iii. Briefly discuss factors that would change in the answers above if you were to model the details of your ISP's network in the cost analysis. Include a description of which aspects of the network would be most beneficial to expose in the model.
- (b) Based on the methods you have devised, your ISP has done well for itself. So well, in fact, that you are asked to join a consortia of ISPs that want to make the world a better and friendlier place. by devising a strategyproof mechanism that incent all ISPs to reveal their private information $c(b)$ and $f_i(b)$ (actually one $f_i(b)$ for every other ISP j). Here, as in the very first part of the question, ISPs are each setting an ISP-wide transit price per unit bandwidth.
- i. Argue that a VCG-based mechanism for this problem is feasible in theory, by outlining a construction. Feel free to make simplifying assumptions as needed.
 - ii. Beyond the fact that simplifying assumptions may be needed, argue for why such a mechanism is impractical for this problem.

3. Recently a lot has been said about the shortcomings of the current Internet albeit its explosive success and growth since the '80s. Many new architectural proposals have come out, including recursive architectures and role-based architectures.
- (a) Identify at least two of the current Internet shortcomings and whether and how exactly these new architectures would address them effectively. Consider at least recursive and role-based architectures.
 - (b) Given a new architecture that addresses such shortcomings, how would you evaluate its effectiveness? Define evaluation metrics and a detailed evaluation methodology that you would follow.
 - (c) Many new wireless technologies, like EV-DO and WiMAX, have become quite sophisticated by employing adaptive error and rate control. For example, some would adapt the service rate based on the channel error rate, so the service rate is decreased for higher channel error rate and vice versa. This kind of link-layer adaptation interacts with the sending-rate adaptations of higher-layer protocols, e.g. transport.
 - i. Develop a dynamical model to capture such interaction between link-layer service-rate adaptation and transport-layer sending-rate adaptation.
 - ii. Define appropriate performance metrics and show plots for them by solving your model analytically (e.g., using basic control theory) and/or numerically (in this case, indicate why your model is intractable).
 - iii. Do new architectural proposals address (or need to address) such protocol interaction scenarios, and how?
 - (d) Heartbeat protocols are essential in discovering new neighbors and confirming old neighbors in a communication network. The frequency of heartbeat messages affects the quality of the resulting topology. This is especially true in a mobile setting where nodes move. High frequency ensures up-to-date topology but incurs high overhead, whereas low frequency has low overhead but may result in an outdated topology.
 - i. To balance such performance-overhead tradeoff, you are asked to design an adaptive heartbeat control so the frequency of heartbeats is increased or decreased as needed, e.g. if the topology is stable, then your scheme should adapt to a low heartbeat frequency. Your scheme should address how the local heartbeat adaptation at a node may interact/influence that of other nodes in the network.
 - ii. Develop a dynamical model to capture such nodal interactions in your protocol.
 - iii. Define appropriate performance metrics and show plots for them by solving your model analytically (e.g., using basic control theory) and/or numerically (in this case, indicate why your model is intractable).
 - iv. Do new architectural proposals address (or need to address) such nodal interaction scenarios, and how?