Optimising node selection probabilities in multi-hop M/D/1 queuing networks to reduce the latency of Tor

S. J. Herbert, S. J. Murdoch and E. Punskaya

Optimising node selection probabilities in multi-hop M/D/1 queuing networks to reduce the latency of Tor

S. J. Herbert, S. J. Murdoch and E. Punskaya

In this paper the expected cell latency for multi-hop M/D/1 queuing networks, where users choose nodes randomly according to some distribution, is derived. It is shown that the resulting optimisation surface is convex, and thus gradient based methods can be used to find the optimal node assignment probabilities. This is applied to a typical snapshot of the Tor anonymity network at 50% usage, and leads to a reduction in expected cell latency from 11.7 ms using the original method of assigning node selection probabilities to 1.3 ms. It is also shown that even if the usage is not known exactly, the proposed method still leads to an improvement.

Introduction: Tor is an anonymity service which routes users traffic through a three-hop network before accessing the Internet [1]. Originally users select a path at random where nodes are chosen with probability proportional to their bandwidth for each hop. This, however, has been shown not to be optimal, apart as the usage tends to 100% [2].

In Tor each cell (packet) is the same size (512 bytes), and treating expected cell latency as a suitable metric.[1] Dingledine and Murdoch have found a method of optimal node assignment using queuing theory, for a one-hop network [3]. The purpose of this work is to generalise this result to the full three-hop representation of Tor.

To achieve this, a general N hop network is optimised. Furthermore arbitrary constraints on which hops a given node may serve are allowed. In Tor there are three types of node: Guard nodes which can serve hops 1 and 2; Normal nodes which can only serve hop 2; and Exit nodes which can serve hops 2 and 3.[2]

Modelling assumptions: In order to formulate the problem mathematically, it is necessary to make some modelling assumptions.

1 Each node has a queuing policy $M/D/1/\infty$/FIFO, according to Kendall's notation [4] (in reality nodes will have finite size queue buffers).
2 The latency is dominated by one direction (the download direction), it is therefore valid to model the network as uni-directional.
3 Each cell is free to probabilistically choose its own path (in reality a user will send all their cells by the same path, however if the number of users is large compared to the number of nodes, which is usually the case in Tor, then the node queues will have behave as if each cell has chosen its path independently).
4 The network status, i.e., nodes online and percentage usage, varies slowly (compared to the time taken for equilibrium queues to be reached) and is well known at all times.
5 Any given cell can be served by the same node for more than one hop (this is not true, but for a network such as Tor with many more nodes than hops it will make a negligible difference to the overall solution).

Optimising node selection probabilities: As previously [3], minimising expected cell latency is deemed a suitable optimisation metric. Let there be n nodes and N hops, the i^{th} node has service rate μ_i and arrival rate $\lambda_{i,j}$ for the j^{th} hop. The total usage is Λ, and the waiting time (i.e., the total waiting time including queuing and service) for the k^{th} hop is W_k. The problem can thus be expressed as a function optimisation.

$$\text{Minimise:} \quad \mathbb{E}\left(W_0 + W_1 + ... + W_{N-1}\right) = \mathbb{E}(W_0) + \mathbb{E}(W_1) + ... + \mathbb{E}(W_{N-1}), \tag{1}$$

$$\text{Subject to:} \quad \sum_{i=0}^{n-1} \lambda_{(i,j)} = \Lambda \qquad \text{all } j, \tag{2}$$

$$\sum_{j=0}^{N-1} \lambda_{(i,j)} = K_i < \mu_i \qquad \text{all } i, \tag{3}$$

$$\lambda_{(i,j)} = 0 \qquad \text{selected } (i,j), \tag{4}$$

where $\mathbb{E}(.)$ denotes expectation, Λ is the total usage and K_i is defined to simplify the notation. (2) states that at each hop, the total arrival rate must equal the usage; (3) states that no node may have a total arrival rate (i.e., for all hops) greater than its service rate; and (4) allows arbitrary constraints on which hops each node may serve.

The Pollaczek-Khinchine formula [5] gives the expected waiting time for a cell at the i^{th} node:

$$\mathbb{E}\left(W_i\right) = \frac{1}{\mu_i} + \frac{K_i}{2\mu_i^2 - 2\mu_i K_i}. \tag{5}$$

Which can in turn be used to find the expected waiting time for the j^{th} hop (note that the probability of the i^{th} node being selected for the j^{th} hop is $\lambda_{i,j}/\Lambda$):

$$\mathbb{E}(W_j) = \sum_{i=0}^{n-1} P(\text{node}_i)\mathbb{E}\left(W_i\right), = \sum_{i=0}^{n-1} \frac{\lambda_{(i,j)}}{\Lambda}\left(\frac{1}{\mu_i} + \frac{K_i}{2\mu_i^2 - 2\mu_i K_i}\right). \tag{6}$$

The shape of the optimisation surface is very important, for convex surfaces a gradient based optimisation algorithm can be used to approach a global maximum. The optimisation surface of $\sum \mathbb{E}\left(W_j\right)$ is indeed convex, and this can be shown by demonstrating that the Hessian matrix is positive semi-definite. First the vector of arrival rates is defined:

$$\boldsymbol{\lambda} = \left[\lambda_{0,0}, \lambda_{0,1}, \cdots, \lambda_{0,N-1}, \lambda_{1,0}, \lambda_{1,1}, \cdots, \lambda_{1,N-1}, \cdots, \lambda_{n-1,0}, \lambda_{n-1,1}, \cdots, \lambda_{n-1,N-1}\right]^T \tag{7}$$

The Hessian, H of $\sum \mathbb{E}\left(W_j\right)$ will have the form:

$$H = \begin{bmatrix} H'_0 & 0 & 0 & \dots \\ 0 & H'_1 & 0 & \dots \\ 0 & 0 & H'_2 & \\ \vdots & \vdots & & \ddots \end{bmatrix}, \tag{8}$$

for H'_0 to H'_{n-1}, each of size $N \times N$.

A necessary and sufficient condition for H to be positive semi-definite, is for H'_i to be positive semi-definite for all i. Noting that differentiation is a linear operator, H'_i is split up into the Hessian for each hop, according to (1), these are named $H'_{i,j}$:

$$H'_i = \sum_{j=0}^{N-1} H'_{i,j}, \tag{9}$$

thus:

$$H'_{i,j}(a,b) = \frac{(\mu_i - K_i)^2 + (\lambda_{i,j} + K_i)(\mu_i - K_i) + \lambda_{i,j}K_i}{\Lambda\mu_i(\mu_i - K_i)^3} \qquad \text{where } a = b = j, \tag{10}$$

$$H'_{i,j}(a,b) = \frac{0.5(\mu_i - K_i)^2 + (\lambda_{i,j} + 0.5K_i)(\mu_i - K_i) + \lambda_{i,j}K_i}{\Lambda\mu_i(\mu_i - K_i)^3} \qquad \text{where } a \neq b = j || b \neq a = j, \tag{11}$$

$$H'_{i,j}(a,b) = \frac{\lambda_{i,j}(\mu_i - K_i) + \lambda_{i,j}K_i}{\Lambda\mu_i(\mu_i - K_i)^3} \qquad \text{where } a, b \neq j, \tag{12}$$

Further splitting $H'_{i,j}$ into the sum of two matrices:

$$H'_{i,j} = H''_{i,j} + H'''_{i,j}, \tag{13}$$

[1] It may in fact be more useful to use cell latency variance or the expected extreme value of latency of a given number of cells as a user will require multiple cells to arrive before a requested piece of data has arrived

[2] It is now possible for nodes to be both Guard and Exit, and thus serve all three hops, however this was not the case when the snapshot of Tor used in this paper was taken

where:

$$H''_{i,j}(a,b) = \frac{\lambda_{i,j}(\mu_i - K_i) + \lambda_{i,j}K_i}{\Lambda\mu_i(\mu_i - K_i)^3} \qquad \text{all } a,b, \tag{14}$$

which is positive semi-definite, because each element is the same and positive in the region $\mu_i > K_i$. Also:

$$H'''_{i,j}(a,b) = \frac{(\mu_i - K_i)^2 + K_i(\mu_i - K_i)}{\Lambda\mu_i(\mu_i - K_i)^3} \qquad a = b = j, \tag{15}$$

$$H'''_{i,j}(a,b) = \frac{0.5(\mu_i - K_i)^2 + 0.5K_i(\mu_i - K_i)}{\Lambda\mu_i(\mu_i - K_i)^3} \qquad a \neq b = j || b \neq a = j, \tag{16}$$

$$H'''_{i,j}(a,b) = 0 \qquad a,b \neq j, \tag{17}$$

The matrix H'''_i is defined by summing $H'''_{i,j}$ over all j:

$$H'''_i = \sum_{j=0}^{N-1} H'''_{(i,j)} \tag{18}$$

Noting that each term on the leading diagonal will equal (15) for exactly one of the matrices $H'''_{i,j}$ and 0 for all the others, and that all terms not on the leading diagonal will equal (16) for two of the matrices $H'''_{i,j}$ and 0 for all the others:

$$H'''_i(a,b) = \frac{(\mu_i - K_i)^2 + K_i(\mu_i - K_i)}{\Lambda\mu_i(\mu_i - K_i)^3} \qquad \text{all } a,b, \tag{19}$$

which is positive semi-definite because each element is the same and positive in the region $\mu_i > K_i$. Finally, expressing the matrix H_i as the sum of the matrix H'''_i and the matrices $H''_{(i,j)}$ i.e., by using (9), (14) and (19):

$$\begin{aligned} H'_i &= \sum_{j=0}^{N-1} H'_{i,j}, \\ &= \sum_{j=0}^{N-1} H''_{i,j} + \sum_{j=0}^{N-1} H'''_{i,j}, \\ &= \sum_{j=0}^{N-1} H''_{i,j} + H'''_i, \end{aligned} \tag{20}$$

which is positive semi-definite, because it is the sum of positive semi-definite matrices, as shown in (14) and (19). Note that the inclusion of the linear equality constraints (2) and (4) does not affect this property.

Numerical example for Tor: The theory is applied to a snapshot of the Tor network. Usage is assumed to be 50%, as is the typical loading of Tor [3]. Fig. 1 shows the optimal node arrival rates for each node type: Guard (G); Normal (N); and Exit (E) for each of the three hops. This leads to an expected cell latency of 1.3 ms compared to 11.7 ms using the original node probability weightings. Notice that many of the lower bandwidth nodes have zero arrival rate (i.e., zero probability of being chosen). Intuitively this can be understood by considering that the minimum waiting time at these nodes (i.e., the service time) is greater than the expected waiting time for the remainder of the node population. Note that the plots in Fig. 1 have a similar shape to that of the one-hop network [3 Fig. 1], providing further evidence that this is likely to be the global minimum.

As well as demonstrating that the proposed method for assigning node selection probabilities significantly outperforms the current method, it is also necessary to consider how robust the proposed method is. This can be achieved by evaluating the expected cell latency for varying usages, given that the selection assignment probabilities have been optimised for 50% (i.e., to mimic the situation where the network status varies and the assignment probabilities haven't been updated accordingly). This leads to the result that the proposed method optimised for 50% usage outperforms the current method between 0 and 60% usage, however for usages greater than approximately 62% the arrival rate exceeds the service rate for some nodes, and thus the expected cell latency tends to infinity.

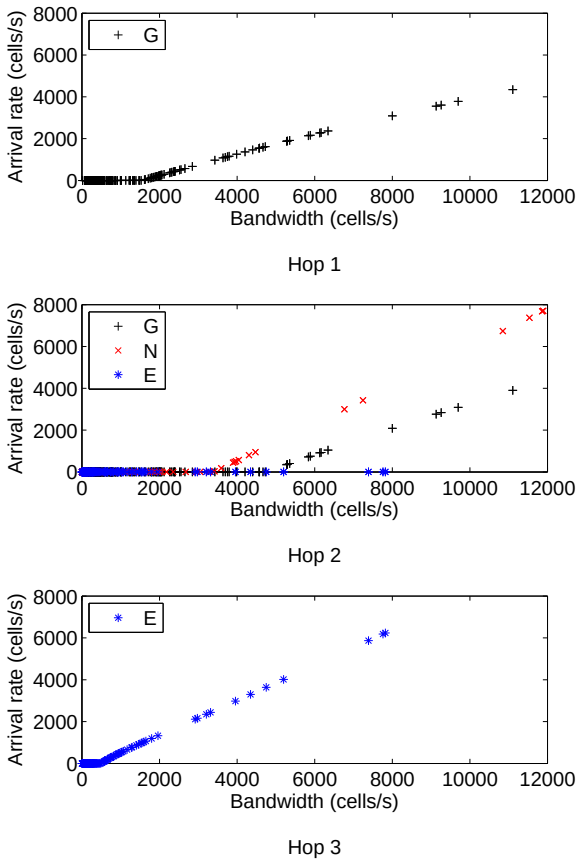

Fig. 1. Optimal arrival rates for snapshot of Tor nodes at 50% usage

Conclusion: A general solution to minimising expected cell latency in multi-hop M/D/1 queuing networks has been derived, and it has been shown that the optimisation surface is convex. This has been applied to the Tor anonymity network at 50% usage, and it has been shown that the expected cell latency can be reduced from 11.7 ms with the original node selection probability method (i.e., the node selection probability is proportional to its bandwidth) to 1.3 ms with the proposed method. Furthermore it has been shown that the proposed method and leads to a reduced latency even if the usage is not known exactly.

The derivation has assumed that nodes have infinite queue capacity, which is not the case in reality. Therefore it would be interesting to run a network simulation with nodes with finite buffer queues to verify that the proposed method would actually lead to improved results in an actual network.

S. J. Herbert, S. J. Murdoch and E. Punskaya (*University of Cambridge, UK*)

E-mail: Steven.Herbert@cl.cam.ac.uk

References

1 Dingledine, R., Mathewson, N., and Syverson, P.: 'Tor: The Second-generation Onion Router', *USENIX security symposium*, 2004, p. 21

2 Murdoch, S. and Watson, R.: 'Metrics for Security and Performance in Low-Latency Anonymity Systems', *International symposium on privacy enhancing technologies*, 2008, p. 115-132

3 Dingledine, R. and Murdoch, S.: 'Performance Improvements on Tor, or why Tor is slow and what we're going to do about it', *Technical report*, 2009

4 Kendall, D., 'Stochastic Processes Occurring in the Theory of Queues and their Analysis by the Method of the Imbedded Markov Chain', *The Annals of Mathematical Statistics*, 1953, p. 338-354

5 Khinchin, A., 'Mathematical theory of a stationary queue', 1932