Minimizing Routing State for Light-weight Network Simulation

Polly Huang
ETH Zurich
ETH-Zentrum, CH-8092
Switzerland
huang@tik.ee.ethz.ch

John Heidemann
USC/ISI
Marina del Rey, CA 90292
USA
johnh@isi.edu

Technical Report
TIK-Nr. 106

Abstract

We introduce an $O(N)$ routing mechanism, referred to as *algorithmic routing*. It minimizes the memory consumption for routing states in network simulations. Our computation shows that for network simulations size of the Internet, algorithmic routing can potentially reduce memory requirement by *three orders of magnitude* to hierarchical routing in Internet-like hierarchy and an *additional three orders of magnitude* to flat shortest path routing. To evaluate the efficiency of algorithmic routing in practice, we implement the mechanism on *ns-2*. With network topologies randomly generated by GT-ITM, we compare the memory and run-time consumption of flat, hierarchical and algorithmic routing and confirm the expected performance improvement by algorithmic routing. We also quantify difference of the three mechanisms in route length (i.e., hop counts). We find the relative difference is below 10% for more than 80% of the routes. Nevertheless, we are cautious about the applicability of algorithmic routing to general network simulations, given that it is not clear to what extent the GT-ITM topologies capture the hierarchical structure of Internet and what impact the different routes may incur. Thus, we further rationalize distortions that algorithmic routing may introduce and recommend users to apply this technique only when suitable.

Keywords: Network Simulation, Abstraction Technique, Routing

1 Introduction

One major bottleneck in large-scale network simulation is the memory requirement for routing states. As of January 2000, we observe at least 284,805 routers in the Internet [20] and some 6,474 of them situated in the backbone [16]. Assuming each routing entry consumes 4 bytes of memory, network simulations of this size will require a minimum of 360GB memory for the Dijkstra all-pair shortest path routing [5]. Its space

complexity is $O(N^2)$. Hierarchical routing scales in O($NlogN$) for balanced networks with $logN$ levels of hierarchy [4]. The Internet however runs in a two-level hierarchy where the backbone runs one routing protocol, e.g, BGP [18], also of space complexity O(N^2), and each domain connected by domain border routers to the backbone runs another routing protocol, e.g., OSPF [15] or RIP [9]. In this case, the memory requirement to maintain routing states for the backbone dominates the total consumption and it is approximately 200MB.

To further minimize memory requirement for routing, we propose an O(N) routing mechanism, which reduces the memory requirement to 120KB. This mechanism was inspired by previous work of Raman et. al. [17], in which the authors evaluated scalability and behavior of a reliable multicast mechanism in large-scale binary tree networks. For their simulations, they used a simple algorithm to *compute* next hop for any source and destination pair without maintaining a routing table, thus the name – *algorithmic routing*. We extend the algorithm to apply to k-ary trees. To map an arbitrary network topology into a k-ary tree, we adopt the Breath First Search (BFS) algorithm [4] and re-assign node addresses in an *orderly* fashion. This order then allows simple next hop computation without maintaining a routing table. The $O(N)$ cost in memory space actually comes from maintaining the address mapping between the original topology and the corresponding k-ary tree. Despite effective in reducing memory requirement by several orders of magnitude, algorithmic routing comes with a price. That is, routes may not always be the shortest. Consider ring components in the networks. In algorithmic routing, certain links in these ring components will be ignored as if they do not exist. This artifact results in sub-optimal routes.

To evaluate algorithmic routing in practice, we implement the mechanism on *ns-2* [3], a commonly used network simulator in the research community. In addition to comparing the actual memory and run-time consumption with other routing mechanisms, we also quantify the degree of sub-optimality introduced by algorithmic routing. We use a set of transit-stub topologies generated by GT-ITM [26] for the evaluation. To ensure simulation results reflect true behavior of Internet, we discuss limitations of algorithmic routing for general network simulations and recommend simulation scenarios that are suitable or not suitable for applying the algorithmic routing technique.

In short, the contribution of this work is (a) a generalization of Raman et al.'s work to arbitrary topologies, (b) an evaluation of the memory and time savings and potential inaccuracies of our approach, and (c) recommendations for simulation scenarios that algorithmic routing is applicable or not applicable.

2 Background and related work

We provide in this section background on Internet routing, routing abstraction in network simulation and related work in scalable simulation techniques.

2.1 Routing in the Internet

The physical Internet is a collection of routers interconnected by links. Users usually access the network through the end systems (or hosts) that are connected to routers

through access media, for example Ethernet. A contiguous collection of routers under one administrative authority is called an administrative domain (or an *Autonomous System*).

OSPF [15] and RIP [9] are two popular protocols to route packets within an administrative domain. They are different in a sense that each OSPF router floods the domain with states of its neighboring links (thus also known as a *link state* routing protocol) whereas each RIP router distributes to neighboring routers its vector of distances to every other routers in the domain (thus also known as a *distance vector* routing protocol). Otherwise, the two protocols both converge to shortest routes for all source and destination pairs in the domain, no matter whether it is distance vector or link state information their computation is based upon.

The domain-wide flooding and the $O(N)$ routing table per router do not scale to the entire Internet. Thus, BGP [18] is introduced to route packets across domains. Border routers running BGP aggregate routing information per domain and exchange this per-domain routing information with the neighboring domains, in a fashion similar to distance vector routing. This routing information aggregation alleviates the scaling problems with message flooding and routing table size. BGP in principle converges to shortest routes unless specific paths are given based on domain policy (thus also known as a *path vector* routing protocol).

2.2 Routing in network simulation

Some routing protocols in simulators implement details of routing information exchange. However many simulators can also compute routes instantly in a centralized fashion [10] when details of message exchange are not crucial and memory or computation resource is scarce. Flat and hierarchical routing in *ns-2* [3] are two examples of such abstraction technique for light-weight network simulations.

Flat routing in *ns-2* performs the Dijkstra shortest route computation. In that, each node maintains one adjacency state to each other node in the topology. Hierarchical routing in *ns-2* is similar to the flat routing in a sense that it also does Dijkstra-style iteration by walking through all nodes and eventually settling on the shortest routes. Thus, this particular form of hierarchical routing also yields shortest routes for all source and destination pairs. The difference is that, if there are three levels of hierarchy, each node maintains one adjacency state to each other node (level-1) in a local cluster, to each other cluster (level-2) in a domain, and to each other domain (level-3) in the topology. This form of hierarchical routing is different from that of the Internet where the backbone and local domains compute shortest routes independently. There is sometimes an inter-operating protocol, e.g., IBGP [2], operating among border routers within a domain. These IBGP routers are both backbone-level and local routers and have both backbone-level and local routing information. They correspond to this cluster level (level-2) in the hierarchical routing in *ns-2*. Below we analyze various forms of routing by their space and time complexity.

The flat routing generates a routing table that each node has the next hop and cost information to every other node. The space complexity is $O(N^2)$. Each node in *ns-2* hierarchical routing maintains routes to nodes within a local cluster, to other clusters within the domain, and so on. This results in $O(kNlog_kN)$ space complexity on av-

Dijkstra Flat	$O(N^2)$
Dijkstra Hierarchical	$O(N\sqrt[3]{N})$
Internet Hierarchical	$O(N\sqrt[3]{N})$
Algorithmic Routing	$O(N)$

Table 1: Space Complexity

	Per Setup	Per Lookup
Dijkstra Flat	$O(N^3)$	$O(1)$
Dijkstra Hierarchical	$O(N^2\sqrt[3]{N})$	$O(1)$
Internet Hierarchical	$O(N\sqrt[3]{N^2})$	$O(1)$
Algorithmic Routing	$O(N)$	$O(logN)$

Table 2: Time Complexity

erage, where k is the rank (i.e., number of sub-domains per domain) and $log_k N$ is the height of the hierarchy. Given that Internet hierarchy in this routing sense is generally 3-level high (i.e., $log_k N = 3$), k is $\sqrt[3]{N}$. Thus the average space complexity for hierarchical routing in *ns-2* is $O(N\sqrt[3]{N})$. As for Internet hierarchical routing, when ignoring IBGP specific states, the space complexity is $O(kN)$, where each node in a k-node group ($\frac{N}{k}$ of these groups) maintains only routing states to other k local nodes. $O(kN)$ equals $O(N\sqrt[3]{N})$ with a fixed 3-level hierarchy. The $O(N)$ space complexity of algorithmic routing comes from maintaining mapping between the original topology and the corresponding k-ary tree. This will become clear in Section 3 where algorithmic routing is explained in detail. Table 1 summarizes the space complexity analysis.

Computational complexity for flat routing, which adopts the Dijkstra algorithm, is $O(N^3)$. In that, each node (N) searches for a best route to each other node (N) through existing routes known by these other nodes (N). Computational complexity for hierarchical routing in *ns-2* is $O(N^2\sqrt[3]{N})$. In that, each node (N) searches for a best route to each of the local nodes, clusters, and domains ($\sqrt[3]{N}$) through existing routes in each other node (N) known to these local nodes, clusters and domains. As for Internet hierarchical routing, when ignoring IBGP specific computation, the time complexity is $O(k^2N)$, where each group of k nodes ($\frac{N}{k}$ of them) does a $O(k^3)$ Dijkstra computation. $O(k^2N)$ equals $O(N\sqrt[3]{N^2})$ with a fixed 3-level hierarchy. Computational complexity for algorithmic routing is $O(N)$. This is because the BFS algorithm and address re-assignment are both $O(N)$ processes. The computation cost of $O(logN)$ per route lookup is shifted to the traffic generation phase and is not included in the route establishment phase (the period of simulation we monitor for results described in Section 4). Table 2 summarizes the time complexity analysis.

2.3 Scalable simulation technique

Parallelism can improve simulation scale in ratio to the number of machines added, but this linear growth is not sufficient to add several orders of magnitude scaling needed.

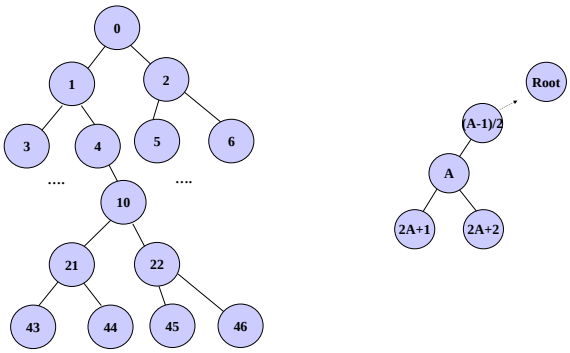

Figure 1: Left: An Example Binary Tree Right: Regular Node Address Assignment

There is increasing interest on taking a complimentary solution. Just as a custom simulator includes only details necessary for the task at hand, a general simulator can support configurable levels of details for different simulation studies. This is also referred to as the *selective abstraction* approach. Although abstract simulations are often more efficient, they are not identical to more detailed ones. It is critical that when proposing abstraction techniques one also explores the potential impact or distortion these abstraction techniques can introduce, so the users can avoid applying techniques that would lead to invalid conclusions.

Most abstraction techniques focus on enabling simulations with a large amount of traffic, by aggregating individual packets into coarser-grained packet trains [1, 11, 8] or fluid flows [14]. We are interested in resolving the scaling problem in the routing element of network simulations. Riley et. al. [19] proposed to compute route on demand when a new packet is generated. Routes are cached per source and destination pair to prevent same routes from being computed again and again. This mechanism reduces time and memory consumption significantly at route setup phase because it does not compute routing table at all. However, when traffic starts, it involves an $O(N)$ time consumption per route computation and potentially $O(N^2)$ memory consumption if traffic needs to be generated across all source and destination pairs. We show in the remaining part of the paper that our abstraction technique, algorithmic routing, effectively reduces memory consumption to $O(N)$ independent of traffic pattern. Algorithmic routing also reduces time consumption to $O(N)$ for one-time route setup and $O(logN)$ per route lookup. We discuss as well simulation scenarios that are suitable or not suitable for applying this particular abstraction technique.

3 Algorithmic routing

There are three components in algorithmic routing. The first component involves a simple algorithm to compute next hop in a binary tree and its extension to a more general k-ary tree. The second component looks to map an arbitrary network topology into a k-ary tree which in turn can be used for algorithmic lookup. The last component

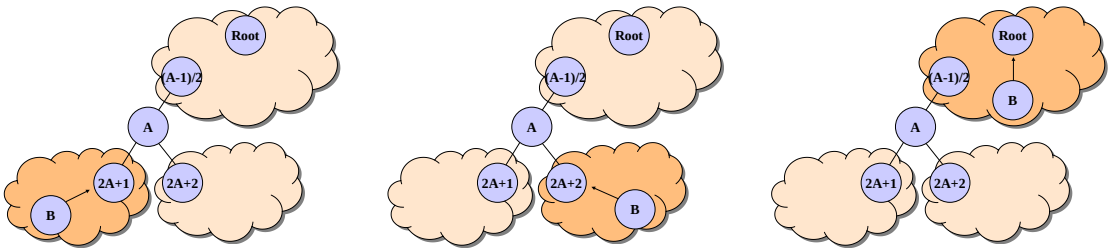

Figure 2: Left: B in The Left-child Domain Middle: B in The Right-child Domain Right: B in The Parent Domain

defines a three-step procedure for general route lookup. For any source and destination pair, we first find their corresponding addresses in the mapped k-ary tree, compute the next hop, and then find the address of the identified next hop in the original topology. We elaborate details of the three components in the subsections to come.

3.1 Route lookup in k-ary tree

The first component is best explained with a binary tree. Figure 1 gives an example of a binary tree (see left plot) and the formula for address assignment (see right plot). For any source A, a destination B can fall in three possible regions. See Figure 2. Node $2A+1$, $2A+2$, or $(A-1)/2$ would be the next hop depending on which region B resides.

We can find the next hop from A to B by walking B to the root, a process of $O(log_2 N)$. If we pass by $2A+1$, the next hop is $2A+1$. If we pass by $2A+2$, the next hop is $2A+2$. If we reach the root without passing by any of the two, the next hop is $(A-1)/2$. Walking from B to the root is a recursive $(B-1)/2$ computation.

This algorithm can be extended to k-ary tree as follows:

```
next_hop (A -> B):
        while (B > 0)
                B_parent = (B-1)/k
                if (B_parent == A) return B
                B = B_parent
        return (A-1)/k
```

With average computational complexity $O(log_k N)$, this simple algorithm let us *compute* next hop for any source and destination pair without maintaining a routing table, and it works as long as the topology is a k-ary tree with regular address assignment. This is not the case for most network topology of interest. Network topology is often arbitrary. We address this problem by applying a tree search algorithm on any arbitrary topology to obtain the corresponding k-ary tree. Details are described in the following subsection.

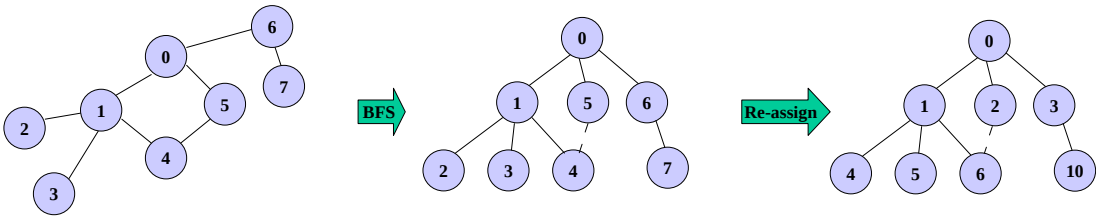

Figure 3: Left: an example arbitrary topology Middle: after BFS mapping Right: after re-assigning node addresses

3.2 Tree mapping

To enable algorithmic lookup, we use a tree search algorithm to define the k-ary tree equivalent for an arbitrary topology. We choose Breath First Search (BFS) for the reason that it produces least height trees, i.e., routes in a least-height tree are likely shorter. This is crucial and will become clear in Section 3.4 and 4.4 when discussing sub-optimal routes as results of algorithmic routing. The tree mapping starts with the lowest addressed node as the root of BFS. BFS algorithm traverses all the immediately connected nodes, and then recursively starting from each of these immediately connected nodes traverses its immediately connected nodes until all the nodes are visited. While the original topology is converted into such a tree, its rank (maximal number of leaves per node) can be recorded and used as the value of k for the k-ary tree. Subsequently, the node addresses are re-assigned to be a k-ary tree with possibly some empty leaves. See Figure 3 for an example. As a result of this tree mapping process, we maintain a one-to-one mapping of the node addresses between the original topology and mapped k-ary tree. This mapping is in the order of $O(N)$ in space complexity and is in fact the dominant factor in the entire memory consumption for algorithmic routing. The BFS and address re-assignment is $O(N)$ in computational complexity.

3.3 Route lookup in general topology

With the k-ary tree mapping and lookup algorithm in hand, general next hop lookup from node A_{orig} to B_{orig} is a three-step procedure. First, find the corresponding addresses of node A_{orig} and B_{orig} to the k-ary tree, say A_{tree} and B_{tree}. Second, compute the next hop node from A_{tree} to B_{tree}, say C_{tree}. Third, find the corresponding address of C_{tree} from the k-ary tree to the original topology, C_{orig}. Through the process, we identify the next hop C_{orig} for any source A_{orig} and destination B_{orig} pair in an arbitrary topology. The process can be repeated until defining the complete route from A_{orig} to B_{orig}.

Combining the three components, we derive the algorithmic routing mechanism which is $O(N)$ in space complexity, and $O(N)$ per simulation setup and $O(log_k N)$ per route lookup in computational complexity. Algorithmic routing also gives exactly

one route for any source and destination pair and this route may not always be optimal, or the shortest path in the original topology. We explain the sub-optimal route problem and its implication to general network simulation in the following subsection.

3.4 Sub-optimal route

In a tree topology, there is exactly one route for any source and destination pair whereas in a topology containing cycles (or is cyclic), there exist multiple routes. During the tree mapping process, certain links in the original topology are ignored. For example, link 4-5 in Figure 3. These links are those used to be part of a cycle in the original topology. Thus, in case the tree mapping happens to cut those links on the shortest routes for certain source and destination pairs, for these source-destination pairs algorithmic routing will yield sub-optimal routes. For example, in Figure 3 the shortest route from node 4 to 5 with algorithmic routing is via node 1, not directly to node 5 anymore.

In other words, algorithmic routing may not capture simulation relevance to some of the links. In particular, for data transfer in between source and destination pairs that links on the original shortest routes are cut, the network delay may be longer. Thus, when studying quality of service control mechanisms that deal with strict end-to-end delay guarantee, simulations using algorithmic routing could result in over-conservative design. Over-conservative designs may leave the network under-utilized but end-to-end delay guarantee would be met. In a sense, although simulations using algorithmic routing do not yield precise results, they are reasonable as worst-case analysis.

In addition, in algorithmic routing every node has exactly one route to every other node. One can expect a higher degree of traffic concentration and sometimes unintended network congestion at the root of the mapped k-ary tree. Thus when studying congestion control protocols, simulations using algorithmic routing could result in over-estimation of congestion level and users may end up with back-off mechanisms that are over-sensitive. Again, although simulations using algorithmic routing may not be precise, they give reasonable worst-case results.

For special cases where there is only one sender in the simulated scenario, we can start BFS tree search at the sender as the root of the k-ary tree. In this case all routes will be exactly the same as those in the shortest path computation. For simulations with few senders, multiple k-ary trees can be maintained, one per sender. In this case, memory consumption will be slightly higher depending on the number of senders there are, i.e, $O(sN)$, where s is the number of senders in the simulations, but simulations using algorithmic routing are exactly the same as those using shortest path routing.

3.5 Summary

In this section, we describe the major components of the algorithmic routing, explain how the simple lookup and BFS algorithms can help us identify next hops in a general topology without maintaining any routing table. The memory complexity of $O(N)$ actually comes from maintaining the mapping between the general topology and the corresponding k-ary tree. Although algorithmic routing introduces an extra $O(log_k N)$

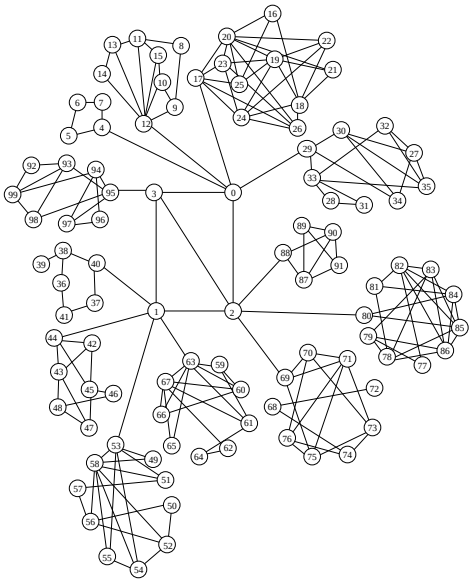

Figure 4: An example GT-ITM transit-stub topology with 100 nodes

computational cost per route lookup, for a serial simulator such as *ns-2*, it would be a reasonable trade-off to be able to run the simulations at all with limited memory space per machine.

Algorithmic routing also introduces distortion where routes for certain source and destination pairs can be sub-optimal. While this distortion may affect simulation studies on end-to-end delay estimation and congestion control, the fact that this sub-optimal route problem consistently over-estimates the level of end-to-end delay and congestion makes simulations using algorithmic routing reasonable worse-case analysis. For special cases where there are one or few sources, algorithmic routing gives exactly the same results as the shortest path routing. To back up our analysis, we quantify in Section 4, the efficiency and accuracy of algorithmic routing in practice using the *ns-2* simulator.

4 Evaluation

We implement algorithmic routing on *ns-2* which has already a set of implementation for flat (or shortest path) and hierarchical routing. Using random topology generated by GT-ITM, we evaluate memory and time usage of the three routing mechanisms. To quantify degree of distortion, we compare path length for all source and destination pairs. Details of the simulation setup and analysis are provided below.

4.1 Methodology

In this subsection, we describe implementation of the three routing mechanisms in *ns-2*, GT-ITM random topologies, measurement metrics, simulation scenarios, and software/hardware platform of our experiments. In *ns-2*, the flat routing is implemented in the `C++` space whereas a major portion of hierarchical routing is implemented in the `OTcl` [23] space. We choose to implement algorithmic routing in `OTcl` space due to development time constraint.

To analyze how each of these routing mechanisms scales to the size of network topology, we generate random topology with increasing sizes, from 100 nodes to 500 nodes. For each size, we simulate one random topology. We use GT-ITM to generate the random topologies. We are particularly interested in the transit-stub type (Figure 4) because it attempts to model the hierarchical structure of Internet. We change parameters proportionally to generate transit-stub topologies in different sizes while holding the average connectivity constant, approximately 1.77.

At the end of the simulation, we measure the total memory and runtime consumption. To quantify the effect of sub-optimal routes, we measure path length for all possible source and destination pairs and compute the relative difference of algorithmic routing to the other two routing mechanisms.

All simulations end when route computation is completed, i.e., no traffic is generated. To highlight the memory and run-time cost in the routing component of simulation, we conduct all simulations in session mode [10], where node and link structures are reduced (or abstracted) to conserve memory. All memory and time consumption data are collected by simulating the scenarios on a PC with a Pentium II 450MHz CPU and 1 GB RAM, running FreeBSD 3.0.

Characteristics of the GT-ITM topology is critical in this study, given that different graph constructs could result in different performance and distortion measurement. We address here whether GT-ITM transit-stub topologies reasonably capture characteristics of the Internet. Previous work [6] has unveiled that topology of the Internet backbone exhibits power-law decay. Analysis from preliminary measurement observed [7] similar power-law decay in router-level Internet topology. However, it is identified [13, 12] that GT-ITM random topologies do not show clear power-law decay. Although the evidences suggest that GT-ITM random topology model might be inadequate, it is not clear yet if the measurement of router-level Internet topology does cover a reasonable base (difficult to validate as well). It is under investigation whether the router-level Internet topology inherits the power-law properties from its backbone component. Even if it is so, simulated topology must contain a large number of nodes and links (1000s of them) to bring out the power-law statistical properties. Thus, for small- to median-sized topology, GT-ITM is deemed still a reasonable tool [25].

4.2 Memory Usage

Figure 5 shows the memory usage for one simulation because after repeating the experiment we find that the results are deterministic. As we increase number of nodes included in the topologies, the memory requirement for flat, hierarchical and algorithmic routing mechanisms increase as well. Although the expected growth of flat routing

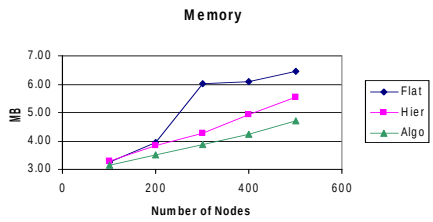

Figure 5: Memory Consumption Comparison

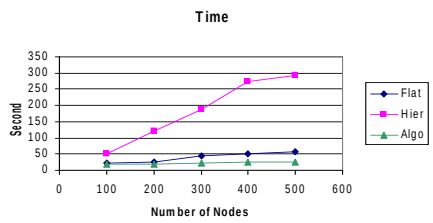

Figure 6: Time Consumption Comparison

should be proportional to N^2, it experiences a faster than N^2 jump every 2^k nodes, and then flattens out between 2^k and 2^{k+1}. This is due to the memory allocation policy in FreeBSD's `C` library which interacts with *ns-2*'s flat routing. The number of entries per node in the routing table is always 2^k, where k is an integer. When the number of nodes increases to $2^k + 1$, *ns-2* will allocate another 2^k entries. In this case, $2^k - 1$ entries are not used. This means the memory requirement increases by $(2^k)^2$, i.e, 4^k, at these jump points. This *power of 2* artifact contributes to the jumping-4^k-and-then-flattening-out behavior. We expect to see another jump between node 500 and 600 (jumping from 512 entries to 1024 entries). Memory requirement for hierarchical routing increases in a significantly slower rate roughly in the order of $N\sqrt[3]{N}$ (Table 1). Memory requirement for algorithmic routing is even lower in the order of N.

4.3 Time Usage

Figure 6 shows surprising results that flat routing, supposedly the least scalable routing mechanism, runs faster than hierarchical routing and only slightly slower than algorithmic routing. By investigating this unexpected phenomenon, we discover that simulation speed is closely related to two important factors. One is analytical complexity and the other is programming language efficiency. Sometimes, language efficiency can be as crucial as analytical complexity in determining simulation speed. From the analytical results, flat routing should scale the worst in the order of $O(N^3)$, hierarchical

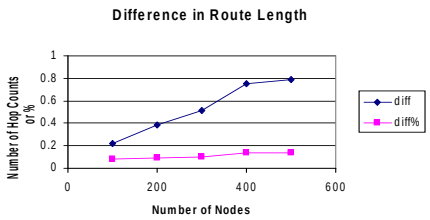

Figure 7: Distortion by Algorithmic Routing: Difference in Route Length

Number of Nodes	Mean %	Std Dev %	Median %	Max %
100	7.954	29.817	0	400
200	9.236	28.395	0	600
300	10.314	25.888	0	700
400	13.368	24.323	12.5	600
500	13.539	21.610	12.5	700

Table 3: Large Maximum and Small Median in Relative Difference

routing second in $O(N^2 \sqrt[3]{N})$, and algorithmic routing the best in $O(N)$ (Table 2). As a matter of fact, simulation results (see Figure 6) for hierarchical and algorithmic routing agree reasonably well with the analytical results. However, flat routing scales surprisingly better than expected, much better than hierarchical routing and only slight worse than algorithmic routing. This is because that flat routing's Dijkstra algorithm is implemented in `C++`, a much more efficient language than `OTcl`, in which hierarchical and algorithmic routing mechanisms are implemented. In this case, efficiency of programming language overtakes and dominates simulation performance. We plan to investigate further the effect of programming language inefficiency.

4.4 Distortion

By mapping arbitrary topologies into trees, we ignore certain links that otherwise are part of the cycles in the original topology. This results in sub-optimal routes (Section 3.4). For each transit-stub topology, we run the same simulation using flat, hierarchical and algorithmic routing. For each source and destination pair, we compute the path length difference in hop counts. Relative difference indicated in percentage is the ratio of the difference to the path length in flat (or shortest path) routing. Due to the nature of the hierarchical routing in *ns-2*, i.e., routes are always the shortest, we do not observe any difference in route length between the flat and hierarchical routing. For the rest of the section, we compare routes obtained by algorithmic to the shortest path routing.

Figure 7 depicts the average absolute (`diff`) and relative difference (`diff%`) in

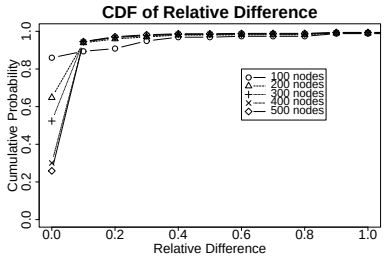

Figure 8: Distribution of Relative Difference in Route Length

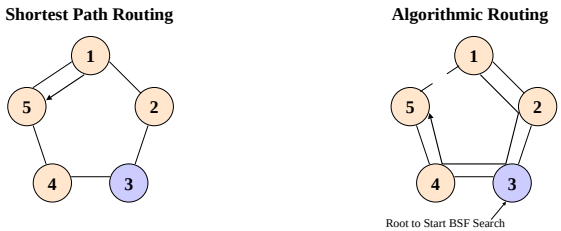

Figure 9: An Example of Route Distortion Using Algorithmic Routing

route length. Each data point is obtained by averaging over absolute or relative differences of all source and destination pairs. The average relative difference increases slowly within the range of 8-14%. Table 3 shows that most routes have the same length (low median) and that few routes have very large differences (large maximum, up to 700%). Although median seems to increase when topology size increases, Figure 8 shows that independent of topology size increase, a high percentage (80-90%) of routes stay within 10% difference.

We hypothesize that the relative difference in path length is related to the size and the amount of ring components (or cycles) in a topology. To be more precise, if the sizes of the ring components are large, the maximal and average relative differences will be large; if the amount of ring components increases, the average relative difference will also increase. Consider a ring topology with 5 nodes. See Figure 9. The shortest path from node 1 to 5 is to go directly through link 1-5 (Figure 9, left). After this ring topology is mapped to a string, the shortest path from node 1 to 5 becomes the path through link 1-2, 2-3, 3-4, and 4-5 (Figure 9, right). The route length difference is 3 hops and the relative difference is 300%. One can imagine that the maximal relative difference would be determined by the largest ring component in the topology.

Additionally, if there are more ring components, there will be more routes having different lengths which brings up the average difference and possibly also the median of differences. With preliminary analysis (See Appendix A), we identify a sound property on the average relative difference of algorithmic to shortest routes in ring topology. That is, for a ring of N nodes, the average relative difference is exactly $\frac{k}{2k+1}$, where

k is the largest integer smaller than $\frac{N-1}{2}$ (Formula 3). From the formula, we obtain a theoretical upper bound of average relative difference for rings – 50%. The formula also indicates that a small ring of three nodes introduces 33% of relative differences by itself. For the transit-stub topologies we experimented, there must be a substantial amount of routes remaining the same or not much different to bring the average down to 8-14%. This explains in Figure 8 that 80-90% of routes fall within 10% relative difference.

Our analysis confirms that algorithmic routing yields sub-optimal routes. However, the degree of sub-optimality (i.e., relative difference) for most routes is small, which means a good portion of communication over the topology is minorly affected. The difference could be negligible for worst-case analysis. The average relative difference potentially allows users to compensate the measured end-to-end delay results from simulations using algorithmic routing. Further analysis has to be done to verify that our hypothesis of larger or more ring components results in larger route length difference. Experiments need to be conducted to understand more on the effect of algorithmic routing to actual Internet topology and to quantify as well the relative difference to Internet-like hierarchical routing which may involve non-trivial (sometimes complicated) domain policy. We address these issues in the future work in Section 6.

4.5 Summary

We implement algorithmic routing on *ns-2* which has already a set of implementation for flat (or shortest path) and hierarchical routing. Using random topology generated by GT-ITM, we evaluate memory and time usage of algorithmic routing to flat and hierarchical routing. This set of results confirm that the three routing mechanisms scale approximately in the order of O(N^2), O($N\sqrt[3]{N}$), and O(N) in terms of memory consumption which demonstrate that algorithmic routing helps achieving the objective of minimizing the routing states for light-weight network simulations. Our time consumption results confirm that the hierarchical and algorithmic routing mechanisms scale in the order of O($N^2\sqrt[3]{N}$) and O(N) respectively. The O(N^3) flat routing surprisingly out-performs hierarchical routing. This is due to the fact that flat routing in *ns-2* is implemented in `C++` whereas the other two in `OTcl`, and in this case, efficiency of programming language becomes the dominant factor.

To quantify degree of distortion, we again experiment with GT-ITM generated topology and compare path lengths for all source and destination pairs in algorithmic and shortest path routing. This set of results show that there are occasional large relative difference for few source and destination pairs, but mass majority of the relative differences are small. Depending on the problem in study, these differences might be negligible (e.g., worst-case analysis) or simply not have any impact at all (e.g., single- or few-sources analysis). Given our preliminary analysis on graphs with tendency of higher degree of distortion, we recommend users to avoid especially applying the algorithmic routing technique on networks that contain large ring-like components.

5 Conclusion

In summary, we have described a routing mechanism whereby large-scale networks can be simulated with only an $O(N)$ amount of routing states. When applying to networks size and structure of the Internet, *algorithmic routing* has the potential of reducing memory consumption by three orders of magnitude to hierarchical routing and six orders of magnitude to flat routing. However, simulations using algorithmic routing are not identical to those using hierarchical or flat routing. Just as other abstraction techniques distort simulation results one way or another, we identify a bounded amount of distortion in route length as well by algorithmic routing. Through analysis and simulation, we come to the following conclusions.

- Algorithmic routing does scale in $O(N)$ in both memory and time consumption.
- Algorithmic routing has an extra $O(logN)$ computational cost per route lookup.
- Average degree of distortion in route length is small (8-14%).
- For mass majority of the routes (80%+), relative differences remain small (10%).

Although there is an $O(logN)$ computational cost per route lookup, algorithmic routing is still desirable for serial simulators such as *ns-2*, where large-scale simulations may not be performed at all with a limited amount of memory. Besides, with elegant data structure [21], this lookup can be done in $O(1)$ time. The last two bullets imply that the degree of distortion scales well to topology size and it is likely that most communication is done through routes that are only slightly different.

For simulation scenarios that have single or few senders, multiple k-ary tree mappings can be maintained, one per sender, so that routes are always the shortest. If the simulated topologies are trees, algorithmic routing yields exactly the same results. Simulations using algorithmic routing consistently over-estimate metrics such as end-to-end delay and congestion level. Therefore, simulations using algorithmic routing qualify as worst-case analysis. With our simulation and analytical results on the degree of distortion, one can compensate the overall simulation results proportionally for more precise understanding of the problem in study. In fact, there exist spanning tree algorithms [24] that compute in polynomial time and give bounds to the degree of distortion (i.e., average relative difference). In short, algorithmic routing is suitable for any of the following cases.

- simulated scenarios contain few senders,
- simulated topologies are trees, or
- simulation objective is to assess the worst-case performance.

Otherwise, we would recommend avoid using algorithmic routing for cases that satisfy any two of the following conditions.

- simulated scenarios contain many senders,
- simulated topologies contain large and many ring-like components.
- results have to be precise

6 Future work

We propose to continue this work in three ways. The first proposal is inspired by the fact that single-sender simulations can start tree mapping from the sender node to avoid distortion at all. One can maintain one tree mapping per sender for few-sender simulations or per node on a large ring component (or alternatively nodes on opposite sides of a large ring). This results in $O(sN)$ memory consumption but is a reasonable alternative to trade off memory for higher degree of accuracy, given s remains small. The second proposal aims to reduce the maximal relative difference. Given that every node has information of directly connected links anyway, we propose to perform a 2-pass route lookup where we check first if the destination is directly connected to current node If not, we look up using algorithmic routing in the ordinary fashion. Thirdly, we expect a lower degree of distortion with a slightly modified tree search algorithm. BFS algorithm results in least-height tree which is already a desirable property to avoid sub-optimal routes, i.e., routes tend to be shorter in general. Current implementation of BFS is in its most naive form. For nodes at the same level in tree, including the root at the highest level, we search based on node address. We propose to give higher precedence to nodes with higher degree of connectivity and search nodes at the same level based on the precedency. This will potentially bring nodes connecting to high-connectivity nodes higher in the tree and reduce average route difference in general.

We plan to incorporate these changes in the near future in the mean time we port the current `OTcl` implementation to `C++`. For the long-term future, formal analysis need to be done to verify the hypothesis of larger or more ring components results in larger route difference. The hierarchical structure of Internet and its, sometimes, policy-based routing are currently subjects under investigation [22]. Once these become clear, experiments can be conducted to understand the effect of algorithmic routing to actual Internet topology and routes.

References

[1] J. Ahn and P. B. Danzig. Speedup vs. simulation granularity. *IEEE/ACM Transactions on Networking*, 4(5):743–757, October 1996.

[2] T. Bates and R. Chandra. BGP Route Reflection - An alternative to full mesh IBGP. RFC 1966, Internet Request For Comments, June 1996.

[3] L. Breslau, D. Estrin, K. Fall, S. Floyd, A. Helmy, J. Heidemann, P. Huang, S. McCanne, K. Varadhan, H. Yu, Y. Xu, and VINT Project. Advances in network simulation. *IEEE Computer*, 33(5):59–67, May 2000.

[4] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. McGraw-Hill, New York, 1989.

[5] E. W. Dijkstra. A Note on Two Problems in Connection with Graphs. *Numerical Mathematics*, 1:269–271, 1959.

[6] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-Law Relationships of the Internet Topology. In *Proceedings of the ACM SIGCOMM*, Cambridge MA, August 1999.

[7] R. Govindan and H. Tangmunarunkit. Heuristics for Internet Map Discovery. In *Proc. IEEE INFOCOM*, 2000.

[8] Y. Guo, W. Gong, and D. Towsley. Time-stepped Hybrid Simulation (TSHS) for Large Scale Networks. In *Proceedings of the IEEE Infocom*, March 2000.

[9] C. L. Hedrick. Routing Information Protocol. RFC 1058, Internet Request For Comments, June 1988.

[10] P. Huang, D. Estrin, and J. Heidemann. Enabling Large-scale simulations: selective abstraction approach to the study of multicast protocols. In *Proceedings of the International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, pages 241–248, Montreal, Canada, July 1998. IEEE.

[11] P. Huang and J. Heidemann. Capturing TCP Burstiness for Lightweight Simulation. In *Proceedings of the SCS Multiconference on Distributed Simulation*, Phoenix, Arizona, USA, January 2001. Society for Computer Simulation.

[12] C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Technical Report CSE-TR-433-00, University of Michigan, 2000.

[13] A. Medina, I. Matta, and J. Byers. On the origin of power-laws in Internet topologies. *ACM Computer Communication Review*, April 2000.

[14] V. Misra, W.-B. Gong, and D. Towsley. Fluid-based Analysis of a Network of AQM Routers Supporting TCP Flows with an Application to RED. In *Proceedings of the ACM SIGCOMM*, pages 151–160, Stockholm, Sweden, August 2000.

[15] J. Moy. OSPF Version 2. RFC 2328, Internet Request For Comments, April 1998.

[16] NLANR. NLANR AS Map Archives. http://moat.nlanr.net/Routing/rawdata/, 2000.

[17] S. Raman, S. Shenker, and S. McCanne. Asymptotic Scaling Behavior of Global Recovery in SRM. In *Proceedings of Performance'98/ACM Sigmetrics'98*, June 1998.

[18] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 1771, Internet Request For Comments, March 1995.

[19] G. F. Riley, M. H. Ammar, and R. Fujimoto. Stateless Routing in Network Simulations. In *In Proceedings of the International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, San Francisco, CA, USA, August 2000.

[20] SCAN. SCAN+LUCENT Internet Map. http://http://www.isi.edu/scan/mercator/maps.html/, 1999.

[21] B. Schieber and U. Vishkin. On Finding Lowest Common Ancestors: Simplification and Parallelization. *SIAM J. Comput.*, 17(6):1253–1262, 1988.

[22] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network Topologies, Power Laws, and Hierarchy. Submitted for publication, 2001.

[23] D. Wetherall. Object-oriented Tcl. ftp://ftp.tns.lcs.mit.edu/pub/otcl/doc/tutorial.html.

[24] B. Y. Wu, G. Lancia, V. Bafna, K.-M. Chao, R. Ravi, and C. Y. Tang. A Polynomial Time Approximation Scheme for Minimum Routing Cost Spanning Trees. In *Proc. Ninth Annual ACM-SIAM Symposim on Discrete Algorithms (SODA'98)*, pages 21–32. ACM, 1998.

[25] E. Zegura. Thoughts on Router-level Topology Modeling. The end2end-interest Mailing List Archives, January 2001.

[26] E. Zegura, K. Calvert, and M. Donahoo. A Quantitative Comparison of Graph-based Models for Internet Topology. *ACM/IEEE Transactions on Networking*, 5(6), December 1997.

A Induction of average relative difference

We derive average relative difference of rings by induction, from 3-node to 8-node rings. The deriving process is slightly different for rings with an odd or even number of nodes; thus discussed separately.

Each table in this section shows in three groups that the route lengths and relative differences. The first two groups indicate the route lengths in shortest (Flat) and algorithmic (Algo) routing for all different source and destination pairs. The third group shows the relative differences (Diff) if the values are not zero.

Ring with odd number of nodes Table 4- 6 illustrate results for 3-node, 5-node, and 7-node rings. From these we induce relative difference $= \sum_{i=1}^{k} 2i - 1$, and number of routes $= \sum_{i=1}^{k} (2i - 1 + 2i)$, where $k = \frac{N-1}{2}$.

Given $\sum_{i=1}^{k} i = \frac{k(k+1)}{2}$, we come to the average relative difference:

$$\frac{2\sum_{i=1}^{k} i - k}{4\sum_{i=1}^{k} i - k} = \frac{k}{2k+1} \tag{1}$$

Ring with even number of nodes Similarly, Table 7- 9 illustrate results for 4-node, 6-node, and 8-node rings. From these we induce relative difference $= \sum_{i=1}^{k} 2i$, and number of routes $= 1 + \sum_{i=1}^{k} (2i + 2i + 1)$, where $k = \frac{N-2}{2}$.

	Flat		Algo		Diff	
from → to	2	3	2	3	2	3
1	1	1	1	1		
2		1		2		$\frac{1}{1}$

Table 4: 3-node Ring, $RelativeDiff = 1 * (1/1)$, $NumberRoutes = 1 + 2$

	Flat				Algo				Diff			
from → to	2	3	4	5	2	3	4	5	2	3	4	5
1	1	1	2	2	1	1	2	2				
2		2	1	2		2	1	3				$\frac{1}{2}$
3			2	1			3	1			$\frac{1}{2}$	
4				1				4				$\frac{3}{1}$

Table 5: 5-node Ring, $RelativeDiff = 2 * (1/2) + 1 * (3/1)$, $NumberRoutes = 1 + 2 + 3 + 4$

Given $\sum_{i=1}^{k} i = \frac{k(k+1)}{2}$, we come to the average relative difference:

$$\frac{2\sum_{i=1}^{k} i}{1 + 4\sum_{i=1}^{k} i + k} = \frac{k}{2k+1} \tag{2}$$

Despite the difference in deriving process, we are able to reach one common formula to describe average relative difference for rings of N nodes. Let k be the largest integer smaller than $\frac{N-1}{2}$. The average relative difference is:

$$\frac{k}{2k+1} \tag{3}$$

	Flat						Algo						Diff					
from → to	2	3	4	5	6	7	2	3	4	5	6	7	2	3	4	5	6	7
1	1	1	2	2	3	3	1	1	2	2	3	3						
2		2	1	3	2	3		2	1	3	2	4						$\frac{1}{3}$
3			3	1	3	2			3	1	4	2					$\frac{1}{3}$	
4				3	1	2				4	1	5				$\frac{1}{3}$		$\frac{3}{2}$
5					2	1					5	1					$\frac{3}{2}$	
6						1						6						$\frac{5}{1}$

Table 6: 7-node Ring, $RelativeDiff = 3 * (1/3) + 2 * (3/2) + 1 * (5/1)$, $NumberRoutes = 1 + 2 + 3 + 4 + 5 + 6$

	Flat			Algo			Diff		
from → to	2	3	4	2	3	4	2	3	4
1	1	1	2	1	1	2			
2		2	1		2	1			
3			1			3			$\frac{2}{1}$

Table 7: 4-node Ring, $Relativ_Diff = 1 * (2/1)$, $NumberRoutes = 1 + 2 + 3$

	Flat					Algo					Diff				
from → to	2	3	4	5	6	2	3	4	5	6	2	3	4	5	6
1	1	1	2	2	3	1	1	2	2	3					
2		2	1	3	2		2	1	3	2					
3			3	1	2			3	1	4					$\frac{2}{2}$
4				2	1				4	1				$\frac{2}{2}$	
5					1					5					$\frac{4}{1}$

Table 8: 6-node Ring, $RelativeDiff = 2 * (2/2) + 1 * (4/1)$, $NumberRoutes = 1 + 2 + 3 + 4 + 5$

	Flat							Algo							Diff						
from → to	2	3	4	5	6	7	8	2	3	4	5	6	7	8	2	3	4	5	6	7	8
1	1	1	2	2	3	3	4	1	1	2	2	3	3	4							
2		2	1	3	2	4	3		2	1	3	2	4	3							
3			3	1	4	2	3			3	1	4	2	5							$\frac{2}{3}$
4				4	1	3	2				4	1	5	2						$\frac{2}{3}$	
5					3	1	2					5	1	6					$\frac{2}{3}$		$\frac{4}{2}$
6						2	1						6	1						$\frac{4}{2}$	
7							1							7							$\frac{6}{1}$

Table 9: 8-node Ring, $RelativeDiff = 3 * (2/3) + 2 * (4/2) + 1 * (6/1)$, $NumberRoutes = 1 + 2 + 3 + 4 + 5 + 6 + 7$