

Design and Analysis of Schedules for Virtual Network Migration

Samantha Lo, Mostafa Ammar, and Ellen Zegura
 School of Computer Science, Georgia Institute of Technology
 Email: {samantha,ammar,ewz}@cc.gatech.edu

Abstract—The Internet faces well-known challenges in realizing modifications to the core architecture. To help overcome these limitations, virtual networks run over physical networks and use Internet paths and protocols as essentially a link layer in the virtual network. Effective use of the underlying network requires intelligent placement of virtual networks so that underlying resources do not incur over-subscription. Additionally, because virtual networks may come and go over time, and underlying networks may experience their own dynamic changes, virtual networks may need to be migrated—re-mapped to the physical network during active operation—to maintain good performance. In this paper we consider the problem of scheduling the sequence of node moves that take a virtual network from an original placement to a new placement. We build on prior work that achieves migration of a single node with minimal disruption to develop a model for the migration cost and latency for a given network migration schedule. We then develop algorithms for determining an optimal single-node-at-a-time sequence of moves to minimize migration cost and consideration of multiple node moves in parallel to minimize migration time and cost. Our algorithms are the first we are aware of to systematically address the virtual network migration scheduling problem.

I. INTRODUCTION

The Internet faces well-known challenges in realizing modifications to the core architecture. The use of network virtualization has been proposed to help overcome these limitations (e.g., [6], [18], [5], [7], [15]). Virtual networks run over physical networks and use Internet paths and protocols as essentially a link layer in the virtual network. We focus in this work on the specific form of virtualization where virtual routers are instantiated in physical router hardware and where multiple virtual routers belonging to different virtual networks may share the same physical router. The structure of such networks is illustrated in Figure 1, where two virtual networks are sharing a common physical network. On each physical router, the figure indicates which virtual routers are instantiated; virtual links are denoted by dashed lines that may traverse more than one physical link. This type of virtualization allows a substrate network provider to be decoupled from (and offer services to) overlay network providers.¹

Virtual networks are attractive because they provide significant flexibility in operations and in mapping their requirements to physical network resources. Numerous studies have investigated mapping virtual networks to physical substrate resources

¹This form of virtualization is distinguished from peer-to-peer virtual overlay networks, e.g. [7], where end users establish a network by using physical paths between end systems.

in a manner that makes effective use of the physical network [21], [16], [20], [8]. Virtualization also allows flexibility in changing the mapping of a virtual network over time. Work in this area has explored reasons for network reconfiguration (or migration) such as changes in the traffic carried on the virtual network [9] or because of other virtual networks that share the same physical infrastructure have arrived or departed and the substrate resources should be remapped [21], [16], [20], [10]. Additionally, virtual network migration can form the foundation of a “moving target defense” [4], where a virtual network evades detection or attack by changing its location in the physical network.




Fig. 1. The system architecture for virtual network migration

The work cited above considers questions of *policy* in virtual network migration and determines when to initiate the migration of an existing virtual network as well as the target new placement based on performance or security objectives. Our focus in this work is on questions of *mechanism*, i.e., how to actually move the network. Keller et al. have performed migration of the entire network on the software-defined networking (SDN) platform [14], [13]. However, they have not considered migration scheduling in their work.

A virtual network migration mechanism needs to take into account the specifics of the router virtualization technology. It also needs to be guided by desired performance objectives such as: 1) how to manage disruption to existing traffic on the virtual network, 2) what is the additional overhead incurred by the physical substrate while the migration is in progress and 3) how long the migration process lasts.

Our work starts by leveraging prior work that designs a *single virtual router* live migration mechanism [19]. That work aims to perform virtual router migration in a manner that does not disrupt current data flows using the migrating router. We build on this single virtual router migration mechanism (outlined in more detail in the next section) and use it as a subroutine in the live migration of an *entire* virtual network. Specifically, we are interested in determining, given a new




Fig. 2. Virtual router migration steps 1-5. The old and new physical routers contain multiple virtual routers. (CP=Control Plane, DP=Data Plane)

target location within the substrate for a virtual network, the best *schedule* of virtual router migrations that implement the desired network move. We maintain the goal of not disrupting current network traffic during the migration process and consider how to design a schedule that minimizes the overhead (cost) of network migration or the time it takes to complete the migration task. Our work is the first we are aware of that considers the question of scheduling a live network migration.

We first give an overview of a single virtual router migration and virtual network migration process. Then we present our model in Section III that describes the cost and duration of a network migration sequence. In Section IV, we develop three algorithms for determining an optimal single-node-at-a-time sequence of moves to minimize migration overhead (cost) and schedules that allow for multiple node moves in parallel to minimize migration cost or time. We evaluate these algorithms in Section V with simulations of virtual networks on a physical network. We conclude our work in Section VI with future work.

II. BACKGROUND

In this section, we first review the process of migrating a single virtual router. Then we describe the process of migrating a virtual network and the system architecture that supports this functionality.

A. Live Migration of a Single Virtual Router

Our work leverages the research in [19], which develops a mechanism for the live migration of a single router without disrupting current traffic traversing the virtual network. That work assumes a router architecture with multiple virtual routers instantiated on a physical router. Three features of virtual router architecture make it “migratable”: independence of the virtual router through OS virtualization support, separation of data and control planes, and the ability to dynamically bind a router’s data plane to physical substrate interfaces. With this router architecture, the work in [19] describes the following sequence of actions that are required to move a virtual router. The numbers in Figure 2 correspond to each step:

- 1) *Setup a tunnel* between the original and the final physical locations of the virtual router.
- 2) *Migrate control plane and copy memory*: Create an image of the virtual router’s control plane and transfer it through the tunnel to the new physical router. At the same time, collect the routing update messages for this virtual router and forward them to the virtual router after the image transfer is completed such that they can be processed in the new location.

- 3) *Clone the data plane* based on the existing control plane on the new physical node and repopulate the FIB. After this is completed, both old and new data planes are running at the same time.
- 4) *Duplicate forwarding links between the virtual router and its neighbors*: Set up the outgoing virtual links from the new physical router to its neighbors. Once these are ready, set up original incoming virtual links to redirect traffic from the neighbors of the original virtual router to the virtual router in the new physical machine. The new traffic flows are assigned to the new forwarding links.
- 5) *Remove old forwarding links*: Once old traffic flows at the old forwarding links are complete, remove the data plane instance in the original physical location.

We use this single virtual router migration process as a subroutine to migrate each virtual router on a virtual network, as described next. In the rest of the paper, we use the term “virtual node” and “virtual router” interchangeably.

B. Network Migration Process

Figure 3 demonstrates a virtual network migration. The example network has three virtual nodes, A, B and C, which are moved one at a time, in three stages, across the substrate. Each stage has two parts. During the *preparation part*, the control plane moves to its new physical node and the data plane is cloned (Step 1-3 in Section II-A). After the preparation part, during the *migration part*, the duplicate forwarding virtual links are created and eventually the old forwarding links are removed (Steps 4-5 in Section II-A).

The migration sequence in the figure is (A, C, B). After A’s control plane is migrated in the preparation part of Stage 0 in Figure 3(b), virtual links between the physical nodes of A and its neighbors, B and C are set up and are working in parallel with the original virtual links in the migration part in Figure 3(c). After the migration of node A, the virtual links between the migrated node A and its neighbors are not finalized until all the neighbors have been migrated. We will return to this issue when we analyze the cost of a move, but for now we simply note that there are migration overhead links that must be maintained in support of a migration in progress. The length (physical network hop count) and the duration (time maintained) of any overhead link is influenced by the migration sequence. Our migration scheduling algorithms aim to reduce length and duration of overhead links.

C. System Architecture for Virtual Network Migration

To enable the network to migrate and determine the migration sequence, we adopt a simple system architecture. Figure 1 shows our system architecture which is similar to some existing overlay network architectures such as GENI [11]. In our architecture, the physical network has a controller to initiate and synchronize the virtual network migrations. The controller connects to all of the physical routers through in-band or out-of-band tunnels. When a request for a virtual network migration is initiated based on reconfiguration goals, the controller is given initial and final virtual network mappings.

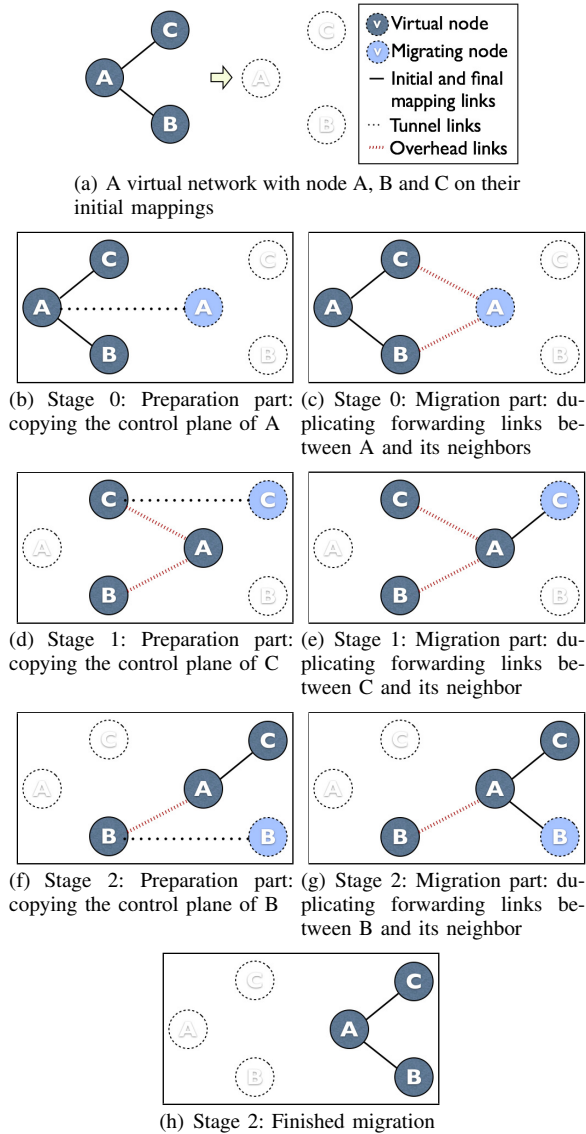


Fig. 3. Migrating a virtual network with virtual node A, B and C. (Physical network is not shown.) The migration sequence is A, C and then B.

The controller decides the node migration sequence using algorithms running locally or remotely, and that are discussed in Section IV. The controller initiates the steps for each single node migration to accomplish the network migration.

III. PROBLEM STATEMENT

We now turn to a formal statement of the migration problem and a model for the cost and time associated with a given virtual node move sequence.

A. The Migration Problem

Let $P = (R, E)$ be the physical network where R is a set of physical routers and E is a set of physical links. Let $G = (V, L)$ be a virtual network where V is a set of virtual nodes and L is a set of virtual links.

Let $m : V \rightarrow R$ be a function that maps the set of the virtual nodes in G to a set of physical routers in P , i.e. $m(v) = r$ means a virtual node $v \in V$ is mapped to a physical router $r \in R$. A virtual link is mapped to a physical path (an ordered

list of connected physical routers) where the endpoints of the virtual link are mapped to the first and last routers on the path and each pair of physical routers in the list are joined by a physical link.

We are given an initial mapping $m_0 : V \rightarrow R$ and a final target mapping $m_f : V \rightarrow R$. The goal of migration is to determine the most efficient schedule to move virtual nodes so as to accomplish the virtual network move from m_0 to m_f , where efficiency may be measured by overhead cost of migration, total migration time, or a combination of the two.

As explained earlier, network migration consists of a sequence of stages, where each stage migrates a subset of the virtual nodes. Let $S = (S_0, S_1, S_2, \dots)$ be a sequence of node migrations that move the virtual network from m_0 to m_f . At stage i , the set of virtual nodes moved is $S_i \subseteq V$ such that $\forall i \neq j S_i \cap S_j = \emptyset$ and $\forall i S_i \neq \emptyset$ and $\bigcup_i S_i = V$.

There is an additional constraint on the sets S_i , namely that a given stage, we cannot move two nodes that are neighbors in the virtual network. This constraint arises because of the way that links are maintained in the single node move procedure. If two neighbors are moved at the same time, the virtual network will enter an inconsistent state. When we devise algorithms to move multiple nodes at a time, we are careful to satisfy the no-neighbors constraint.

We are interested in two important performance metrics: 1) the completion time of a virtual network migration, and 2) the cost of a network migration representing the bandwidth impact on the physical network of the migration process. We describe each of these formally in the next two subsections.

B. Network Migration Time

Recall that each stage consists of a preparation part and a migration part, as illustrated in Figure 4. The preparation part consists of setting up the tunnels, migrating the control planes and cloning the data planes by repopulating the FIB. The migration part consists of duplicating forwarding links, waiting for all old data flows to complete, and finally removing old forwarding links. After stage i , the network has fully completed all moves in sequence S up to and including S_i .




Fig. 4. Illustration of the time of a virtual network migration

We denote the time taken in the *preparation part* and the *migration part* at stage i by $t^p(i)$ and $t^m(i)$, respectively. The total time for a particular migration sequence S is then

$$\sum_{i=0}^{|S|-1} (t^p(i) + t^m(i)). \quad (1)$$

The time spent in the preparation part $t^p(i)$ depends on the following characteristics of the migrating nodes in S_i : the available bandwidth of the tunnels, the size of the control plane images, and the time to repopulate each FIB. For simplicity of exposition, we assume for a given virtual network that

the control plane image sizes are the same for all the nodes. When the sizes of the control planes are the same and the data plane technology of the physical routers are the same, the FIB repopulation times for all the nodes are also the same. An interesting effect takes place, however, when nodes migrate together at the same stage. As more nodes move together, there is a higher probability that the tunnels used to move the control planes may share a bottleneck physical link. Thus while it is generally advantageous to move multiple nodes at a time, too much interference in the control plane tunnels can reduce or eliminate the benefit.

We model the time in the preparation part $t^p(i)$ as a function of the nodes in the set S_i and the time to migrate a control plane and repopulate a FIB, $f(S_i, \alpha)$, where α is the time for migrating the control plane and populating the FIB of a single node. We explore different functions $f(S_i, \alpha)$ in Section V.

The time in the migration part $t^m(i)$ mainly depends on the current traffic flows on the migrating nodes' old forwarding links. Once these flows end, the migrating nodes can completely switch to the new forwarding links. We assume that the time for these flows to end is constant over all stages and independent of the number of nodes moved at a stage.

Given the time computation, it should be clear that the primary method for reducing the migration time is to move as many nodes as possible together, thus reducing the number of stages, while taking care not to overload physical links that are needed to move the control planes.

C. Network Migration Cost

We model the impact of the virtual network migration on the physical network as a cost function that is divided into the cost during the preparation part ($c^p(i)$) and the cost during the migration part ($c^m(i)$) of each migration stage i . Our measure of cost is the number of physical links that carry migration-related traffic multiplied by the duration these links are used (stage).

In the preparation part, there is a bandwidth impact on the physical network resulting from the transfer of the control plane across the tunnel established from the original location to the final location. Regardless of the migration sequence, each control plane must be transferred, and the cost is simply the sum of all control plane sizes. Hence this cost component is a constant. The bandwidth impact can be different under different migration sequences when tunnels possibly share some physical links with different sets of overhead virtual links. We will address this issue in two of our proposed algorithms.

We are interested in the cost implications that derive from overhead virtual links. In both the preparation and the migration part, a set of virtual links are present only because of the migration process. These links are dependent on the migration sequence. To understand the impact of these overhead virtual links, we divide the nodes into three sets at stage i :

- 1) Nodes that will move in this stage (i.e., $S_i \in V$),
- 2) Nodes that have already moved in prior stages, denoted $A_i \in V$, and




Fig. 5. Illustration of types of links at stage i when v is migrating after w has moved at stage $i - 1$.

- 3) Nodes that have not yet moved, denoted $Y_i \in V$.

where $S_i \cup A_i \cup Y_i = V$ and $S_i \cap A_i \cap Y_i = \emptyset$.

We classify overhead virtual links at stage i into three sets:

- 1) $L_1(i)$: Virtual links between the original mappings of the moving nodes (S_i) and final mappings of their already moved neighbor nodes (A_i), i.e. $l = (v, w)$ which the physical path is $(m_0(v), m_f(w))$ where $v \in S_i, w \in A_i$,
- 2) $L_2(i)$: Virtual links between the final mappings of the moving nodes (S_i) and the original mappings of their not moved neighbor nodes (Y_i), i.e. $l = (v, x)$ which the physical path is $(m_f(v), m_0(x))$ where $v \in S_i, x \in Y_i$, and
- 3) $L_3(i)$: Virtual links between the final mappings of the already moved nodes (A_i) and the original mappings of their not moved neighbor nodes (Y_i), i.e. $l = (w, x)$ which the physical path is $(m_f(w), m_0(x))$ where $w \in A_i, x \in Y_i$.

Figure 5 illustrates a simple three-node virtual network with nodes x, v , and w . In the stage illustrated, node v is moving, node w has already moved, and node x has yet to move. The virtual link (v, w) which is mapped to the physical path $(m_0(v), m_f(w))$ is of type L_1 , the virtual link (x, v) which is mapped to the physical path $(m_f(v), m_0(x))$ is of type L_2 , and the virtual link (w, x) which is mapped to the physical path $(m_f(w), m_0(x))$ is of type L_3 .

The links in $L_1(i)$ and $L_3(i)$ are present throughout both the preparation and migration parts of stage i . The links in $L_2(i)$, on the other hand, are present only during the migration part of stage i . Indeed they are created to maintain connectivity when the nodes in S_i are moving. We use $p(l)$ to denote the bandwidth impact of virtual link l . Later we simplify by assuming $p(l)$ is the number of physical links in path l .

The cost of the virtual network in the preparation part of stage i is then

$$c^p(i) = t^p(i) \sum_{l \in (L_1(i) \cup L_3(i))} p(l). \quad (2)$$

The cost of the virtual network in the migration part of stage i is

$$c^m(i) = t^m(i) \sum_{l \in (L_2(i) \cup L_3(i))} p(l). \quad (3)$$

The total cost for a particular migration sequence S is the sum of the cost in all preparation and migration stages:

$$\sum_{i=0}^{|S|-1} c^p(i) + c^m(i). \quad (4)$$

IV. ALGORITHMS FOR VIRTUAL NETWORK MIGRATION

We present three algorithms for migrating virtual networks. These algorithms can be divided into two types: moving one node at a time and moving multiple nodes at a time. We design the algorithms with two goals. One of the goals is to minimize the total migration time of the migration, which is described in Equation 1. The other goal is to minimize the total cost, which is described in Equation 4.

In the move-one-node-at-a-time scheme, a virtual node can start moving to its final placement only when the previous virtual node has finished its move. We find a move sequence $S^* = (S_0, S_1, \dots)$ where $\forall i |S_i| = 1$.

In the move-multiple-nodes-at-a-time scheme, multiple nodes can be migrated at the same stage. However, we are constrained not to move nodes that are neighbors in the virtual network at the same stage.

In this section, we first introduce the algorithm Local Minimum Cost First (LMCF) which produces a move-one-node-at-a-time schedule to minimize the migration cost. Then we present two move-multiple-nodes-at-a-time algorithms: Maximal Independent Set-Size Sequence (MIS-SS) and Maximal Independent Set-Local Minimum Cost First (MIS-LMCF) which minimize the time and the cost of the migration. Note that we do not have an algorithm in the category of move-one-node-at-a-time with the goal of minimizing the time of migration because the number of stages required for any one-node-at-a-time scheme is the same.

A. Local Minimum Cost First (LMCF)

The goal of Local Minimum Cost First (LMCF) algorithm is to minimize the migration cost under the move-one-node-at-a-time scheme. LMCF is described in Algorithm 1. First we compute the initial migration cost of the migration part $c^m(0)$ of stage 0 for each virtual node and choose the cheapest one to migrate first. Then from the rest of the unmoved virtual nodes, we recompute the migration cost c^m for each of them based on the current link mapping at that stage. Again, we choose the one that gives the cheapest cost and migrate it. We repeat until the network is completely migrated. If there are nodes that give the same cost, we pick the node with the lowest cost of the preparation part c^p . If both of their migration and preparation parts' costs are the same, we break the tie by randomly picking one of the cheapest nodes to be moved next.

B. Maximal Independent Set-Size Sequence (MIS-SS)

The goal of MIS-SS is to minimize the migration time. In MIS-SS, we simplify the time as the number of stages involved in a virtual network migration. We consider minimizing the number of migration stages by migrating multiple nodes at a time. We first determine the nodes which can be migrated at the same stage. We call the set of nodes to be migrated at a stage as *migration set*. Because of the no-neighbor constraint, the nodes to be migrated at the same stage cannot be neighbor among each other. The set of these non-neighbor nodes is called *independent set*. We find the maximum sets among these independent sets such that each set is not a subset of another

Algorithm 1 Local Minimum Cost First

```

1: function LMCF( $G = (V, P), P = (R, E), m_0, m_f$ )
2:   set sequence  $S = []$  ▷ Result sequence
3:   for  $i = 0 \rightarrow |V| - 1$  do
4:      $j = 0$ 
5:     for  $v \in V - (S_0 \dots S_{i-1})$  do
6:        $S_i = v$ 
7:        $cost[v] = c^m(i)$  ▷ Find the cost for  $v$  to be migrated at
      stage  $i$ 
8:       if  $j = 0$  then
9:          $minv = v$ 
10:         $mincost = cost[v]$ 
11:       else if  $cost[v] < mincost$  then
12:          $minv = v$ 
13:          $mincost = cost[v]$ 
14:       end if
15:        $j = j + 1$ 
16:     end for
17:      $S_i = minv$  ▷ Put the minimum cost node to move at stage  $i$ .
18:   end for
19:   return  $S$ 
20: end function

```

set. These sets are called *maximal independent sets* (MISs). We use an algorithm from [2] to generate MISs. This MIS algorithm generates a maximal independent set with inputs of a graph and a node in the graph. We use this algorithm and input each node to the algorithm to generate the maximal independent sets of the graph.

With the migration sets produced by the MIS algorithm, we design an algorithm to determine the migration sequence of these sets to minimize the time. MISs are not *maximum* independent sets. MISs can have intersections. Because of their intersections, when we migrate a MIS, some of the nodes in other MISs may be migrated as well. Thus, the remaining sets shrink. The number of stages involved in a virtual network migration can be highly reduced from LMCF, to the number of MISs, then to the number of migration sets in a set sequence which eliminates more nodes at the earlier stages.

In MIS-SS (Algorithm 2), first we use the MIS algorithm (FindMIS($G=(V,P)$)) to determine migration sets of a virtual network. Then, MIS-SS determines the migration sequence by picking the largest MIS to be migrated first. At each stage, MIS-SS picks the largest remaining set to be migrated. After each stage, MIS-SS removes the nodes which have been migrated from the remaining sets. The remaining sets shrink when more nodes are migrated. If the sets are highly intersected, the time for migration can be reduced significantly.

C. Maximal Independent Set-Local Minimum Cost First (MIS-LMCF)

In MIS-LMCF, we attempt to minimize a combination of cost and time. When we shorten the time of the migration, we can minimize the time of the overhead links staying at the migrating state and further minimize the cost of the migration. In Algorithm 3, we move multiple nodes at a time to minimize the migration time. We also consider the migration cost of each migration set to schedule the migration sequence.

We start with the sets of nodes $MSet[]$ from the MIS algorithm. We use a modified version of LMCF to determine

Algorithm 2 MIS-Size Sequence (MIS-SS)

```

1: function MIS-SS( $G = (V, P)$ ,  $P = (R, E)$ ,  $m_0$ ,  $m_f$ )
2:    $MSets[] = FindMIS(G = (V, P))$  ▷ Call the MIS algorithm to
   generate all maximal independent sets (MISs)
3:    $S_0 = MaxSize(MSets[])$  ▷ First migrate the largest MIS
4:   for  $i = 1 \rightarrow |MSets[] - 1$  do
5:     Remove  $\forall v \in S_0 \dots S_{i-1}$  from all  $MSets[]$  ▷ Remove the
   migrated nodes from  $MSets[]$ 
6:     if  $MaxSize(MSets[]) = \emptyset$  then
7:       break
8:     end if
9:      $S_i = MaxSize(MSets[])$  ▷ Migrate the set
   with the most number of nodes left. If there are equal number of nodes
   in more than one sets, pick one set randomly.
10:  end for
11:  return  $S$ 
12: end function

```

the set migration sequence. We migrate the set with the lowest cost first. After each stage, we remove the migrated nodes from the remaining sets and recompute the migration cost for each remaining set. Again, we choose the set with the lowest cost to be moved next.

Algorithm 3 MIS-Local Minimum Cost First (MIS-LMCF)

```

1: function MIS-LMCF( $G = (V, P)$ ,  $P = (R, E)$ ,  $m_0$ ,  $m_f$ )
2:    $MSets[] = FindMIS(G = (V, P))$  ▷ Call the MIS algorithm to
   generate all MISs
3:   for  $i = 0 \rightarrow |MSets[] - 1$  do
4:      $j = 0$ 
5:     Remove  $\forall v \in S_0 \dots S_{i-1}$  from all  $MSets[]$  ▷ Remove the
   migrated nodes from  $MSets[]$ 
6:     if  $MaxSize(MSets[]) = \emptyset$  then
7:       break
8:     end if
9:     for  $set \in MSets[] - \{S_0, \dots, S_{i-1}\}$  do
10:       $S_i = set$ 
11:       $cost[set] = c^m(i)$  when  $S_i = set$  ▷ Find the cost for the
   set of nodes to be migrated at stage  $i$ 
12:      if  $j = 0$  then
13:         $minset = set$ 
14:         $mincost = cost[set]$ 
15:      else if  $cost[set] < mincost$  then
16:         $minset = set$ 
17:         $mincost = cost[set]$ 
18:      end if
19:    end for
20:     $S_i = minset$  ▷ Return the set of nodes with minimum cost to
   move at stage  $i$ 
21:  end for
22:  return  $S$ 
23: end function

```

V. EVALUATION

We evaluate the migration algorithms with simulations. We are interested in a set of basic questions about algorithm performance relative to optimal, to worst case, and to one another. Specifically, (Q1) how important is it to have an intelligent migration algorithm? (Q2) how close do the proposed algorithms come to optimal? (Q3) how do the algorithms compare to one another? We next turn to two characteristics of the virtual and physical networks that play a role in algorithm performance. (Q4) What is the influence on cost of migration distance, i.e., how far apart in the physical network are the

original and final virtual network (VN) locations? (Q5) What is the influence of VN node degree on the time and cost of migration? Based on the results, we make some preliminary recommendations about the operation of networks that support migration.

A. Simulation Setup

We evaluate the algorithms with 50 different 10-node VNs and a 100-node physical network. The 10-node VNs are randomly generated by the Python NetworkX package [3] with average node degree distributions from 2.4-5.4 and diameter from 2-7. The 100-node physical network is generated by GT-ITM [12] with the transit-stub model with 3.7 average node degree and diameter of 11.

In our simulations, each VN is mapped to different initial and final mappings on the physical network. These mappings are randomly chosen with the constraint that there are no physical nodes in common in the initial and final mappings. We simulate the migration sequences produced by the algorithms. We also exhaustively try every one-node-at-a-time move sequence. We calculate the cost and time for each migration sequence generated by the exhaustive search and the three algorithms. We use the control plane transfer time from [19] for the preparation part ($t^p = f(S_i, \alpha)$). We assume that each of the virtual routers on the 10-node VNs has less than 1k routes. The memory copy time is around 1 second with 10k routes. The FIB repopulating time of the software data plane is 2.1 seconds for 1k routes. Thus, $\alpha = 1 + 2.1 = 3.1$ seconds. We experiment with two choices for the function $f(S_i, \alpha)$. In the first case, we assume we will migrate one control plane at a time, hence $f_1(S_i, \alpha) = \alpha |S_i|$. In the second, we assume that we can migrate all control planes in parallel with no bandwidth interference, hence $f_2(S_i, \alpha) = \alpha$.

For the time in the migration part, we model the longest finishing time of any current flow traversing the migrating router(s). This is not easy to estimate, since in general it depends on the uses of the VN, the capacities of the physical and virtual links, and the degree of sharing among flows and VNs. Further, long running flows might lead to policy decisions to avoid migrating them altogether or to disrupt these flows so that migration can be accomplished more quickly. In the absence of specific information, and consistent with our desire to use values derived from [19], we use the longest control plane transfer time as our longest flow finishing time. Figure 8 in [19] shows a time of about 3 seconds to transfer a 124 MB control plane. We choose to use $t^m = 3.5$ seconds, providing a bit of cushion above the 3 second measurement. In future work, we will examine models for flow finishing time and adaptations to the migration process with long flows.

B. Baseline Results

We first consider basic algorithm performance (Q1). Table I shows the optimal and worst case costs of all possible one-node-at-a-time migration sequences for VN 1-5, as well as the ratio of worst cost to optimal cost. It is clear that the ratio of worst cost to optimal cost can be substantial,

exceeding 2.5 for all five networks and as high as 3.37 for VN 1, which has lowest average node degree and highest diameter. We conclude from this experiment that there is value to investigating migration algorithms that can achieve cost that drives towards optimal.

TABLE I

THE COSTS OF THE BEST CASES, THE WORST CASES AND LMCF OF VN 1, 2, 3, 4 AND 5 (MOVE ONE-NODE-AT-A-TIME SEQUENCES) WITH $t^p = f_1(S_i, \alpha = 3.1)$ SEC AND $t^m(i) = 3.5$ SEC

VN	Optimal S^*	Worst	Ratio	LMCF
1	890.42	3000.64	3.37	890.42
2	1116.06	3263.55	2.92	1179.68
3	1413.76	3585.74	2.53	1635.49
4	1739.87	4864.23	2.80	1897.57
5	2751.34	6984.34	2.54	2985.42

We compare the results from our greedy LMCF algorithm with the best cost. The rightmost column of Table I shows that LMCF does very well, finding migration sequences with cost close to the optimal solution on all five VNs. Without an exhaustive search on all $10! \approx 3.6\text{M}$ combinations of migration sequences, LMCF can give a migration sequence with cost between 1 and 1.16 times the optimal migration sequence.

We next turn to the performance of the two algorithms that can move multiple nodes at a time. Recall that MIS-SS attempts to minimize time by moving as many nodes at a time in each stage and hence reduce the total number of stages. MIS-LMCF attempts to minimize a combination of cost and time by moving large sets but choosing which set to move based on cost considerations. Both algorithms rely on the MIS algorithm to generate candidate migration sets. It is computationally prohibitive to exhaustively try all possible multiple-node-at-a-time sequences, hence we focus our attention on the comparison of these algorithms to one another and then to the LMCF algorithm. Table II shows the results assuming one control plane is migrated at a time, while Table III shows the results assuming all control planes in a stage can migrate together without interference.

The goal of the MIS-SS algorithm is to minimize the

TABLE II

RESULTS OF MIS-SS, MIS-LMCF AND LMCF WITH FIVE OF THE 10-NODE VNS WITH $t^p = f_1(S_i, \alpha = 3.1) = |S_i|\alpha$ (THE BEST COSTS ARE UNDERLINED)

VN	MIS-SS		MIS-LMCF		LMCF	
	Cost	Time	Cost	Time	Cost	Time
1	1087.42	44.88	1093.18	44.88	<u>890.42</u>	65.88
2	1415.38	41.38	<u>1162.36</u>	48.38	1179.68	65.88
3	1714.10	44.88	<u>1402.41</u>	51.88	1635.49	65.88
4	1603.54	44.88	<u>1482.29</u>	44.88	1897.57	65.88
5	<u>1710.82</u>	44.88	3335.73	44.88	2985.42	65.88

TABLE III

RESULTS OF MIS-SS, MIS-LMCF AND LMCF WITH FIVE OF THE 10-NODE VNS WITH $t^p = f_2(S_i, \alpha = 3.1) = \alpha$ (THE BEST COSTS ARE UNDERLINED)

VN	MIS-SS		MIS-LMCF	
	Cost	Time	Cost	Time
1	797.15	29.44	<u>657.78</u>	29.44
2	881.15	25.94	<u>872.09</u>	32.94
3	<u>1102.67</u>	32.53	1105.96	39.53
4	1090.94	32.53	<u>1056.14</u>	32.53
5	1710.82	26.35	1501.46	26.35

migration time. It successfully achieves the lowest migration time of any of the three algorithms on each VN, and it improves on the LMCF algorithm by about 1.5 times when the control planes move one at a time and by more than 2 times when the control planes move in parallel. If migration time is the key consideration, the MIS-SS algorithm provides a good solution.

The goal for MIS-LMCF is also to reduce the number of stages but to do so taking cost into consideration. The MIS-LMCF algorithm successfully achieves relatively low migration time as compared to the LMCF algorithm, and frequently matches the time performance of MIS-SS. Turning to cost, the MIS-LMCF algorithm achieves the lowest cost more often than MIS-SS or LMCF, but it does not always have lowest cost. For example, in Table II, for VN 1, the lowest cost is achieved by the LMCF algorithm rather than either of the MIS algorithms, while for VN 5 the lowest cost is achieved by MIS-SS. If migration cost is the key consideration, it is important to understand whether control planes can move in parallel with limited interference. If they can, one of the MIS-based algorithms will improve cost. If they cannot, the MIS algorithms generally offer only modest improvement over LMCF, though there is a notable exception for VN 5 where LMCF cost is twice the cost of the MIS-based algorithms.

C. Effect of Migration Distance

We now dig deeper into the effect of virtual and physical network characteristics on the performance of the algorithms. In particular, we examine the effect of migration distance and then the effect of VN node degree (Q4). Migration distance is determined by the initial and final mappings of a VN to the physical network. We use VN 2 and simulate 90 migrations with 10 different mappings on the 100-node network. For each pair of mappings m_0 and m_f , we calculate the average migration distance by averaging the physical path lengths of the virtual node locations between $m_0(v)$ and $m_f(v)$. For our VN locations and physical network characteristics, we see average migration distances in the range of roughly 3.5 to 6.5 physical node hops. Recall that the physical network has diameter 11, so these distances are consistent with that constraint on the maximum distance a node can move.

Our model for time is based on the time to move one control plane, the interference between control plane moves, and the time to complete the longest outstanding flow. None of these depend on migration distance in our current model, though a refinement of the model might take into consideration that longer distance moves may take more time and that distance might influence how much interference exists between control plane moves.

We expect cost to be influenced by migration distance, since overhead links must traverse the physical network between original node locations to final node locations. In Figure 6, we show the costs of the migration sequences from LMCF, MIS-SS and MIS-LMCF as a function of the average migration distance. The LMCF plot contains 90 points, one for each initial-final pair, while the MIS-SS and MIS-LMCF plots




Fig. 6. Cost for migrating VN 2 among 10 different initial and final mappings (90 migrations) with algorithm LCMF, MIS-SS and MIS-LMCF

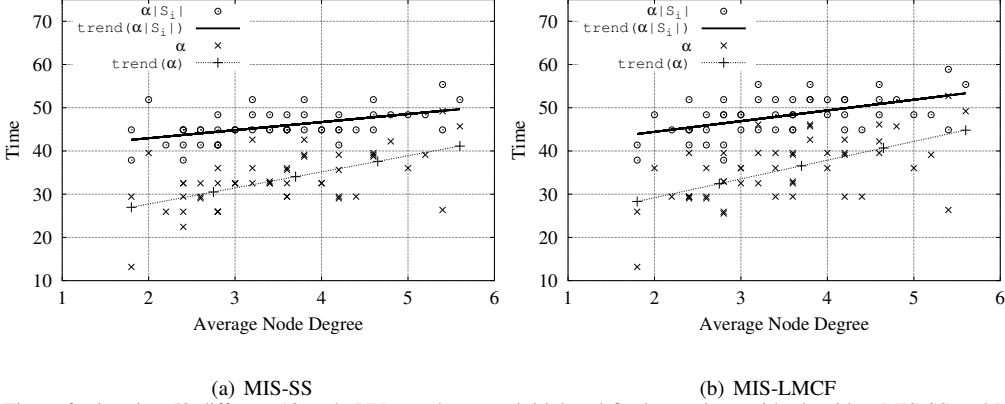


Fig. 7. Time of migrating 50 different 10-node VNs on the same initial and final mappings with algorithm MIS-SS and MIS-LMCF

contain two sets of 90 points, one for each model of control plane interference ($f_1(S_i, \alpha) = \alpha|S_i|$ and $f_2(S_i, \alpha) = \alpha$). In all three cases, the cost increases with average migration distance, as expected. The MIS-SS and MIS-LMCF algorithms tend to achieve comparable cost that improves noticeably over LCMF when the control planes can move in parallel. The impact of higher average migration distance is somewhat modest—the highest migration distance settings have cost approximately 1.4x the lowest migration distances. The cost advantage of MIS-LMCF is more pronounced when control plane moves can interfere than when they cannot, and this persists across migration distances.

D. Effect of Node Degree

Finally we turn to Q5, the question of the impact of VN node degree on the time and cost performance of migration. In our initial experiments, we observed variability in cost across the five VNs, with cost growing as VN node degree increased. Higher node degree means more neighbors, with two cost effects. First, more dense networks have more smaller size MISs, hence more stages and more time is required. Second, more dense networks require more overhead links to connect nodes that have moved to neighbors that have not. Both effects will increase cost. In time, we saw very little variability as average node degree increased.

To explore more deeply, we use all 50 different 10-node VNs. We migrate each from and to the same initial and final mappings so as to eliminate the effect of migration distance.

We compare the time for the two MIS-based algorithms and the cost for all three algorithms. Figure 7 shows the migration time of MIS-SS and MIS-LMCF. With increasing average node degree, fewer nodes can be grouped into the same migration set. Thus, the time for migration increases with the average node degree, albeit slowly. Overall, MIS-SS achieves its goal to minimize the migration time for networks with different average node degrees. Because the time is dominated by the migration sets, both algorithms give similar migration times, but the time grows more slowly when the control planes interfere than when they do not.

Figure 8 shows the costs of sequences from LCMF, MIS-SS, and MIS-LMCF. As expected, an increase in average node degree has a substantial effect on the cost. The LCMF algorithm is most significantly affected by the increase in average node degree, since there are more overhead links and they must remain in place until the affected nodes have been moved. The MIS-based algorithms also see increases in cost with average node degree, with slightly slower cost growth when control planes do not interfere. In the limit, when the VN is a full mesh, the MIS-based algorithms must move single nodes at a time and hence cannot do any better than LCMF.

E. Summary

From our simulations, all algorithms perform reasonably well. They all give migration sequences with reasonable costs and time. We also observe that no one algorithm consistently




Fig. 8. Cost of migrating 50 different 10-node VNs on the same initial and final mappings with algorithm LCMF, MIS-SS and MIS-LMCF

exhibits the lowest cost across all the average distance and node degree cases considered in our experiments.

VI. CONCLUSIONS

We believe our work makes the application of re-mapping policies in virtual networks possible, especially in architectures where virtual networks require reconfiguration. Our work considers the problem of the mechanism to migrate an entire virtual network from its initial mapping to its final mapping. We design scheduling algorithms for virtual network migrations that minimize the disruption to the current data traffic and the overhead traffic in the migration process. We propose three algorithms, LCMF, MIS-SS, and MIS-LMCF to discover the virtual router migration sequences for migrating an entire virtual network with lower overhead (cost) and shorter time. The LCMF algorithm produces a move-one-node-at-a-time schedule that exhibits cost close to the optimal solution. We then consider move-multiple-nodes-at-a-time schedules to minimize the time and the cost of migration. Our simulations of the MIS algorithms show that move-multiple-nodes-at-a-time schedules effectively minimize the time and cost.

The interactions among multiple virtual network migrations have not been studied in this paper. In future work, the scheduling problem of migrating multiple virtual networks that share the same substrate becomes important. The scheduling problem of multiple virtual network migrations involves the dynamic of migrating multiple virtual routers belonging to different virtual networks. When multiple virtual networks are migrating at the same time, some of the physical links and physical routers are fully occupied and may exceed the capacities of the links and physical routers. Thus, scheduling multiple virtual network migrations is more complicated than a single virtual network migration.

Our work further raises the questions of how we design the applications and system architectures for virtual networks. A possible application of virtual network migration is defending against attack traffic on virtual networks. Future work on extending our virtual network migration scheduling to other technologies such as architecture with FlowVisor [1] and OpenFlow enabled switch networks [17] will be very useful as well.

ACKNOWLEDGMENTS

This work was supported by the NSF Grant 1017237.

REFERENCES

- [1] FlowVisor. <http://www.openflowswitch.org/wk/index.php/FlowVisor>.
- [2] Maximal independent set algorithm implementation. <https://networkx.lanl.gov/trac/ticket/282>. URL retrieved Apr 2012.
- [3] Networkx: A python package. <http://networkx.lanl.gov/>.
- [4] E. Al-Shaer. Toward network configuration randomization for moving target defense. In *Moving Target Defense*, volume 54 of *Advances in Information Security*, pages 153–159. Springer New York, 2011.
- [5] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *Proc. 18th ACM SOSP*, pages 131–145, Canada, Oct. 2001.
- [6] S. Banerjee, B. Bhattacharjee, and C. Kommareddy. Scalable application layer multicast. In *Proc. ACM SIGCOMM*, Pittsburgh, PA, 2002.
- [7] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. Wallach. Secure routing for structured peer-to-peer overlay networks. In *Proc. 5th USENIX OSDI*, Boston, MA, Dec. 2002.
- [8] M. Demirci and M. Ammar. Fair allocation of substrate resources among multiple overlay networks. In *Proc. of MASCOTS*, volume 0, pages 121–130, Los Alamitos, CA, USA, 2010. IEEE Computer Society.
- [9] J. Fan and M. H. Ammar. Dynamic topology configuration in service overlay networks: A study of reconfiguration policies. In *Proc. IEEE INFOCOM*, 2006.
- [10] N. Farooq Butt, M. Chowdhury, and R. Boutaba. Topology-awareness and reoptimization mechanism for virtual network embedding. In *Proc. IFIP, NETWORKING'10*, 2010.
- [11] GENI: Global Environment for Network Innovations. <http://www.geni.net/>.
- [12] GT-ITM. <http://www.cc.gatech.edu/projects/gtitm/>.
- [13] E. Keller, D. Arora, D. P. Botero, and J. Rexford. Live migration of an entire network (and its hosts). *Princeton University Computer Science Technical Report*, TR-926-12, June 2012.
- [14] E. Keller, S. Ghorbani, M. Caesar, and J. Rexford. Live migration of an entire network (and its hosts). In *Proc. ACM HotNets-XI*.
- [15] A. D. Keromytis, V. Misra, and D. Rubenstein. SOS: Secure overlay services. In *Proc. ACM SIGCOMM*, pages 61–72, Pittsburgh, PA, 2002.
- [16] J. Lu and J. Turner. Efficient mapping of virtual networks onto a shared substrate. *Tech Report WUCSE2006*, 35(2006-35):1–11, 2006.
- [17] OpenFlow Switch Consortium. <http://www.openflowswitch.org/>.
- [18] L. Subramanian, I. Stoica, H. Balakrishnan, and R. H. Katz. Overqos: offering internet qos using overlays. *SIGCOMM Comput. Commun. Rev.*, 33:11–16, January 2003.
- [19] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford. Virtual routers on the move: Live router migration as a network-management primitive. In *Proc. ACM SIGCOMM*, Seattle, WA, 2008.
- [20] M. Yu, Y. Yi, J. Rexford, and M. Chiang. Rethinking virtual network embedding: Substrate support for path splitting and migration. *ACM SIGCOMM Computer Communications Review*, Apr. 2008.
- [21] Y. Zhu and M. Ammar. Algorithms for Assigning Substrate Network Resources to Virtual Network Components. In *Proc. IEEE INFOCOM*, Barcelona, Spain, Mar. 2006.