

# A Verification Based Flow Space Management Scheme for Multi-Tenant Virtualized Network

Shun Higuchi

Graduate School of Computer and Information Science  
Hosei University  
Tokyo, Japan  
Email: shun.higuchi.6j@stu.hosei.ac.jp

Toshio Hirotsu

Faculty of Computer and Information Science  
Hosei University  
Tokyo, Japan  
Email: hirotsu@hosei.ac.jp

**Abstract**—Cloud services that virtualize existing IT infrastructures at data centers are widely used by governments, universities, and companies. Multi-tenancy is required for data centers to provide a large number of isolated networks to each organization. OpenFlow is a core technology of software defined networking (SDN) and is useful for centrally managing and controlling these networks; however, SDN is used only at the management level. It is desirable to make the flexible features of SDN/OpenFlow available to users' virtual networks. FlowVisor [3] virtualizes multi-tenant OpenFlow networks by coordinating multiple controllers, but it is unable to deal with conflicts of control rules among individual virtual networks. Administrators of each tenant thus need to design the control rules of their networks carefully. This paper describes a verification-based scheme for coordinating multiple tenants' OpenFlow networks. The scheme enables administrators to design each tenant network without having to worry about conflicts with other tenants. It ensures isolation of virtual networks among multiple tenants transparently. It manages the address space overlaps and resolves conflicts in the flow entries.

**Keywords**—OpenFlow; Virtualization; Multi-tenant Network.

## I. INTRODUCTION

With the development of server virtualization technology, cloud computing services, such as Infrastructure as a Service (IaaS), have become popular. Server virtualization technology virtualizes an organization' s IT infrastructure at a data center and provides it through the Internet. In multi-tenant networks, one physical network is divided into many tenant virtual networks. The traffic in each virtual network is isolated from the traffic in other networks. Virtual LAN (VLAN) is a popular virtualization technology. IaaS providers divide one physical network into many layer 2 networks by assigning a VLAN-ID to each tenant virtual network, and the tenant users can then freely construct their own layer 3 network on the allocated tenant virtual network. When the providers of an IaaS cloud using VLAN technology change the configurations of the virtual networks, they need to change the VLAN settings of all the network devices. However, in a cloud environment where the number of virtual networks and virtual machines change rather dynamically, a more flexible virtual network construction and management method is required.

OpenFlow [2], which is a core technology of software-defined networking (SDN) [1], has the features that satisfy these requirements. OpenFlow enables flexible routing control and centralized management of networks by separating

the control plane from the data transfer plane. A controller controls the routing of packet forwarding, and the data plane switches transfer packets in accordance with the instruction of the controller. Since this technology has the ability to recognize and rewrite the VLAN-ID of each packet, IaaS providers can aggregate VLAN management functionalities into one controller. The OpenFlow based network architecture also enables flexible virtual network management; however, a tenant network may accidentally disable OpenFlow functionalities when the IaaS provider and user tenants have different control policies. The administrators of each tenant network may thus have difficulty gaining the benefits of OpenFlow, if the provider uses OpenFlow technology to manage its IaaS platform. The idea of coordinating multiple OpenFlow networks on a physical network would enable individual virtual networks to be managed by a single tenant.

FlowVisor [3] is a technique that handles requests from multiple OpenFlow controllers. In FlowVisor, a proxy is placed between the OpenFlow controller and the switches, and it exchanges and manages each tenant' s control messages sent between the controllers and switches. This enables OpenFlow switches to be individually controlled by multiple controllers on one physical OpenFlow network. FlowVisor expresses a tenant network space in a way that is called a flow space, and the administrator of each tenant writes flow entries belonging to the allocated flow space definition by using their own controller. This mechanism can be used to construct a plurality of virtual OpenFlow networks, and it enables each tenant controller to control each tenant' s virtual network individually. FlowVisor assumes that there is no overlap between flow spaces. When applying it to a multi-tenant network, each tenant network must define its flow entries within the flow space provided by the IaaS provider. This problem becomes more difficult because the flow spaces are not always discrete. In the case of monitoring one tenant' s flow space from another flow space it owns, the flow spaces must overlap. The IaaS provider needs to define them very carefully so as not to cause unintended traffic control.

In this research, we propose verification-based OpenFlow network virtualization based on FlowVisor that enables the network to be freely designed by each tenant. To guarantee traffic separation, we propose a conflict management that uses verification of flow space definitions. If a conflict occurs, it can be resolved by rewriting a flow entry. Our approach verifies

and manages overlapping parts between flow spaces defined by individual tenants, detects conflicts between flow spaces and flow entries, and rewrites the entries to avoid conflict in the FlowVisor. This paper describes the method of flow-space verification among multiple tenants and its implementation. Section II is an overview of OpenFlow/SDN technology. Section III explains the mechanism and problems of FlowVisor. Section IV outlines the proposed virtualization method based on flow entry verification, and Section V describes the method for avoiding flow entry conflicts in more detail. Section VI describes the details of our prototype implementation and its performance evaluation. Section VII discusses our method in relation with other research. Section VIII is a conclusion that mentions future work.

## II. OPENFLOW/SDN

OpenFlow is a representative architecture of software-defined networking, and it is currently being standardized. It is a next-generation network technology for cloud computing environments. An OpenFlow network consists of an OpenFlow controller responsible for routing control and an OpenFlow switch for transferring packets according to flow entries written by the controller. Hence, it is a centralized control architecture that enables centralized management of networks by separating the traditional network system into a control plane and data plane.

The controller is software, and a pair of matching fields, such as a MAC address, an IP address, a transport number, a VLAN-ID, and actions to be performed on a packet are defined as a flow entry. Flexible routing control is enabled by transferring packets according to flow entries in the switch. If the switch has to be reconfigured in response to a change in the network configuration, the change is applied to all the switches by describing the change settings as new flow entries in the controller. This improves the manageability of the network. The controller and switch are connected by an OpenFlow channel, which is a control network using TCP/IP that is constructed separately from the data network, and they exchange control messages called OpenFlow messages through it. Through OpenFlow messages, the controller controls switches such as for writing the flow entry. In OpenFlow, since the controller controls all the switches and knows the network topology, it is possible for it to control routing flexibly such as through source routing and multi-path forwarding. Virtualizing a physical network by using OpenFlow makes it possible not only to improve the manageability of VLAN-IDs but also to ensure logical division of the network by using the packet headers of layers 1 to 4 that can be specified as a match field. OpenFlow enables its users to create a number of virtual networks beyond the usual limits of VLAN-IDs by dividing up the used address space in advance.

However, the conventional OpenFlow technology has some problems when it comes to virtualizing and controlling the OpenFlow network itself. For example, it is not possible to control each switch individually from multiple controllers in one OpenFlow network, and there is no mechanism to logically divide one OpenFlow network into multiple virtual OpenFlow networks, etc. These problems make it impossible for a tenant to construct and control each controller or devise a virtual OpenFlow network in a multi-tenant data center that provides IaaS.

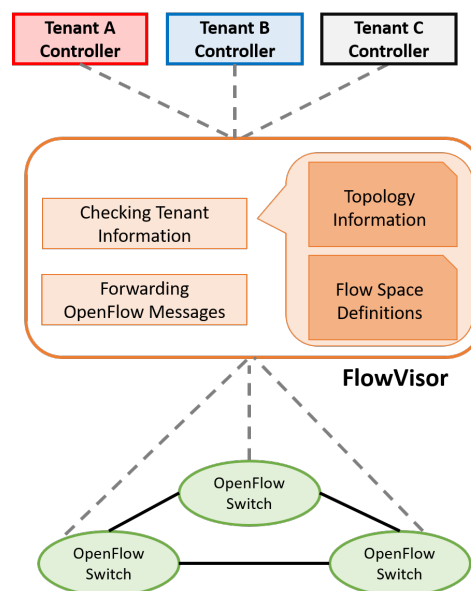


Figure 1. FlowVisor

## III. FLOWVISOR

A FlowVisor is placed in an OpenFlow channel that connects the controller and switches, as shown in the Figure 1. It operates as a proxy that transfers the OpenFlow messages necessary to control the switch from the controller. The administrator of FlowVisor defines the available network space to each tenant as a flow space and presents flow space information to each tenant user in some way. Each tenant user creates flow entries and a controller for writing them in accordance with the network topology and flow space information of the virtual OpenFlow network presented by the FlowVisor administrator. A tenant user can control the tenant network by connecting his controller to FlowVisor.

### A. Flow Space

It is necessary for the administrator of FlowVisor to define the available network space in each tenant as a “flow space” in advance. As shown in Table I, a flow space has a slice name indicating the name of the tenant network, a DPID that indicates the OpenFlow switch ID, and a MAC address, IP address, transport number, etc., as an available match field from layer 1 to 4 in a flow entry and priority. In addition, each flow space is based on the premise that the defined network space is independent and has no overlaps. Therefore, there is no mechanism for checking whether flow space conflicts exist in FlowVisor, and hence, the administrator needs to define each flow space carefully.

### B. FlowVisor Mechanism

FlowVisor functions as a proxy on the OpenFlow channel and controls the transfer of OpenFlow messages between multiple controllers and switches. This function differs between the case of transferring messages from the switch to the controller, such as when sending Packet-In and Port-Status messages, and the case of forwarding messages from the controller to the switch, such as when sending the Flow-Mod message.

TABLE I. EXAMPLES OF FLOW SPACE

Slice	DPID	Priority	VLAN	Src MAC	Dst MAC	Src IP	Dst IP	Src TCP	Dst TCP
Tenant A	1	100	50	*	*	*	*	80, 22	*
Tenant B	1	100	50	*	*	10.0.1.0/24	*	80	*
Tenant C	1	100	50	*	*	10.0.2.0/24	*	80	*

First, we describe the messages that are transferred from the switch to the controller. In this case, it is necessary to specify the controller to which the message pertains before transferring the message to it. As an example, a Port-Status message notifying that the physical port state of the switch has changed will affect all the tenant controllers using that port. Accordingly, FlowVisor searches for all target controllers from the topology information of each tenant network and transfers the Port-Status message to all of them. In the case of a Packet-In message, FlowVisor searches the flow space definition to specify which tenant network the packet belongs to and forwards the message to the tenant controller of the corresponding flow space.

Next, we describe the messages that are transferred from the controller to the switch. In this case, FlowVisor refers to the topology information of all the tenant networks; then it transfers the message to the target switch; it performs the same operation on every message. If a tenant user tries to send a message to the switch that does not belong to its own tenant network, the send operation fails and a message transfer error is returned to the controller.

### C. FlowVisor Problem

FlowVisor is based on the premises that the flow spaces allocated to each tenant network are independent and the tenant controller sets flow entries within the allocated flow space. If a FlowVisor administrator defines an unintended or incorrect content flow space, an unexpected network control will be executed. In contrast, if IaaS providers want to enable each tenant user to freely design their own tenant network as way of a providing a multi-tenant network, the flow space should be able to be freely defined by each tenant user. There is a problem that unintended traffic control can occur when a flow entry is written that conflicts with the flow space of another tenant. Hence, it is necessary to implement a mechanism that can check for conflicts in flow spaces and flow entries in a multi-tenant network.

Table I shows an example of conflicting flow entries, wherein if tenant user A tries to write a flow entry that prohibits the SSH session such as by sending “Src TCP = 22, action = DROP” to the switch with DPID = 1. In Table I, match fields of tenant A are defined as wildcard values “ \* ” with the exception of Src TCP; thus tenant user A can freely use this value. However, if the flow entry such as what is mentioned above is written, it will be applied to all packets that are transferred through this switch with source TCP port number 22. Since all the packets are dropped, all SSH connections are closed even in other tenant networks. In this case, the packet was dropped unintentionally, however, it is possible to rewrite the packet header as a specified action and transfer it in OpenFlow. It is also possible to act in dubious or illegal ways, such as eavesdropping by transferring traffic of other tenants that are not permitted to use a server on their tenant

network. In particular, it is also possible to transfer the traffic of other tenants to a server on one’s own tenant network for the purpose of sniffing packets.

If a FlowVisor administrator allows each tenant user to freely design their tenant network and flow space definition, a flow space that has overlaps will cause unintended behavior because the flow entries conflict. This is due to OpenFlow’s ability to flexibly set values such as wildcards about L1-L4 headers in the match field. In the example mentioned above, since tenant user can write a flow entry with wildcards other than the source TCP port number to the switch, he can control the traffic in unassigned flow spaces.

## IV. VIRTUALIZATION BASED ON FLOW ENTRY VERIFICATION

We propose a virtualization method for an OpenFlow network that enables a network to be designed for each tenant. In particular, we propose a verification and management system of duplications in the flow space allocated to each tenant and a conflict verification and rewriting method for the flow entries written by tenant controllers. As shown in Figure 2, the verification is implemented in FlowVisor. First, this system verifies and manages the overlapping address spaces in each flow space. A tenant user defines the combination of address spaces that s/he will use in each tenant network as a “flow space”, and this system verifies and manages duplications. It is possible to avoid conflicts of flow entries among tenants as much as possible. In addition, when a flow entry in a flow space includes overlapping address spaces with others was written, it checks for a conflict of the flow entry and rewrites the match field to guarantee the separation of traffic between each tenant network. This minimizes the amount of rewriting of flow entries by applying verification and management on the flow space in advance.

Our system uses a new definition of a flow space. It is constructed by restricting the elements and combinations of match fields against the existing definition. It is possible to set arbitrary values for all elements of the OpenFlow match field in the existing definition. In this case, tenant users can write flow entries that cause unintended traffic control when using wildcards. On the other hand, our method restricts tenants to using only a combination of address spaces that have pre-specified range as a match field. We make it so that a tenant user can control only the allocated network and traffic; furthermore, we make it unnecessary to verify fields that are not specified in a practical network.

### A. Flow Space Definition

This flow space is different from the definition of FlowVisor in Section III-A. In previous work, a flow space was defined for each switch that the tenant can control; however, here, a new flow space is defined as a combination of address spaces that the tenant can use for one tenant network. A flow space is

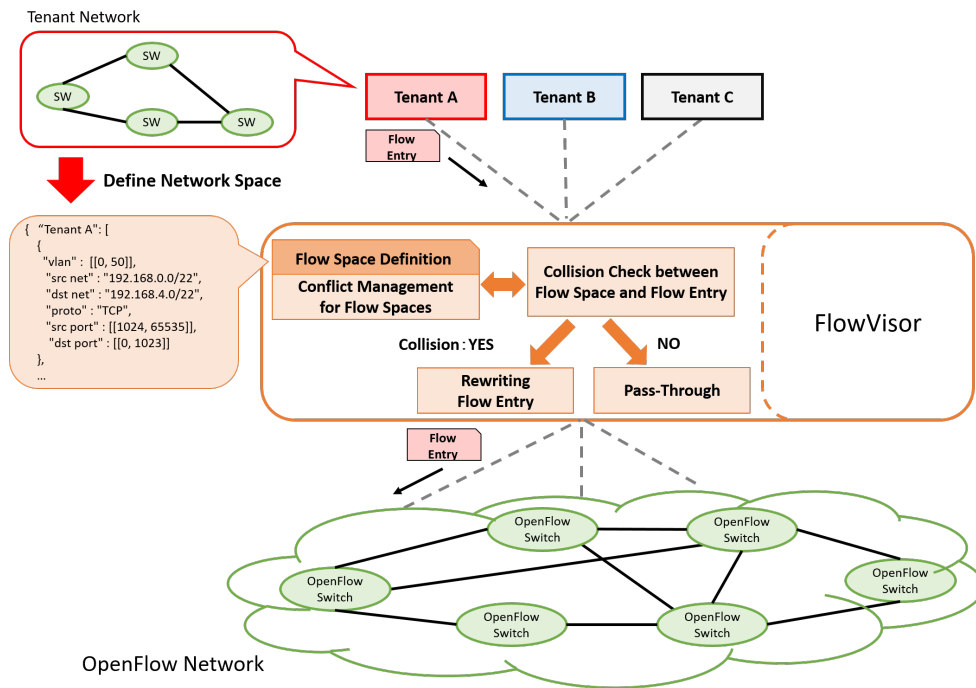


Figure 2. Proposed architecture

composed of multiple rules, where each rule consists of rule IDs, flow space names, and a matching field that is available to the tenant, as shown in Table II. In the matching field, it is possible to set five kinds of header information from L2 to L4 as VLAN ID, Src/Dst IP address and Src/Dst TCP port, which are necessary for network operations. These definitions are described in JSON format, as shown in Figure 3. Each flow space describes a flow space name and a set of flow definitions. A flow definition is described for each element of a match field, and it is defined as conjunctions of fields. Since one flow space is represented by one or more flow definitions, multiple flow definitions are defined as disjunctions to allow flow entries that match any one. Each tenant uses only the combination of address spaces specified in this flow space. Definition example 2 in Table II, which summarizes the examples of Figure 3, shows the following address space:

- VLAN ID = 100, Src IP = 192.168.64.0/20, Dst IP = 192.168.64.0/20, Src TCP = 80
- VLAN ID = 101, Src IP = 192.168.64.0/20, Dst IP = 192.168.64.0/20, Src TCP = 80

The tenant assigned this flow space can control the network by using these two different combinations as a match field of the flow entry. The top row of Table II shows the available address space as the match field, but the upper limit of the VLAN ID is half the original limit of 4096. This is due to securing the independent address space as management space for managing duplications of flow spaces and resolving conflicts in advance. VLAN-IDs are allocated from this management space to the flow space when necessary.

**B. Duplicate Flow Space Verification and Flow Entry**

Now let us explain the overlap verification between flow spaces and conflicts of flow entries on the basis of the

```

{
  "Example 1" : [
    {
      "vlan" : [[0, 50]],
      "src net" : "192.168.0.0/22",
      "dst net" : "192.168.4.0/22",
      "proto" : "TCP",
      "src port" : [[1024, 65535]],
      "dst port" : [[0, 1023]]
    },
    {
      "vlan" : [[0, 50]],
      "src net" : "192.168.4.0/22",
      "dst net" : "192.168.0.0/22",
      "proto" : "TCP",
      "src port" : [[0, 1023]],
      "dst port" : [[1024, 65535]]
    }
  ],
  "Example 2" : [
    {
      "vlan" : [[100], [101]]
      "src net" : "192.168.64.0/20",
      "dst net" : "192.168.64.0/20",
      "proto" : "TCP",
      "src port" : [[80]],
      "dst port" : [[80]]
    }
  ]
}
    
```

Figure 3. JSON Format for Flow Space

definition in the previous section. Table III lists examples of flow spaces defined for three tenants A, B, and C. Since the flow space definition of the tenant A at the top row completely includes the flow spaces of the following tenants B and C, flow space A overlaps B and C and is not independent. On the other hand, in the flow spaces of tenants B and C that are independent in Table II, independent values are specified for

TABLE II. FLOW SPACE LIMIT AND DEFINITION EXAMPLES

Rule ID	Space Name	VLAN	Src IP	Dst IP	Src TCP	Dst TCP
1	Maximum usage	0 ~ 2047	0.0.0.0 ~ 255.255.255.255	0.0.0.0 ~ 255.255.255.255	0~65535	0~65535
2	Example 1	0 ~ 50	192.168.0.0/22	192.168.4.0/22	1024 ~ 65535	0 ~ 1023
3	Example 1	0 ~ 50	192.168.4.0/22	192.168.0.0/22	0 ~ 1023	1024 ~ 65535
4	Example 2	100, 101	192.168.64.0/20	192.168.64.0/20	80	*

any of the match fields, such as Src IP address. Since only the conjunction of the combination of the address spaces is used as a match field in our definition of the flow space, we can detect for duplications by verifying the inclusion relation for each combination of address spaces.

If the flow spaces have a complete inclusion relation, one must detect and avoid conflicts of flow entries after managing any flow space duplication. In flow spaces such as in Table III, the flow entry at the top of the Table IV written from tenant A's controller will collide with the flow entries of other tenants. Table IV shows two examples, i.e., one that conflicts with other flow space definitions and another that does not conflict with others. In the example of flow entry at the upper row, the value of Src IP is a wildcard and it is based on the flow space of tenant A. Since it includes the range of flow spaces in other tenants B and C, it conflicts with their flow entries, and their traffic is also controlled by this conflicting flow entry. On the other hand, in the example of the flow entry in the lower row, Src IP = 10.0.0.1, which is an independent value against the flow space of other tenants is set in the match field. This flow entry does not cause a conflict. As mentioned above, we must verify the inclusion relation of the value specified in the match field for each flow space. If the value includes other tenant's flow spaces, it is possible to verify and avoid conflict by extracting a new value from the free independent address space and setting it to a conflicting flow entry.

## V. CONFLICT VERIFICATION OF FLOW ENTRY

To avoid conflicts between flow entries, we propose a two-step verification method. The first step involves checking the consistency between the address space defined in the match field of the flow entry and its own flow space. In the second step, for the match field in the flow entry, the part of the wildcard including the value defined in the flow space of the other tenant is automatically expanded into a free independent address. As a result, conflicts due to flow entries using wildcard values are detected and avoided while at the same time different flow entries are prohibited from the defined flow space. These measures guarantee that traffic of the different tenant networks is separated.

### A. Consistency Check with Flow Spaces

A consistency check is made of the flow entry in the Flow-Mod message from the tenant controller as to whether the match field deviates from the tenant's flow space definition. The consistency check simply compares the range of the address space for values other than wildcards in the match field to see if they go beyond the range defined in the flow space. If a flow entry with a value beyond that of the flow space definition is written, the Flow-Mod message is discarded and a transmission error for the Flow-Mod message is sent to the tenant controller.

### B. Expanding Wildcard Parts

We rewrite the flow entries that passed the consistency check of Section V-A so that the wildcard part of the match field does not conflict with the address space defined in the other flow space. Here, as with the example in Section III-C, we will explain the case of writing a flow entry such as "drop all packets with the source TCP port number 22" from the tenant A controller in Table I. In this case, because the value of the source TCP port is within the flow space first, it passes the consistency check of the Section V-A. Next, all the match fields except for the source TCP port are filled in as wildcards, but these include tenant B's flow space for VLAN-ID and source IP address as well as tenant C's space for the source/destination IP address in Table I. For avoiding conflicts between flow entries, one or more independent values are set for each of these wildcards. The result of rewriting the flow entry using the free address space is shown on the lower row of Table V. Our method rewrites the wildcard part the flow entry so that the match field does not conflict with others and transfers it in a Flow-Mod message. In so doing, it is guaranteed that the flow entry will not incorrectly control the traffic on another tenant network.

## VI. IMPLEMENTATION

Our core methods consisted of two flow space conflict verification systems, i.e., the "flow space manager" and "flow translate engine", and we implemented a prototype flow space manager. This section describes the implementation and initial performance evaluation. The flow space manager holds definitions of the given flow space and investigates in advance the flow space where flow entries can collide. Here, the flow space is defined as shown in Figure 3; the manager analyzes it and holds flow definitions for each flow space. At this time, in each flow definition, a flow definition that is a duplicate of one of the other flow spaces in all match fields may cause a conflict.

The flow definition is managed by hashing the source IP address space with the network address of 24 bit prefix as the key. In this case, if the source IP address space of the flow definition is narrower than /24, the network address of the /24 network including it is used as the key. If it is larger than /24, network addresses of all /24 networks are registered as multiple entries.

This manager was implemented in Ruby 2.3, and the initial performance evaluation measured the overhead of flow registration. We measured the change in the time taken from the 1st to 5000th in two cases of 5000 flow spaces that did not contain any conflicts and 5000 flow spaces that completely contained conflicts. The results, as measured by a computer with Intel Core-i7 2.8GHz, 16 GB memory, are shown in Figure 4. The flow spaces where conflicts occurred are slower, but could be processed at a rate of about 0.2 ms per entry. Considering that

TABLE III. EXAMPLES OF DUPLICATE FLOW SPACES

Rule ID	Space Name	VLAN	Src IP	Dst IP	Src TCP	Dst TCP
1	Tenant A	50	*	*	80, 22	*
2	Tenant B	50	10.0.1.0/24	*	80	*
3	Tenant C	50	10.0.2.0/24	*	80	*

TABLE IV. EXAMPLES OF FLOW ENTRIES IN TABLE III

Entry	Match Field	Action
Conflicting Flow Entry	VLAN ID = 50 Src TCP = 80	Output: port 2
Non-Conflicting Flow Entry	VLAN ID = 50 Src IP = 10.0.0.1 Src TCP = 80	Output: port 2

TABLE V. REWRITING WILDCARD PARTS

Entry	Match Field	Action
Conflicted Flow Entry	Src TCP = 80	DROP
Rewritten Flow Entry	Src TCP = 80 Src IP = 10.0.0.0/24 VLAN ID = 2048	DROP

the flow space registration is relatively infrequent, this result indicates sufficient practical performance. The flow translate engine is currently being implemented; however, it examines only the definitions of target flow spaces, accordingly, the engine searches fewer flow space definitions than in the flow space manager.

VII. DISCUSSION

We proposed an OpenFlow network virtualization scheme that allows each tenant to freely use OpenFlow technology in a multi-tenant network environment. The features of this scheme include virtualization of an OpenFlow network by using conflict management in a flow space abstracting individual tenant networks and conflict verification of each flow entry. The designer of each tenant network can freely design the network configuration by defining the network address field such as the IP address.

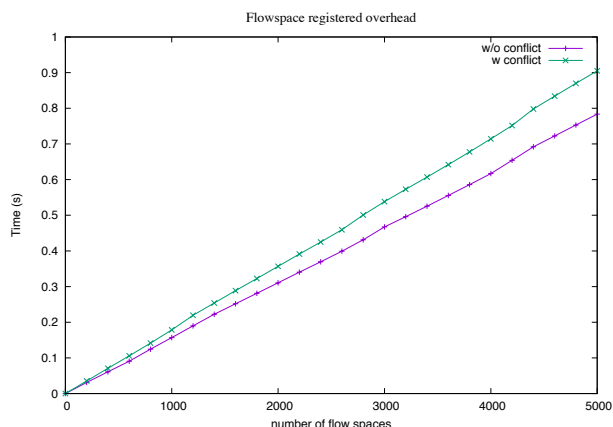


Figure 4. Overhead of Flow Registration

FlowVisor requires that each flow space never overlap; thus, it cannot verify whether conflicts occur between flow spaces. This means that conflict avoidance among flow spaces is left to the operator’s responsibility. From this point of view, it seems reasonable to view it as a network partitioning technique rather than a virtualization. Sköldström [4] et al. propose virtualization method that uses FlowVisor as a relay network of a wide area network. They focus on resource management, whereas our research mainly deals with mapping to lower-layer network separation technology such as MPLS. Yamanaka et al.’s [5] virtualization method works by assigning and tagging a specific MAC address for each virtual network at the edge of the network. This method restricts flow definitions to those that can be described by each tenant.

Our method can freely define a virtual network for each tenant and realizes a control that maintains its independence. As a result, based on the design and construction of the established TCP/IP network, users can introduce flexible controls by using OpenFlow technology. Even when the backend of the IT infrastructure of the current organization is moved to the cloud environment, it will be possible to provide both flexible network control and ease of design like that of a conventional network.

VIII. CONCLUSION

We proposed a virtual network management system that maximizes the ability of OpenFlow virtualization by using verification of the flow space definition. The method enables individual tenant networks to be freely designed in a multi-tenant network environment and ensures isolation among them. This makes it possible for IaaS providers to provide a flexible tenant network in which OpenFlow technology is freely used for and by each tenant user. A preliminary evaluation of a prototype shows that the proposed flow space management has sufficient performance.

ACKNOWLEDGMENT

This work was supported by JSPS KAKENHI Grant Number JP15K00138.

REFERENCES

- [1] N. McKeown, "Software-defined networking," INFOCOM keynote talk, vol. 17, no. 2, pp. 30-32, 2009.
- [2] N. McKeown et al., "OpenFlow: enabling innovation in campus networks," ACM SIGCOMM Computer Communication Review, vol. 38, Issue 2, pp. 69-74, April 2008.
- [3] R. Sherwood et al., "FlowVisor: A Network Virtualization Layer," Tech. Rep. OPENFLOW-TR-2009-01, OpenFlow Consortium, October 2009.
- [4] P. Sköldström and K. Yedavalli, "Network Virtualization and Resource Allocation in OpenFlow-based Wide Area Networks," IEEE International Conference on Communications (ICC), pp. 6622-6626, June 2012.
- [5] H. Yamanaka, S. Ishii and E. Kawai, "Realizing Virtual OpenFlow Networks by Flow Space Virtualization," IEICE Technical Report, Network Systems, vol. 112, no. 85, pp. 67-72, June 2012.