


WIRELESS DATA CENTER NETWORKING

YONG CUI AND HONGYI WANG, TSINGHUA UNIVERSITY

XIUZHEN CHENG, THE GEORGE WASHINGTON UNIVERSITY AND KING SAUD UNIVERSITY

BIAO CHEN, UNIVERSITY OF MACAU



The authors analyze the challenges of DCNs and articulate the motivations of employing wireless in DCNs. We also propose a hybrid Ethernet/wireless DCN architecture and a mechanism to dynamically schedule wireless transmissions based on traffic demands.

ABSTRACT

Data centers play a key role in the expansion of cloud computing. However, the efficiency of data center networks is limited by oversubscription. The typical unbalanced traffic distributions of a DCN further aggravate the problem. Wireless networking, as a complementary technology to Ethernet, has the flexibility and capability to provide feasible approaches to handle the problem. In this article, we analyze the challenges of DCNs and articulate the motivations of employing wireless in DCNs. We also propose a hybrid Ethernet/wireless DCN architecture and a mechanism to dynamically schedule wireless transmissions based on traffic demands. Our simulation study demonstrates the effectiveness of the proposed wireless DCN.

INTRODUCTION

Cloud computing has been widely used to provide services such as search, email, and distributed storage systems. In recent years, more and more data centers are built to support the expansion of cloud applications. As the key infrastructure, data center networks (DCNs) are required to scale to accommodate a large number of servers and supply adequate bandwidths for various applications such as Map-Reduce.

Evolved from large-scale enterprise networks, current DCNs are typically constructed based on a hierarchical topology. As shown in Fig. 1, servers are arranged in racks, and a top-of-rack switch (ToR) connects all the servers in a rack. These racks, and the switches of the aggregation and core layers in a DCN form a multi-root tree. Since only a few switches serve as root nodes at the core layer, it is obvious that these root nodes would become the bottleneck if the global traffic volume is large. Therefore, transmissions between servers of different racks suffer from oversubscription, which means that the actual throughput is much lower than the available bandwidth of the servers' network interfaces.

In order to tackle these challenges, works have been done to improve the structures of DCNs as well as to design new routing mechanisms [1–5]. The common feature of these works

is to add new paths, especially disjoint paths, so as to increase the potential end-to-end capacity. These approaches alleviate the problems of DCNs to some degree.

Despite the enhancement mentioned above, it is still difficult to handle the servers with a large volume of outburst traffic in Ethernet structures, due to the static links and finite network interfaces. These servers cause congestions and thus have a negative effect on other servers. Furthermore, as servers in DCNs usually collaborate to perform joint tasks, the servers suffering from high traffic may further affect the global performance.

Therefore, approaches beyond Ethernet have been investigated to address these issues. Wireless data center networking, based on the state-of-the-art wireless technologies, provides a promising solution. Wireless links can be set up without the cost of wiring; the flexibility of wireless connections makes it much more convenient to adapt the topology to the requirements of real-time transmissions; direct wireless links between servers also alleviate the load of core switches.

Wireless data center networking is first introduced by [6], in which flyways are established by adding wireless links between ToRs to alleviate the congestion problem of hot racks to minimize the maximum transmission time. However, no technical approaches are detailed to address various issues faced by a practical realization of a wireless DCN. In this article, we present a novel approach to realize a wireless DCN. A hybrid Ethernet/wireless architecture is proposed, and a distributed scheduling is investigated to arrange the wireless links. In our scheduling mechanism, we model the new infrastructure and formulate an optimization problem for which a heuristic algorithm is designed. We also conduct a simulation study to demonstrate the effectiveness of utilizing wireless transmissions in data centers.

The rest of the article is organized as follows. We provide the motivations for bringing wireless to DCNs. Our wireless DCN architecture together with its scheduling mechanism is depicted, and the channel allocation problem is elaborated. The methodology and evaluation results are presented. At last, we conclude this article.

MOTIVATIONS FOR USING WIRELESS IN DCNS

CHALLENGES IN DCN

As mentioned earlier, the tree-based structures in current data centers suffer from oversubscription. The problem worsens if a few servers generate high traffic as these servers, usually called *hot servers*, may become the bottleneck of the system.

However, the capacity provided by the infrastructure is not always inadequate. According to the study on traffic distributions of DCNs, only a few servers generate a high volume of traffic, and the traffic matrix is in general quite sparse [6]. Moreover, the traffic distributions of DCNs are highly dynamic, and the links are not always utilized at their full capacity. At any time, about 60 percent of the core links and edge links are active, while the utilization rate is much lower for aggregation links, of which the 95th percentile utilization is below 10 percent [7]. These statistics indicate that the average loads of DCNs are not very high, and the oversubscription mainly results from the high outburst traffic at hot servers.

These phenomena motivate the existent effort to handle the problem of oversubscription by either extending tree-based structures [4, 5] or designing new topologies [4]. Fat-tree [1] is a typical tree-based structure, in which servers are grouped into pods, and core switches connect pods together. Servers belonging to the same pod share several paths, and the traffic of hot servers can be assigned to multiple available paths. An exemplary server-based topology is BCube [5], which constructs a recursive server-centric topology. By involving servers in data forwarding, it avoids the oversubscription problem of core switches. Moreover, by employing multiple disjoint end-to-end paths, servers in BCube are able to distribute flows to achieve load balance.

Although these approaches are effective to some degree, the high loads of core switches in tree-based topologies are difficult to avoid, and the low efficiency of server-based data forwarding restricts the performance of server-centric structures. Fundamentally, it is necessary to set up more links to increase the capacity for hot servers. However, the non-deterministic nature of the traffic distributions of DCNs makes it impossible to solve the problem by adding extra links to a certain group of servers. And it is also inadvisable to add links to all the servers due to high cost and the difficulty of wiring. In brief, traditional Ethernet-based solutions get stuck in handling the oversubscription problem in DCNs. In order to address these challenges, it is imperative to design novel approaches that are able to flexibly provide additional capacity for hot servers and can easily be realized with current hardware technologies. That is why we turn to wireless.

EXTREMELY HIGH FREQUENCY WIRELESS COMMUNICATIONS

In constructing data centers, wireless has unique advantages over Ethernet. First, it brings convenience to the deployment and maintenance of a DCN. For a large-scale data center, it usually




Figure 1. A typical DCN structure.

takes a great amount of manual effort to wire the huge number of servers, which is inherently hard and error-prone. This problem is especially severe for those improved DCN solutions as they introduce more wires than conventional DCN architectures. By using wireless, these difficulties can be considerably reduced. Second, wireless can enhance the flexibility of a DCN. Since wireless links can be dynamically established, it is possible to perform adaptive topology adjustment. In other words, the network can be rearranged to fulfill the real-time traffic demands of hot servers. Furthermore, since wireless connections no longer rely on switches, they are free of the problems caused by these centralized devices, such as single-point failures and limited bisection bandwidth.

Nevertheless, there are some significant challenges in introducing wireless to data centers. One concern is whether wireless can provide high-speed transmission support for DCN applications. Fortunately, advances in wireless technologies have set the stage for high-data-rate communications. Extremely high frequency (EHF), which ranges from 30 to 300 GHz, is a choice for high-speed wireless solutions. In particular, the 60 GHz spectrum provides a 7 GHz (57–64 GHz) waveband and is able to support a data rate over 1 Gb/s. The small wavelength of radio signals also supports highly directional communications to increase frequency reuse. Although the transmission range is not very large (typically about 10 m), it is adequate for short distance indoor wireless communications [8].

Another problem is the instability of wireless communications, especially in data centers full of metal devices. As for 60 GHz communications, the dominant factor in signal attenuation is the effect of oxygen absorption. Experiments have been performed to evaluate the characteristics of 60 GHz channels in various indoor environments. The results indicate that there is not much difference between the delay of 60 GHz omnidirectional antennas in an office and that in a laboratory with highly reflective metallic equipments [9]. Furthermore, as diffraction and reflection can hardly apply to high-frequency radio signals, the 60 GHz communications are mainly based on line-of-sight (LOS) propagation. For LOS transmissions, the path loss in a laboratory is only a little higher than that in other indoor environments [10]. In fact, in order to address the high path loss, 60 GHz radios usually utilize beamforming to achieve a high gain [11]. Therefore, the servers and racks in data centers would

A wireless DCN architecture should satisfy the basic requirements of DCNs, including scalability, high capacity, and fault tolerance. Nevertheless, it is difficult for a wireless network alone to meet all the demands.

not cause much of a problem to the channel environment as long as they do not obstruct LOS transmissions.

In addition to performance considerations, there are also some concerns about the cost of a wireless DCN. For example, using wireless would aggravate the energy consumption of data centers as wireless communications requires more power than wired; wireless devices are usually more expensive than wired network interface cards (NICs), especially for 60 GHz equipment. However, as long as the increment is acceptable compared with the total cost of DCN, wireless is still an advisable approach. With progress in the standardization and manufacture of 60 GHz communications, the cost issues would be gradually resolved. In fact, IEEE 802 has started its research in EHF communications, and 802.11ad, which standardizes very high throughput transmission at 60 GHz, will be ratified in 2012. Some proposals in IEEE 802 study the collaboration of 2.4/5 GHz and 60 GHz. One of the use cases is to employ a 2.4/5 GHz-based management mechanism to discover and schedule 60 GHz beamforming. This use case provides a feasible method to manage the EHF transmissions and improve the performance of 802.11ad over a large range. Moreover, WirelessHD, a specification for wireless high-definition signal transmissions for consumer electronics products, is also based on 60 GHz communications. The core technology of the specification allows a data rate as high as 25 Gb/s theoretically. In fact, prototype devices of WirelessHD have already been produced [12].

REQUIREMENTS OF WIRELESS DCNS

There are two issues we need to handle in order to take advantage of wireless transmissions in a DCN: the design of the hybrid (wireless and wired) network architecture and the scheduling of wireless links.

In designing the network architecture, the basic requirements of a DCN, including scalability, high capacity, and fault tolerance, should be addressed. The limited transmission range of EHF and the interference among wireless communications further complicate the problem. Moreover, how to assign radios among servers is also an important issue.

Wireless scheduling should be performed according to the traffic demands as we use wireless transmissions to alleviate the congestion of hot servers. Additionally, the activities of the wireless links should adapt to the changing traffic distributions. On the other hand, centralized scheduling mechanisms leading to high control overhead should be employed with caution. For example, it is prohibitive to exchange information among all the servers, which is typically required by a centralized controller.

OUR APPROACH TO WIRELESS DCNS

Although wireless technology provides an alternative to improve the performance of a DCN, it still requires careful designs to compose a feasible and effective technical approach. In this section, we present our exploratory research on wireless DCN design. We first elaborate the

basic architecture under our consideration, then propose a mechanism to schedule the wireless links to improve the performance of a wireless DCN.

WIRELESS DCN ARCHITECTURE

As mentioned above, a wireless DCN architecture should satisfy the basic requirements of DCNs, including scalability, high capacity, and fault tolerance. Nevertheless, it is difficult for a wireless network alone to meet all the demands. For example, the capacity of wireless links is usually limited due to the interference and high transmission overhead. Thus, wireless networks could not be employed to entirely substitute the Ethernet. Since our motivation to introduce wireless transmissions is to alleviate the congestion of hot servers, we take wireless communications as the supplement to wired transmissions. The main idea of our approach is to add extra wireless links to the existing Ethernet topology to construct a hybrid Ethernet/wireless architecture.

One prerequisite of utilizing wireless communications in DCNs is to equip servers with radios. An intuitive approach is to assign radios to each server. However, this leads to a large amount of radios, resulting in not only high cost but also waste of wireless devices as the limit of wireless channels allows only part of the radios to transmit simultaneously. Therefore, it is more reasonable to assign radios to groups of servers. In the following, we use the term *wireless transmission unit* (WTU) to refer to a group of servers supported by the same set of radios for communicating to the servers out of the group.

In reality, data centers are mainly constructed by connecting racks of servers via Ethernet. Therefore, it is reasonable to consider each rack as a WTU, as illustrated by Fig. 2. Note that the racks do not block the LOS transmissions as the radios are located on top of them. As for other interconnection mechanisms, many of them share the feature that servers are organized in groups. For example, Fat-tree is constructed based on Pods [1]; the topology of BCube is made up of a number of BCube0s [5]. Obviously, it is appropriate to take the basic architecture of these mechanisms as a WTU.

In short, our hybrid wireless DCN is constructed by adding wireless infrastructure to the existing Ethernet-based architecture. It can be applied to various DCN topologies without the cost of rearranging the servers.

SCHEDULING MECHANISM

According to the analysis earlier, we design a distributed scheduling algorithm that can be applied periodically to adapt to the traffic distributions. Note that in our design, feasibility and effectiveness are considered as the objectives with the highest priorities.

Collecting Traffic Demands — As the scheduling of wireless links should adapt to traffic demands, it is necessary to first collect traffic information. We adopt the idea of 2.4/5 GHz assisted 60 GHz communications by utilizing the broadcast at 2.4/5 GHz to disseminate traffic information. In order to avoid high control overhead caused by

the large number of servers, we employ a hierarchical architecture for transmitting control information.

Specifically, a server within a WTU is appointed as the *unit head* for the WTU. The unit head is responsible for collecting local traffic information and executing the scheduling algorithm. Since each unit head only manages the servers in its unit, the overhead is acceptable. Each unit head is equipped with a control radio, and all the units broadcast their traffic demands over a common 2.4/5 GHz channel in a polling manner. Consequently, all the units learn the global traffic distribution and can perform wireless link scheduling independently. Figure 3 illustrates the dissemination of the traffic demands.

Allocating Channels — After collecting the traffic demand information, the head servers need to assign channels to wireless transmissions. This procedure needs to address interference, limit of radios, and the cooperation with Ethernet. We study this problem and depict an algorithm later.

Overall Scheduling Procedure — The scheduling procedure runs periodically in order to adapt wireless links to the dynamic traffic distributions (we assume that the clocks of all the servers are synchronized so that the servers can cooperate to perform scheduling). In each period, each unit head first collects the traffic demand information inside its WTU. Then all the unit heads broadcast their traffic demands in a polling manner. After receiving global traffic information, the unit heads execute the channel allocation algorithm (described later) independently and schedule wireless transmissions based on the channel allocation scheme.

THE CHANNEL ALLOCATION PROBLEM

The channel allocation problem plays a key role in our scheduling procedure because the effect of congestion alleviation mainly depends on whether wireless resources are properly scheduled. In this section, we formulate the problem and present a heuristic to tackle it.

PROBLEM FORMULATION

Wireless Transmissions — A wireless transmission graph $G = (V, E)$ is used to model the architecture of a wireless DCN. In this graph, each node $v \in V$ corresponds to a WTU, and each directed edge $e = (v_1, v_2) \in E$ denotes the transmission from v_1 to v_2 . Since only inter-unit transmissions are assigned to wireless links, there is no self-loop in G .

Each edge e is assigned a weight $t(e)$, which stands for the traffic demand of the corresponding transmission. Let C be the set of channels. Each edge e is associated with a subset of C , denoted C_e , for wireless transmissions. If C_e is \emptyset , no wireless link is set up for that transmission. If C_e has multiple elements, several links are used to accelerate the transmission.

Interference — We adopt a conflict-edge model to formalize the interference relationship between transmissions. Each edge e is associated with a conflict edge set $I(e)$. If a wireless link on chan-

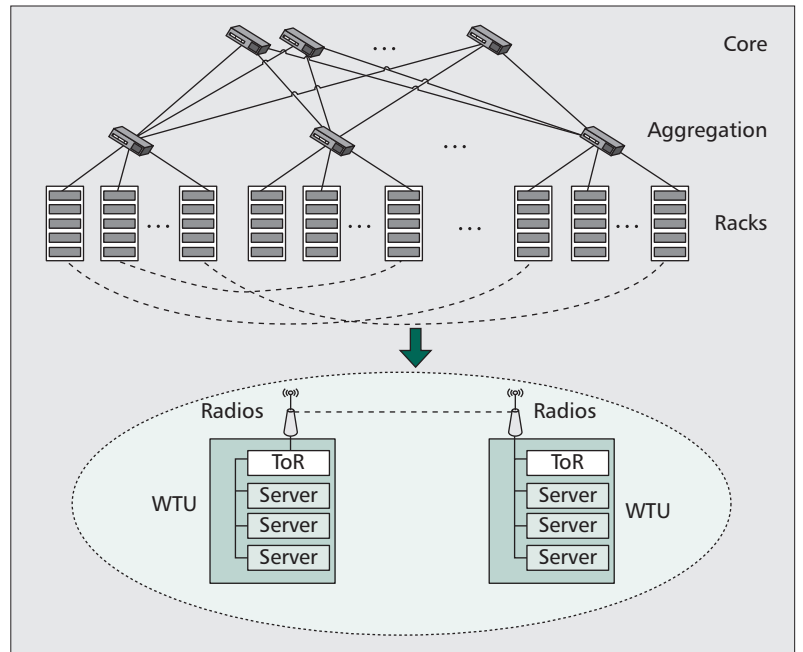


Figure 2. An example wireless DCN architecture, where each rack corresponds to a WTU, and dashed lines denote dynamic wireless links between WTUs.

nel c is set up for an edge e , the transmission of that link succeeds only if no edge in $I(e)$ is active on c .

Several geometric models such as the disk model and the node-exclusive model are commonly used to determine the conflict edge set of an edge [13]. Some of them are based on the interference radius while others are based on the hop count. However, we take EHF communications, which is highly directional, for direct wireless links between WTUs. Therefore, we adopt the node-exclusive model, in which edges sharing a common endpoint interfere with each other.

Let E_v denote the set of edges that take node v as an endpoint. Any two edges e_1 and e_2 in the same E_v are considered as the conflict edges of each other, and hence would cause interference if assigned the same channel. If multiple wireless links are established for an edge e , they take different channels in C_e to guarantee interference-free transmissions.

Limit of Radios — In addition to interference, the number of radios available to each node is also an important limitation factor. A radio cannot transmit data on multiple links at the same time; it is also impossible for a radio to send and receive data simultaneously. Therefore, the number of active wireless links related to a node v cannot exceed the number of radios assigned to v , denoted $r(v)$.

Utility of Transmissions — Channel allocation intends to assign channels to all the edges in a wireless transmission graph. As mentioned above, it plays a key role in network performance enhancement. In order to evaluate the contribution of a channel allocation scheme (i.e., the collection of all the C_e) to the global network performance, several issues should be addressed.

Based on the definition of the wireless transmission utility, we formulate the problem of channel allocation into an optimization problem that maximizes the total utility of all the wireless transmissions.

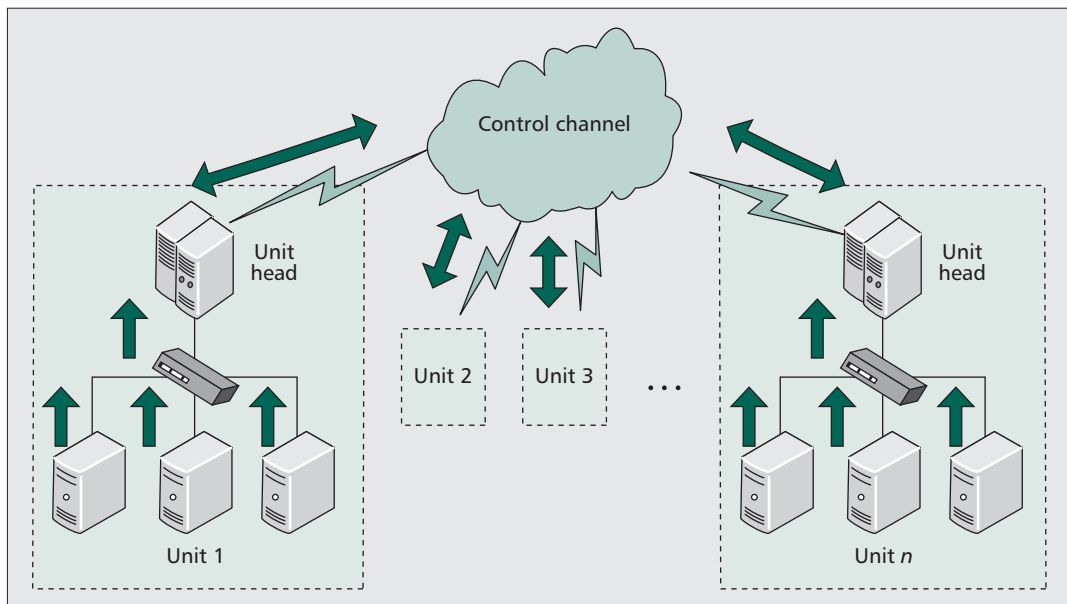


Figure 3. The procedure of collecting traffic demands, in which the arrows stand for the propagation of the information.

First, the amount of traffic to be transmitted via wireless is an important factor. In periodical scheduling, wireless links are set up according to the channel allocation scheme. These links stay active during the whole period whether they actually carry traffic or not. If the traffic of a transmission is low, the corresponding wireless links can finish the transmission in a short time and remain idle after that, causing wireless resource waste. In such a case, it is not reasonable to assign the transmissions to wireless links. Second, the hop distance between the source and the destination of each transmission also matters. Flows with longer wired paths usually incur a larger transmission latency and aggravate the load of higher-layer switches, which results in congestion. Assigning these flows to wireless links is more helpful to improve the global performance.

Taking these factors into consideration, we define the *utility* of a wireless transmission to be the product of the amount of transmitted traffic in a period and a distance factor. Let $u(e)$ be the utility, $\Delta t(e)$ be the transmitted traffic, and $d(e)$ be the hop distance. The utility can be expressed by Eq. 1. Note that $\Delta t(e)$ is determined by the total amount of traffic and the number of wireless links performing the transmission, which is equal to $|C_e|$. We assume that all the wireless links have the same data rate and let Δt_0 denote the maximum amount of traffic that a wireless link can transmit in a period. Thus, $\Delta t(e)$ is the smaller one of $t(e)$ and $|C_e| \Delta t_0$.

$$u(e) = \Delta t(e)d(e) \quad (1)$$

This formalization is inspired by the definition of *footprint* in network redundancy elimination, which is used to measure the amount of resources consumed to forward packets [14]. This definition indicates that the bandwidth resources consumed to transmit a packet is directly proportional to the packet size as well as

the hop distance. We adapt it to our problem by replacing the packet size with the traffic amount to denote the Ethernet bandwidth resources saved by wireless transmissions.

The Optimization Problem for Channel Allocation —

Based on the definition of the wireless transmission utility, we formulate the problem of channel allocation into an optimization problem that maximizes the total utility of all the wireless transmissions. Let $S(e, c)$ be a binary variable denoting the channel allocation scheme, where $S(e, c) = 1$ if and only if $c \in C_e$. This optimization problem is expressed by Eq. 2, in which the first constraint indicates that a feasible channel allocation incurs no interference, while the second constraint ensures that the number of wireless links connecting to a WTU does not exceed the number of radios of the WTU.

$$\max \sum_{e \in \mathcal{E}} u(e) \quad (2)$$

subject to

$$\sum_{e \in \mathcal{E}_v} S(e, c) \leq 1 \quad \forall v \in V, \forall c \in C$$

$$\sum_{c \in C} \sum_{e \in \mathcal{E}_v} S(e, c) \leq r(v) \quad \forall v \in V$$

As all the elements in S are 0-1 integer variables and all the constraints are linear inequalities, Eq. 2 is a 0-1 integer programming whose decision version is one of Karp's 21 NP-complete problems. Therefore, it is impractical to efficiently search for an optimal solution.

A HEURISTIC ALGORITHM

In order to efficiently address the NP-complete problem in Eq. 2, we design a heuristic based on the Hungarian Algorithm [15], which is used to solve the maximum weighted matching problem in graph theory. Our algorithm is motivated by

the similarity between these two problems in spite of the fact that the channel allocation problem allows each edge of the wireless transmission graph to be selected more than once to set up multiple links. To handle the distinction, we employ a dynamic programming approach. In each iteration, the maximum weighted matching of the graph is taken as the candidate wireless links to reach the optimal solution greedily; the weight of each edge is adjusted according to the traffic distribution as well as the limitations of channels and radios; the solutions of all the iterations are combined into the overall channel allocation scheme.

Specifically, we define a two-dimensional *utility matrix* U , in which each entry $u(v_1, v_2)$ denotes the utility of the wireless transmission from v_1 to v_2 . Initially, U is computed based on Eq. 1. The flow chart of the heuristic is shown in Fig. 4. We perform maximum weighted matching on U with the Hungarian Algorithm to select a group of node pairs. For each selected pair (v_1, v_2) , if its corresponding entry in U is not 0, a wireless link with an allocated channel is set up from v_1 to v_2 , and the traffic of the transmission $e = (v_1, v_2)$ is decreased by Δt_0 (if the remaining traffic turns out to be negative, simply set it to 0). Then the entries in U are updated accordingly, which is detailed in the next paragraph. The algorithm terminates if all the entries of U are 0, which indicates that either there is no remaining traffic, or no wireless links can be added due to interference or the limit of radios.

When updating U from iteration to iteration, we have to consider the two constraints elaborated in Eq. 2. First, if there is no idle channel to set up links from v_1 to v_2 , $u(v_1, v_2)$ should be set to 0. Second, if a node v has established $r(v)$ links, all the entries relevant to v (i.e., the entries in either the same row or the same column of $u(v, v)$) should be set to 0.

PERFORMANCE EVALUATION

NETWORK SETTINGS AND SIMULATION SETUP

In order to demonstrate the effectiveness of our wireless DCN, we perform a series of simulations. The experiment platform is a simulator implemented in C++. The network scenario under our consideration is a DCN with a tree topology, in which servers are grouped into 64 racks. We adopt a tree topology simply because it is popular and exemplifies the features of typical current DCN structures.

The simulations employ two types of traffic distributions as the inputs: one is unbalanced, in which 10 percent of the racks generate 90 percent of the total traffic; the other is uniform, where each rack exchanges the same amount of data with all the other racks. These scenarios represent typical traffic distributions in current DCNs.

We evaluate the effectiveness of wireless transmissions by measuring the completion time of the input traffic. A short completion time indicates that the effectiveness achieved by utilizing wireless transmissions is significant. However, it does not make sense to compare the completion time of different traffic distributions

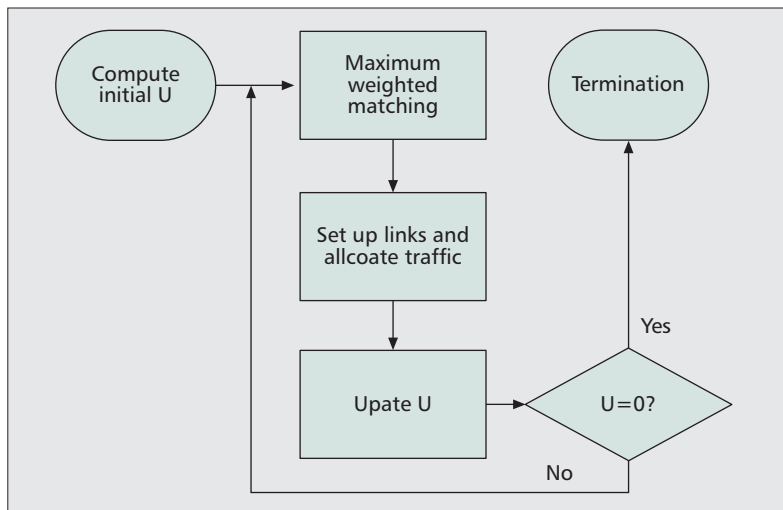


Figure 4. The flow chart of the heuristic algorithm.

directly. Therefore, we take the normalized completion time as the metric, which is defined to be the ratio between the completion time taken by the wireless DCN and that taken by the original wired DCN.

The impacts of several factors are considered in our simulation study, including the number of radios, the total number of wireless links, and the bandwidth of a wireless link. The number of radios determines the maximum number of wireless links that can be added to a WTU. We also limit the total number of wireless links by selectively assigning channels to links in order to investigate how many wireless links are adequate to achieve satisfactory improvement. Besides, the bandwidth of a wireless link, which has an impact on Δt_0 , is also an important factor because we intend to test whether our solution is still efficient if the data rate of wireless links is not as high as that of wired links. Note that we do not take the number of channels into consideration because we adopt directional antennae in our node-exclusive model, in which the limit of radios is usually a stricter constraint than channels.

SIMULATION RESULTS

The simulation results are reported in Fig. 5. In the legend, *unbalanced* and *uniform* stand for the two traffic distributions, respectively; while *high* means that the bandwidth of wireless links is the same as that of wired links, and *low* indicates that the bandwidth of wireless links is only 10 percent that of the wired links.

Number of Radios — In this experiment, we investigate the impact of the number of radios of each WTU. It can be seen from the results that wireless links lead to a great improvement for both types of distributions and the improvement increases as the number of radios grows.

The maximum number of links that can be added to a node is limited by the number of radios. In the unbalanced traffic distribution, when most nodes complete their transmissions, wireless links can only be added to a few nodes with a high volume of traffic. Thus, some radios




Figure 5. Simulation results: a) impact of the number of radios; b) impact of the number of wireless links; c) impact of the bandwidths of wireless links.

would become idle. On the other hand, in the uniform traffic distribution, the utilization ratio of the radios can stay at a higher level for a longer time because there are still a lot of nodes having traffic to transmit. Therefore, the normalized transmission time of unbalanced traffic distribution is higher than that of the uniform traffic distribution.

As the bandwidth of wireless links becomes higher, wireless transmissions play an increasingly important role, and thus the effect has remarkable growth. Notice that the completion time is shortened considerably when the bandwidth of wireless links is the same as that of wired links. This is because the wireless scheduling provides dozens of wireless links equivalent to wired links, which can significantly improve the network capacity.

Number of Wireless Links — We study the impact of the number of wireless links by removing some of the wireless links obtained with our algorithm. For the unbalanced traffic distribution, we observe that the network with only 30–40 additional wireless links achieves almost the same result as that obtained by the network with 100 wireless links. This is because 30–40 wireless links are enough to maintain high utilization of the radios belonging to the nodes with high traffic. Moreover, the addition of a few wireless links to the wired network can make great contributions even if the wireless bandwidth is low. This is due to the fact that these additional links are employed to resolve the bottleneck of the network by our scheduling algorithm. As a result, the completion time of the unbalanced traffic distribution is lower than that of the uniform traffic distribution.

As for the uniform traffic distribution, the performance keeps increasing as the number of wireless links rises because there are always nodes that need extra bandwidth to accelerate transmissions.

Bandwidth of a Wireless Link — Simulations are also performed to investigate the impact of the bandwidth of a wireless link. The results are reported in Fig. 5c, where normalized bandwidth stands for the ratio between the bandwidth of a wireless link and that of a wired link. The trends of both traffic distributions are similar: the completion time drops at a fast speed when the normalized bandwidth increases from 0.1 to 0.5; the decrement is small when the normalized bandwidth becomes higher than 0.8. These phenomena result from the fact that if the bandwidth of wireless links is very low, the transmissions assigned to wireless links would suffer from long delay and become a bottleneck. Solving this problem by increasing the bandwidth can produce a much better result.

CONCLUSION

Wireless networking is proposed as a feasible approach to handle the limitations of Ethernet-based DCN architectures. In this article, the motivations and challenges of wireless data center networking are carefully articulated: a

wireless DCN should meet the performance requirements of data center applications as well as provide efficient scheduling schemes for the large number of servers. To this end, we present a hybrid Ethernet/wireless architecture and elaborate an effective wireless scheduling mechanism. Our simulation results indicate that the congestion of hot servers can be greatly alleviated. Despite the benefits from introducing wireless transmissions to Ethernet-based DCNs, a number of issues arise, including medium access control (MAC) coordination for directional transmissions, cooperation between beamforming and the broadcast nature of wireless communications, hybrid-architecture-based multipath routing, and so on. Further research on these topics is needed to improve the feasibility and efficiency of wireless DCNs.

ACKNOWLEDGMENT

This work is supported by 973 Program of China (2011CB302800, 2009CB320503), NSF of China (61120106008, 60911130511, 60873252), and U.S. NSF grant (CNS-0831852).

REFERENCES

- [1] M. Al-Fares, A. Loukissas, and A. Vahdat, "A Scalable, Commodity Data Center Network Architecture," *ACM SIGCOMM '08*, 2008, pp. 63–74.
- [2] R. Niranjan Mysore et al., "Portland: A Scalable Fault-Tolerant Layer 2 Data Center Network Fabric," *ACM SIGCOMM '09*, 2009, pp. 39–50.
- [3] A. Greenberg et al., "VL2: A Scalable and Flexible Data Center Network," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 39, no. 4, 2009, pp. 51–62.
- [4] C. Guo et al., "DCCell: A Scalable and Fault-Tolerant Network Structure for Data Centers," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 38, no. 4, 2008, pp. 75–86.
- [5] C. Guo et al., "Bcube: A High Performance, Server-Centric Network Architecture for Modular Data Centers," *ACM SIGCOMM '09*, 2009, pp. 63–74.
- [6] J. P. S. Kandula and P. Bahl, "Flyways to de-Congest Data Center Networks," *8th ACM Wksp. Hot Topics in Networks*, 2009.
- [7] T. Benson et al., "Understanding Data Center Traffic Characteristics," *Proc. 1st ACM Wksp. Research on Enterprise Networking*, 2009, pp. 65–72.
- [8] P. Smulders, "Exploiting the 60 GHz Band for Local Wireless Multimedia Access: Prospects and Future Directions," *IEEE Commun. Mag.*, vol. 40, no. 1, 2002, pp. 140–47.
- [9] T. Zwick, T. Beukema, and H. Nam, "Wideband Channel Sounder with Measurements and Model for the 60 GHz Indoor Radio Channel," *IEEE Trans. Vehic. Tech.*, vol. 54, no. 4, 2005, pp. 1266–77.

- [10] P. Smulders, "Statistical Characterization of 60-GHz Indoor Radio Channels," *IEEE Trans. Antennas and Propagation*, vol. 57, no. 10, 2009, pp. 2820–29.
- [11] P. Smulders, H. Yang, and I. Akkermans, "On the Design of Low-Cost 60-GHz Radios for Multigigabit-per-Second Transmission over Short Distances," *IEEE Commun. Mag.*, vol. 45, no. 12, 2007, pp. 44–51.
- [12] Sibeam, <http://www.sibeam.com/>.
- [13] B. Han et al., "Distributed Strategies for Channel Allocation and Scheduling in Software-Defined Radio Networks," *IEEE INFOCOM '09*, 2009, pp. 1521–29.
- [14] A. Anand et al., "Packet Caches on Routers: the Implications of Universal Redundant Traffic Elimination," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 38, no. 4, 2008, pp. 219–30.
- [15] H. Kuhn, "The Hungarian Method for the Assignment Problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1–2, 1955, pp. 83–97.

BIOGRAPHIES

YONG CUI (cuiyong@tsinghua.edu.cn) received his Ph.D. degree in computer science from Tsinghua University. He is currently an associate professor in the Department of Computer Science and Technology at Tsinghua University. He serves as co-chair of the IETF Software Working Group. He has published more than 100 papers in journals and conferences including *IEEE/ACM Transactions on Networking*, *IEEE Internet Computing*, *IEEE Transactions on Parallel and Distributed Systems*, *IEEE Transactions on Multimedia*, and *IEEE INFOCOM*. His major research interests include mobile wireless networks, computer network architecture, and cloud computing.

HONGYI WANG (wanghongyi09@mails.tsinghua.edu.cn) received his B.S. degree in computer science from Tsinghua University in 2009. He is pursuing his Master's degree in the Department of Computer Science and Technology at Tsinghua University. His research interests include data center networking, wireless networking, and mobile computing.

XIUZHEN CHENG (cheng@gwu.edu) received her Ph.D. degree in computer science from University of Minnesota — Twin Cities in 2002. Currently she is an associate professor at the Department of Computer Science, The George Washington University. She has served on the editorial boards of several technical journals and the technical program committees of various professional conferences/workshops. She has also chaired several international conferences. Her research interests include wireless networking, wireless and mobile security, and algorithm design and analysis.

BIAO CHEN [M] (bchen@umac.mo) received his B.S. in computer science from Fudan University in China, and M.S. in mathematics and Ph.D. in computer science from Texas A&M University. After graduation, he joined the Department of Computer Science at the University of Texas at Dallas as an assistant professor. Currently, he is a visiting professor in the Department of Computer and Information Science of the University of Macau. He is a member of Sigma Xi and ACM. His research interests include distributed systems, networking, and security.

These phenomena result from the fact that if the bandwidth of wireless links is very low, the transmissions assigned to wireless links would suffer from a long delay and become bottleneck. Solving this problem by increasing the bandwidth can produce a much better result.