

PRIMAL: PProfit Maximization Avatar pLacement for Mobile Edge Computing

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

Citation:

X. Sun and N. Ansari, "PRIMAL: PProfit Maximization Avatar pLacement for mobile edge computing," 2016 IEEE International Conference on Communications (ICC), Kuala Lumpur, 2016, pp. 1-6, DOI: 10.1109/ICC.2016.7511131.

URL:

<https://ieeexplore.ieee.org/document/7511131/>

PRIMAL: PROfit Maximization Avatar pLacement for Mobile Edge Computing

Xiang Sun, *Student Member, IEEE*, and Nirwan Ansari, *Fellow, IEEE*

Advanced Networking Lab., Helen and John C. Hartmann Department of Electrical & Computer Engineering

New Jersey Institute of Technology, Newark, NJ 07102, USA.

Emails: {xs47, nirwan.ansari}@njit.edu.

Abstract—We propose a cloudlet network architecture to bring the computing resources from the centralized cloud to the edge. Thus, each User Equipment (UE) can communicate with its Avatar, a software clone located in a cloudlet, and can thus lower the end-to-end (E2E) delay. However, UEs are moving over time, and so the low E2E delay may not be maintained if UEs' Avatars stay in their original cloudlets. Thus, live Avatar migration (i.e., migrating a UE's Avatar to a suitable cloudlet based on the UE's location) is enabled to maintain the low E2E delay between each UE and its Avatar. On the other hand, the migration itself incurs extra overheads in terms of resources of the Avatar, which compromise the performance of applications running in the Avatar. By considering the gain (i.e., the E2E delay reduction) and the cost (i.e., the migration overheads) of the live Avatar migration, we propose a PROfit Maximization Avatar pLacement (PRIMAL) strategy for the cloudlet network in order to optimize the tradeoff between the migration gain and the migration cost by selectively migrating the Avatars to their optimal locations. Simulation results demonstrate that as compared to the other two strategies (i.e., Follow Me Avatar and Static), PRIMAL maximizes the profit in terms of maintaining the low average E2E delay between UEs and their Avatars and minimizing the migration cost simultaneously.

I. INTRODUCTION

Recent mobile applications, such as augmented reality, image processing and speech recognition, become resource intensive and drain User Equipments' (UEs') batteries very quickly. Mobile Cloud Computing (MCC) [1], [2] has been proposed to offload applications' workloads from UEs to the cloud in order to not only reduce energy consumption of UEs but also accelerate the execution time of the applications. MCC reduces the UE's computational cost at the expense of the communications cost, i.e., the UE frequently interacts with the cloud by offloading its application workloads. Thus, it is not efficient to perform the application offloading if the End-to-End (E2E) delay between a UE and the cloud is unbearable. The cloudlet architecture is introduced to reduce the E2E delay, i.e., the computing resources are moved from the remote cloud to the local cloudlet, which is a tiny version of a data center residing close to UEs, so that UEs can access the computing resources with the low E2E delay [2].

To reap benefits of the cloudlet, we propose the cloudlet network architecture, as shown in Fig. 1, in order to provide ubiquitous computing resources to UEs and at the same time maintain the low E2E delay. Since the existing LTE network infrastructure can provide seamless connection between a UE and a base station (BS), each BS is connected to a cloudlet




Fig. 1. The cloudlet network architecture.

via high speed fibers so that UEs can utilize computing resources in the cloudlets with one wireless hop delay. Moreover, each UE subscribes one Avatar, a high performance Virtual Machine (VM) in the cloudlet, which provides extra computing resources and storage space. Avatars are software clones of their UEs and always available to UEs when UEs are moving from one coverage area to another. Assigning a specific Avatar to each UE in the cloudlet provides hardware isolation by securely running each UE's application workloads on a shared physical hardware. On the top of the cloudlets, Software Defined Network (SDN) based cellular core network has been proposed in the cloudlet network architecture to provide efficient and flexible communications paths between Avatars in different cloudlets as well as between UEs in different BSs [3], [4]. Moreover, every UE and its Avatar in the cloudlet can communicate with public data centers (e.g., Amazon EC2) and Storage Area Networks (SANs) via the Internet in order to provision scalability, i.e., if cloudlets are not available for UEs because of the capacity limitation, UEs' Avatars can be migrated to the remote data centers to continue serving their UEs.

The locations of cloudlets depend on the UE density, i.e., the BS in the hotspot area can own a specific cloudlet (e.g., BS1 connects to a specific cloudlet, say, Cloudlet A in Fig. 1), or the BSs in the rural or suburban area can share the same cloudlet. Thus, the cloudlet can be placed among the BSs (e.g., cloudlet B is deployed between BS2 and BS3 in Fig. 1) or the cloudlet is directly connected to the switch at the edge of the SDN based cellular core (e.g., cloudlet C connects to the

edge switch so that BS4 and BS5 can share the computing and storage resources of cloudlet C .

The cloudlet network architecture not only helps UEs offload their application workloads to their Avatars with lower latency but also facilitates real time big mobile data analysis. Smart UEs, embedded with a rich set of sensors, become a data stream generator producing their users' information over time. Analyzing these massive amount of mobile data is not only extremely valuable for market applications, but also potentially benefits the society as a whole [5]. Traditionally, these big mobile data are analyzed within a data center [6], [7] by utilizing the distributed computing framework. However, transmitting the big mobile data from UEs to the data center through the Internet suffers from the long latency and increases the traffic load of the network. Meanwhile, the potential value of the mobile data is decreasing as time passes by. Thus, rather than bringing the mobile data to the computing resources, the cloudlet network architecture is proposed to bring the computing resources to the mobile data. In other words, each Avatar locally collects, filters, classifies or even analyzes the raw data stream of its user so that the volume of the mobile data, which need to be transmitted to the remote data center for further analysis, can be reduced substantially or eliminated and the communications latency can be reduced as well.

The rest of the paper is organized as follows. In Sec. II, we propose the live Avatar migration among cloudlets to maintain the low E2E delay between UEs and their Avatars when UEs are moving over time. We design the live Avatar migration gain and cost models to calculate the gain and cost of each Avatar migration, respectively. In Sec. III, we formulate the novel Avatar placement strategy, referred to as P_{RO}fit Maximization Avatar pLacement (*PRIMAL*), to maximize the migration profit in terms of optimizing the tradeoff between the migration gain and the migration cost. In Sec. IV, we demonstrate the performance of the proposed strategy via extensive simulations. The conclusion is presented in Sec. V.

II. LIVE AVATAR MIGRATION

UEs frequently communicate with their Avatars by transmitting their mobile data over time. Thus, deploying the Avatar close to its UE will essentially reduce the E2E delay between the UE and its Avatar, and thus facilitates big mobile data analysis as well as MCC applications. However, UEs are moving over time, and so the E2E delay between a UE and its Avatar might become unbearable. Therefore, it is necessary to optimize each Avatar's placement. By taking advantages of the live VM migration in a data center, Avatars can also be migrated among cloudlets over the SDN based cellular core network to alter their locations, and thus the E2E delay between a UE and its Avatar can be reduced by optimizing the Avatars' locations. Nevertheless, Avatar migrations are expensive operations because they incur additional overheads [8], [9], i.e., Avatar migrations consume extra resources (e.g., CPU, memory, network, disk I/O resources), which affect the performance of applications running in the Avatars. While migrating the Avatar close to its UE potentially improves the

E2E delay, it may introduce humongous migration overheads. In order to measure the profit of the migration, it is important to consider the gain and the cost of the Avatar migration simultaneously.

A. Live Avatar Migration Gain Model

Different from the traditional live VM migration in a data center, the benefit of the live Avatar migration is to reduce the E2E delay between a UE and its Avatar, which comprises three parts: first, T^{access} , i.e., the E2E delay between a UE and its BS (to which the UE is associated with); second, T^{core} , i.e., the E2E delay between the UE's BS and the UE's cloudlet (in which the UE's Avatar is located); third, $T^{cloudlet}$, i.e., the E2E delay within the cloudlet. Since changing the placement of UEs' Avatars does not significantly affect the values of T^{access} and $T^{cloudlet}$, we consider the gain of live Avatar migration as the reduction of T^{core} , which is the most important parameter affecting the E2E delay between a UE and its Avatar. In other words, if a UE's Avatar is migrated to the cloudlet which has a lower E2E delay to the UE's BS, then the gain of the live Avatar migration is defined as the reduction of T^{core} .

Denote \mathcal{I} as the set of UEs/Avatars (note that one UE is associated with one specific Avatar, and thus we equate the set of UEs to the set of Avatars), and i is used to index the UEs and their corresponding Avatars. Denote \mathcal{J} and \mathcal{K} as the set of cloudlets and BSs in the network, respectively, and j and k are used to index the cloudlets and BSs, respectively. Denote $x_{i,j}$ as the binary variable to indicate whether Avatar i is located in cloudlet j (i.e., $x_{i,j} = 1$) or not. Meanwhile, $y_{i,k}$ is used to indicate whether UE i is associated with BS k (i.e., $y_{i,k} = 1$) or not. By taking the advantage of the SDN network, the E2E delay between the j th ($j \in \mathcal{J}$) cloudlet and the k th ($k \in \mathcal{K}$) BS in the cloudlet network, denoted as $d_{j,k}$, can be measured by the SDN controller in each time slot [10], [11]. Thus, the E2E delay between UE i 's BS and UE i 's cloudlet can be derived as follows:

$$T_i^{core} = \sum_{j=1}^{|\mathcal{J}|} \sum_{k=1}^{|\mathcal{K}|} x_{i,j} y_{i,k} d_{j,k}. \quad (1)$$

Assuming that reducing one unit of the E2E delay increases one unit of the gain, then given UE i 's location (denoted as $y_{i,k}^{t+1}$) in the next time slot, the migration gain model is defined as the amount of the E2E delay reduction achieved by the case that UE i 's Avatar migrates to the location of $x_{i,j}^{t+1}$ in the next time slot as compared to the case that UE i 's Avatar stays in the current location (i.e., $x_{i,j}^t$) in the next time slot, i.e.,

$$r_i = \sum_{j=1}^{|\mathcal{J}|} \sum_{k=1}^{|\mathcal{K}|} (x_{i,j}^t - x_{i,j}^{t+1}) y_{i,k}^{t+1} d_{j,k}. \quad (2)$$

B. Live Avatar Migration Cost Model

Although live Avatar migration can facilitate the communications between a UE and its Avatar, the available resources for running the Avatar's applications diminish, thus compromising the QoS for using the Avatar. Suppose the migration overheads are fixed during the migration process (i.e., the migration consumes the same amount of resources in each time slot

during the process); if the migration consumes less time, it generates fewer overheads to affect the applications in the Avatar. We build the live Avatar migration cost model as follows:

$$c_i = \kappa_i T_i^{mig} \sum_{j=1}^{|\mathcal{J}|} \frac{1}{2} (x_{i,j}^t - x_{i,j}^{t+1})^2, \quad (3)$$

where c_i is the cost for migrating Avatar i , T_i^{mig} is the total migration time of Avatar i , κ_i is the cost coefficient that maps the migration time to the cost, and $\sum_{j=1}^{|\mathcal{J}|} \frac{1}{2} (x_{i,j}^t - x_{i,j}^{t+1})^2$ indicates whether Avatar i is migrated to another cloudlet (i.e., the summation equals to 1) or not (i.e., the summation equals to 0).




Fig. 2. The pre-copy live migration procedure.

1) *Total Migration Time*: We apply the pre-copy live migration technique [12] for migrating Avatars among cloudlets. As shown in Fig. 2., there are two phases during pre-copy live migration, i.e., the pre-copy phase and stop© phase [12], [13]. In the pre-copy phase, the whole memory of the source Avatar is transmitted to the destination Avatar in the initial round. For the rest of each round, source Avatar sends the dirty pages, which are generated from the previous round to the destination. Until the number of generated dirty pages is no larger than a predefined threshold, the migration proceeds to the stop-and-copy phase, i.e., the source Avatar stops serving its UE, transmits the rest of the dirty memory pages and informs the destination Avatar to resume services [9].

Lemma 1. *If the bandwidth provisioning for performing migration is constant, given the amount of Avatar i 's memory (denoted as M_i) and Avatar i 's average memory page dirtying rate (i.e., the average number of dirty memory pages generated in each time slot, denoted as D_i) during the migration, the time required for executing Avatar i 's migration is:*

$$T_i^{mig} = \frac{M_i}{R - D_i} \left[1 - \left(\frac{D_i}{R} \right)^{\left\lceil \log_{D_i/R} \left(\frac{M_i^{th}}{M_i} \right) + 1 \right\rceil + 1} \right], \quad (4)$$

where R is the bandwidth provisioning in terms of the data rate for doing migration ($R > D_i$) and M_i^{th} is the threshold of dirty pages generated.

Proof: Suppose there are N_i number of rounds during the migration process and the SDN network provider can guarantee a constant bandwidth in terms of a fixed bit rate for doing migration, the amount of time consumed in the current round n depends on the amount of dirty memory generated in the previous round $n - 1$. Thus, we have [14]:

$$T_{i,n}^{mig} = \frac{D_i}{R} T_{i,n-1}^{mig} = \left(\frac{D_i}{R} \right)^n T_{i,0}^{mig}, \quad 1 \leq n \leq N_i, \quad (5)$$

where $T_{i,n}^{mig}$ is the time consumption of round n during Avatar i 's migration, and $T_{i,0}^{mig}$ is the time consumption of the initial round (i.e., round 0) during which the whole memory of the source Avatar is transmitted to the destination, i.e., $T_{i,0}^{mig} = \frac{M_i}{R}$. Thus, the time consumption of round n is:

$$T_{i,n}^{mig} = \left(\frac{D_i}{R} \right)^n \frac{M_i}{R}, \quad 0 \leq n \leq N_i. \quad (6)$$

The total time consumption of the migration is the sum of the time consumption in each round, i.e.,

$$T_i^{mig} = \sum_{n=0}^{N_i} T_{i,n}^{mig} = \frac{M_i}{R - D_i} \left[1 - \left(\frac{D_i}{R} \right)^{N_i + 1} \right]. \quad (7)$$

As mentioned before, once the number of the generated dirty pages are no larger than a predefined threshold (i.e., M_i^{th}) in the previous round, then the source Avatar would stop serving its UE and transmit the rest of the dirty pages to the destination Avatar in the last round, i.e., $T_{i,N_i-1}^{mig} R \leq M_i^{th}$. Based on Eq. 6, we have $T_{i,N_i-1}^{mig} = \left(\frac{D_i}{R} \right)^{N_i-1} \frac{M_i}{R}$, and so we can derive that if $R > D_i$, $N_i \leq \log_{D_i/R} \left(\frac{M_i^{th}}{M_i} \right) + 1$ (note that if $R \leq D_i$, then $N_i \rightarrow +\infty$). Since the number of the migration rounds should be an integer value, we have:

$$N_i = \left\lceil \log_{D_i/R} \left(\frac{M_i^{th}}{M_i} \right) + 1 \right\rceil, \quad R > D_i. \quad (8)$$

By substituting Eq. 8 into Eq. 7, we have Eq. 4 and thus have proved Lemma 1. ■

2) *Cost Coefficient κ_i* : The cost coefficient κ_i in Eq. 4 may also vary among Avatars because even if different Avatars generate the same migration overheads in terms of the same migration time, the performance degradation of applications running in different Avatars are also different. The reason is that the live Avatar migration itself can be considered as an I/O intensive application, and so the Avatar migration would have more negative effect on the performance of applications which have higher I/O footprints as compared to the pure CPU intensive applications [8], i.e., the Avatars running higher I/O applications have higher cost coefficient than the Avatars running lower I/O applications. Based on the above observation, we model the value of κ_i to be proportional to the weighted sum of utilization of different resources [8]:

$$\kappa_i = \alpha \left(w^{net} u_i^{net} + w^{disk} u_i^{disk} + w^{mem} u_i^{mem} + w^{cpu} u_i^{cpu} \right), \quad (9)$$

where u_i^{net} , u_i^{mem} , u_i^{disk} and u_i^{cpu} denote the bandwidth, memory, disk I/O and CPU resource utilization of Avatar i , respectively; w^{net} , w^{mem} , w^{disk} and w^{cpu} are the migration impact factor of the bandwidth, memory, disk I/O and CPU resource utilization, respectively, indicating the degree of impact of different resources (note that the values of w^{net} , w^{mem} , w^{disk} and w^{cpu} can be derived through experiments [8]); α is the penalty coefficient that maps the weighted sum of utilization of resources to the cost coefficient. α is an important parameter in the system. Increasing the value of α would increase the ratio of the migration cost to migration gain, and discourage Avatars from doing live migrations. Consequently,

the E2E delay would increase if Avatars are not incentivized to do live migration. Thus, α is a parameter to adjust the tradeoff between the E2E delay and the cost for doing live Avatar migration, and can be chosen via experiments by testing users' QoE for utilizing their Avatars. Also, altering the value of α can adjust the traffic in the SDN cellular core, i.e., increasing the value of α would reduce the traffic generated by the migrations and mitigate the traffic load in the SDN cellular core consequently.

III. PROFIT MAXIMIZATION AVATAR PLACEMENT (PRIMAL)

In order to increase the gain by facilitating the communications between a UE and its Avatar, the UE's Avatar can be placed in the cloudlet which has a lower E2E delay to the UE's BS. However, changing the placement of the UE's Avatar involves live Avatar migration which would degrade the performance of applications running in the Avatar. Thus, we need to design an optimal Avatar placement strategy to optimize the tradeoff between the migration gain and the migration cost by estimating whether it is worth to do the live Avatar migration or not. Denote f_i as the profit of migrating Avatar i , i.e., migration gain minus migration cost:

$$\begin{aligned} f_i &= r_i - c_i \\ &= \sum_{j=1}^{|\mathcal{J}|} \left\{ -0.5\kappa_i T_i^{mig} (x_{i,j}^{t+1})^2 + \left(\kappa_i T_i^{mig} x_{i,j}^t - \sum_{k=1}^{|\mathcal{K}|} y_{i,k}^{t+1} d_{j,k} \right) x_{i,j}^{t+1} \right\} \\ &+ \sum_{j=1}^{|\mathcal{J}|} \left\{ x_{i,j}^t \sum_{k=1}^{|\mathcal{K}|} y_{i,k}^{t+1} d_{j,k} - 0.5\kappa_i T_i^{mig} (x_{i,j}^t)^2 \right\}. \end{aligned} \quad (10)$$

We assume that the hardware configuration of each Avatar is the same, and each cloudlet can only host a limit number of Avatars, denoted as s_j ($j \in \mathcal{J}$). Then, we formulate PRIMAL as follows:

$$\arg \max_{x_{i,j}^{t+1}} \sum_{i=1}^{|\mathcal{I}|} f_i \quad (11)$$

$$s.t. \quad \forall i \in \mathcal{I}, \quad \sum_{j=1}^{|\mathcal{J}|} x_{i,j}^{t+1} = 1, \quad (12)$$

$$\forall j \in \mathcal{J}, \quad \sum_{i=1}^{|\mathcal{I}|} x_{i,j}^{t+1} \leq s_j, \quad (13)$$

$$\forall i \in \mathcal{I}, \forall j \in \mathcal{J}, \quad x_{i,j}^{t+1} \in \{0, 1\}, \quad (14)$$

where the objective is to maximize the total profit of Avatar live migrations. The first constraint imposes that every UE's Avatar should be allocated in only one cloudlet. The second constraint means that the total number of Avatars assigned to the cloudlet cannot exceed the cloudlet's capacity.

Theorem 1. *The PRIMAL problem is NP-hard.*

Proof: The proof of the theorem is shown in our technical report [15]. ■

We use the Mixed-Integer Quadratic Programming tool in the CPLEX solver to find the heuristic solution of PRIMAL.

IV. SIMULATION RESULTS

We simulate the proposed PRIMAL strategy in the cloudlet network. For comparisons, we select other two live Avatar migration decision strategies, i.e., the Follow me AvataR (FAR) strategy and the Static strategy. The idea of the FAR strategy is to minimize the E2E delay between an Avatar and its UE by assigning the Avatar to the available cloudlet (i.e., the cloudlet has enough space to hold the Avatar), which yields the lowest E2E delay to its UE's BS [4]. The Static strategy is to avoid the migration cost, i.e., the locations of Avatars do not change over time after they are initially deployed.

We set up a network with the topology that includes 25 cloudlet-eNB combinations (5×5) in a square area of 100 km^2 . Each cloudlet connects one eNB and the coverage area of each eNB is a square area of 4 km^2 . There are 1000 UEs, each associated with one Avatar, in the network. The UE's mobility model follows the random way point model and each cloudlet has the same capacity of 50 Avatars.

The resource capacity of each Avatar is homogeneous; each Avatar is configured with 2-core CPU, 4GB memory, and 500 Mbps bandwidth. The Google cluster data trace [17] is applied to emulate the CPU, memory and disk I/O utilization of each Avatar (we select the machines with CPU and memory capacity of 0.5 (normalized) in the Google cluster data trace, and calculate their CPU, memory and disk I/O utilization in each time slot; then, the resources of Avatars are emulated to be the same as those of the machines). Since the Google cluster data trace does not publish the bandwidth resource utilization and memory page dirtying rate traces of the machines, we emulate the bandwidth demand of each Avatar as a stochastic process which follows a normal distribution $N(\mu_i, \sigma_i^2)$ [18], [19]. Note that the value of μ_i and σ_i^2 are different among different Avatars, and thus we randomly select $\mu_i \in [0, 350 \text{ Mbps}]$ and $\sigma_i^2 \in [0, 100 \text{ Mbps}]$ for each Avatar. Furthermore, each Avatar's memory page dirtying rate depends on different types of applications running in the Avatar, i.e., some memory-intensive applications (e.g., in-memory data analytics) may generate more dirty memory pages as compared to the CPU-intensive and network I/O-intensive applications. In the simulation, the memory page dirtying rate is randomly chosen between 0 and 10K pages per time slot. Moreover, the values of w^{net} , w^{mem} , w^{disk} and w^{cpu} are selected to be 0.8, 0.6, 0.4, and 0.1, respectively [8]. The data rate for doing migration is 200 Mbps.

First, we set up the penalty coefficient $\alpha = 5$ and run the simulation. Fig. 3 shows the profit trace by applying three different Avatar placement strategies. PRIMAL achieves the highest profit as compared to the other two strategies indicating that PRIMAL can choose the valuable migrations (i.e., $f_i > 0$) to maximize the profit. In order to demonstrate the benefit for maximizing the profit, we further test the average Round Trip Time (RTT) in terms of the E2E delay between UEs and their Avatars, as shown in Fig. 4, PRIMAL and FAR yield the similar average RTT, which is much lower than that of Static because Static does not dynamically adjust

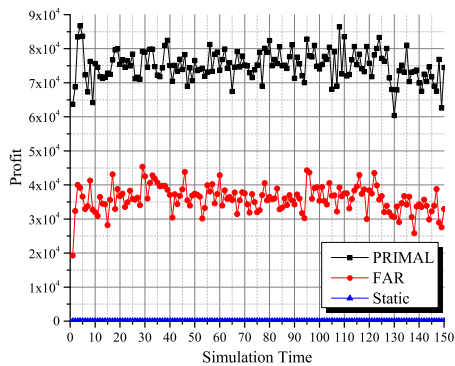


Fig. 3. Profit generated by different strategies.

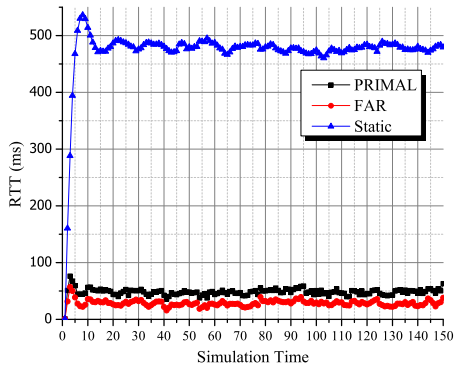


Fig. 4. The average RTT between UEs and their Avatars during the simulation.




Fig. 5. Migration cost incurred by different strategies.




Fig. 6. The average profit differences among different strategies by choosing different values of α .




Fig. 7. The average RTT incurred by PRIMAL over different values of α .




Fig. 8. The average number of migrations and the average migration time incurred by PRIMAL over different values of α .

Avatars' placement even if the E2E delays between UEs and their Avatars are unbearable. On the other hand, as shown in Fig. 5, the migration cost of PRIMAL is much less than that of FAR, and thus PRIMAL achieves higher profit than FAR. We further test the average number of migrations and the average migration time as shown in Table I. Obviously, PRIMAL reduces the average number of the migrations as well as the average migration time, indicating that PRIMAL avoids some migrations with long migration time. In other words, PRIMAL selectively migrates Avatars that improve their RTT time (in terms of the migration gain) significantly but consume less migration time (in terms of the migration cost).

TABLE I
SIMULATION RESULTS

Strategies	Average number of migrations	Average migration time
PRIMAL	472.2 migrations/slot	59.7 seconds
FAR	606.1 migrations/slot	67.4 seconds

Second, we try to analyze the performance of PRIMAL by choosing different values of α . Fig. 6 shows the differences of the average profits between PRIMAL and FAR as well as between PRIMAL and Static during the simulation by choosing different values of α . When $\alpha = 0$, PRIMAL performs exactly the same as FAR (i.e., the average profit difference is zero) as both try to minimize the RTT only.

However, as the value of α increases, the migration cost increases and PRIMAL enables less migrations to maximize the profit, thus increasing the average profit gap between PRIMAL and FAR. On the other hand, as the value of α increases, the average profit gap between PRIMAL and Static is decreasing since more Avatars remain static to avoid the migration cost. We believe the performance of PRIMAL and Static is the same as $\alpha \rightarrow +\infty$. We further test the trend of the average RTT, the average number of migrations and the average migration time for running PRIMAL as the value of α varies. As shown in Fig. 7 and Fig. 8, when the value of α is small, PRIMAL triggers more migrations to improve the RTT even if the migrations consume more time. However, as α increases, the cost for doing migration increases and PRIMAL avoids more worthless migrations (those that improve the RTT a little bit at the expense of longer migration time). Therefore, there are tradeoffs between the average RTT and the average number of migrations, and between the average RTT and the average migration time, i.e., in order to reduce the average RTT, more migrations are triggered and more migration time is consumed. Changing the value of α can adjust these tradeoffs.

V. RELATED WORKS

Various VM placement strategies have been proposed for the resource management in data centers. Wood *et al.* [20] proposed *Sandpiper* to detect hotspots and mitigate them by

migrating the VMs to lightly loaded servers. Piao and Yan [21] considered that VMs and their data may be located at different physical servers in the cloud, and thus proposed a virtual machine placement strategy to minimize the data access time, i.e., placing VMs close to their data so that the access time is minimized. Shrivastava *et al.* [22] proposed the VM placement strategy to place the dependent VMs (i.e., the VMs with heavy interaction among them) close to each other so that the network traffic can be reduced. Rather than only considering the gain of VM placement by migrating VMs to suitable servers, many studies argued that the cost of VM migration involved in the VM placement cannot be neglected when the resource management is applied. Verma *et al.* [23] proposed the VM placement strategy to minimize the total power while taking the migration cost into account. The migration cost is depicted as the throughput of the migration. Hossain *et al.* [24] also tried to minimize the total energy consumption by utilizing the VM placement, but they modeled the migration cost as the migration energy consumption at the destination and source servers.

Our previous work [4] tried to maximize the green energy utilization of the cloudlets in the network by utilizing the Avatar migrations to adjust the energy demands among cloudlets. In this paper, we try to optimize the Avatars' placement by maximizing the profit of Avatar migrations in terms of optimizing the tradeoff between the migration gain and the migration cost. To the best of our knowledge, none of the previous works considered reducing the E2E delay between a user and its VM as the gain of VM migrations, and reducing the E2E is very important in the proposed cloudlet network since it can substantially facilitate communications between a UE and its Avatar in meeting the QoS of MCC applications and provisioning big mobile data analysis.

VI. CONCLUSION

In this paper, we have proposed the cloudlet network architecture to reduce the E2E delay between a UE and its Avatar so as to meet the QoS of MCC applications and to provision big mobile data analysis. However, UEs are moving in the network, and so the E2E may become worse. The live Avatar migration is triggered to adjust the location of the UE's Avatar. However, the migration process consumes extra resources of the Avatar that may degrade the performance of applications running in the Avatar. Therefore, we have proposed PRIMAL to maximize the profit of the Avatar migration in terms of optimizing the tradeoff between the gain and the cost of the Avatar migration. We have also demonstrated that PRIMAL achieves the highest profit as compared to the other two Avatar placement strategies, i.e., FAR and Static.

REFERENCES

[1] X. Sun and N. Ansari, "Energy-optimized bandwidth allocation strategy for mobile cloud computing in LTE networks," in *IEEE Wireless Communications and Networking Conference (WCNC)*, New Orleans, LA, Mar. 9-12, 2015, pp. 2120-2125.

[2] S. Mahadev, P. Bahl, R. Caceres, and N. Davies, "The Case for VM-Based Cloudlets in Mobile Computing," *IEEE Pervasive Computing*, vol. 8, no. 4, pp. 14-23, Oct.-Dec. 2009.

[3] A. Basta, *et al.*, "A Virtual SDN-Enabled LTE EPC Architecture: A Case Study for S-/P-Gateways Functions," in *2013 IEEE SDN for Future Networks and Services*, Trento, Italy, Nov. 11-13, 2013, pp. 1-7.

[4] X. Sun, N. Ansari, and Q. Fan, "Green Energy Aware Avatar Migration Strategy in Green Cloudlet Networks," in *7th IEEE International Conference on Cloud Computing Technology and Science (Cloudcom)*, Vancouver, Canada, Nov. 30-Dec. 3, 2015, accepted.

[5] M. Musolesi, "Big Mobile Data Mining: Good or Evil?," *IEEE Internet Computing*, vol. 18, no. 1, pp. 78-81, Jan.-Feb. 2014.

[6] Y. Zhang and N. Ansari, "On Architecture Design, Congestion Notification, TCP Incast and Power Consumption in Data Centers," *IEEE Communications Surveys and Tutorials*, vol. 15, no. 1, pp. 39-64, 2013.

[7] J. Liu, F. Liu, and N. Ansari, "Monitoring and Analyzing Big Traffic Data of a Large Scale Cellular Network with Hadoop," *IEEE Network*, vol. 28, no. 4, pp. 32-39, July/Aug. 2014.

[8] A. Anand, J.Lakshmi, and S.K. Nandy, "Virtual Machine Placement Optimization Supporting Performance SLAs," in *2013 IEEE 5th International Conference on Cloud Computing Technology and Science (CloudCom)*, Bristol, UK, Dec. 2-5, 2013, vol. 1, pp. 298-305.

[9] L. Haikun, H. Jin, C.Z. Xu, and X. Liao, "Performance and energy modeling for live migration of virtual machines," *Cluster computing*, vol. 16, no. 2 pp. 249-264, 2013.

[10] N.L.M. Adrichem, C. Doerr, and F. Kuipers, "Opennetmon: Network monitoring in openflow software-defined networks," in *2014 IEEE Network Operations and Management Symposium (NOMS)*, Krakow, Poland, May. 05-09, pp. 1-8.

[11] C. Yu, *et al.*, "Software-Defined Latency Monitoring in Data Center Networks," in *Passive and Active Measurement*, vol.8995, pp. 360-372, Mar., 2015.

[12] C. Clark, *et al.*, "Live migration of virtual machines," in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation*, Boston, MA, May. 2-4, 2005, vol. 2, pp. 273-286.

[13] A. Strunk, "A lightweight model for estimating energy cost of live migration of virtual machines," in *2013 IEEE Sixth International Conference on Cloud Computing (CLOUD)*, Santa Clara Marriott, CA, Jun. 27-Jul. 2, 2013, pp. 510-517.

[14] V. Mann, *et al.*, "Remedy: Network-aware steady state VM management for data centers," in *NETWORKING*, Prague, Czech Republic, May. 21-24, 2012, pp. 190-204.

[15] X. Sun and N. Ansari, "PRIMAL: PROfit Maximization Avatar pLacement for Mobile Edge Computing," NJIT Advanced Networking Lab., Tech. Rep. TR-ANL-2015-007; also archived in Computing Research Repository (CoRR), arXiv:1510.04340, 2015.

[16] E.M. Loiola, *et al.*, "A survey for the quadratic assignment problem," *European Journal of Operational Research*, vol. 176, no. 2, pp.657-690, 2007.

[17] Google Cluster Data. [Online]. Available: https://github.com/google/cluster-data/blob/master/ClusterData2011_2.md.

[18] H. Jin, D. Pan, J. Xu, and N. Pissinou, "Efficient VM placement with multiple deterministic and stochastic resources in data centers," in *2012 IEEE Global Communications Conference (GLOBECOM)*, Anaheim, CA, Dec. 3-7, 2012, pp. 2505-2510.

[19] M. Wang, X. Meng, and L. Zhang, "Consolidating virtual machines with dynamic bandwidth demand in data centers," in *2011 Proceedings IEEE INFOCOM*, ShangHai, China, Apr. 10-15, pp. 71-75.

[20] T. Wood, P. Shenoy, A. Venkataramani, and M. Yousif, "Sandpiper: Black-box and gray-box resource management for virtual machines," *Computer Networks*, vol. 53, no. 17, pp. 2923-2938, 2009.

[21] J.T. Piao and J. Yan, "A network-aware virtual machine placement and migration approach in cloud computing," in *IEEE 9th International Conference on Grid and Cooperative Computing (GCC)*, Nanjing, China, Nov. 1-5, 2010, pp. 87-92.

[22] V. Shrivastava, *et al.*, "Application-aware virtual machine migration in data centers," in *2011 Proceedings IEEE INFOCOM*, Shanghai, China, Apr. 10-15, 2011, pp. 66-70.

[23] A. Verma, P. Ahuja, and A. Neogi, "Mapper: power and migration cost aware application placement in virtualized systems," *Middleware*, vol. 5346, pp. 243-264, 2008.

[24] M.M. Hossain, J.C. Huang, and H.H.S. Lee, "Migration energy-aware workload consolidation in enterprise clouds," in *2012 IEEE 4th International Conference on Cloud Computing Technology and Science (CloudCom)*, Taipei, Taiwan, Dec. 3-6, 2012, pp. 405-410.