

Migration towards Cloud-Assisted Live Media Streaming

Feng Wang, *Member, IEEE*, Jiangchuan Liu, *Senior Member, IEEE*, Minghua Chen, *Member, IEEE*, and Haiyang Wang, *Member, IEEE*

Abstract—Live media streaming has become one of the most popular applications over the Internet. We have witnessed the successful deployment of commercial systems with CDN- or peer-to-peer based engines. While each being effective in certain aspects, having an all-round scalable, reliable, responsive and cost-effective solution remains an illusive goal. Moreover, today’s live streaming services have become highly globalized, with subscribers from all over the world. Such a globalization makes user behaviors and demands even more diverse and dynamic, further challenging state-of-the-art system designs.

The emergence of *cloud computing* however sheds new lights into this dilemma. Leveraging the elastic resource provisioning from cloud, we present CALMS (Cloud-Assisted Live Media Streaming), a generic framework that facilitates a migration to the cloud. CALMS adaptively leases and adjusts cloud server resources in a fine granularity to accommodate temporal and spatial dynamics of demands from live streaming users. We present optimal solutions to deal with cloud servers with diverse capacities and lease prices, as well as the potential latencies in initiating and terminating leases in real world cloud platforms. Our solution well accommodates location heterogeneity, mitigating the impact from user globalization. It also enables seamless migration for existing streaming systems, e.g., peer-to-peer, and fully explores their potentials. Simulations with data traces from both cloud service providers (Amazon EC2 and SpotCloud) and live streaming service provider (PPTV) demonstrate that CALMS effectively mitigates the overall system deployment costs and yet provides users with satisfactory streaming latency and rate.

Index Terms—Live media streaming, migration, cloud computing, user/demand globalization.

I. INTRODUCTION

IN the past decade, live media streaming has become one of the most popular applications over the Internet [14][6]. We have witnessed a number of successful commercial deployments with CDN (Content Delivery Network)- or peer-to-peer based engines. The former achieves high availability and short startup latencies, but suffers from excessive costs for deploying dedicated servers. This is particularly severe if the user demand fluctuates significantly and the servers have to be over-provisioned for peak loads. The peer-to-peer solution generally incurs lower deployment cost and is more scalable, but the reliability and hence service quality can hardly be guaranteed. There have also been efforts toward synergizing dedicated servers with peer-to-peer [7]. Unfortunately,

Feng Wang is with the Department of Computer and Information Science, the University of Mississippi, MS, 38677 (e-mail: fwang@cs.olemiss.edu).

Jiangchuan Liu is with the School of Computing Science, Simon Fraser University, British Columbia, V5A 1S6 Canada (e-mail: jcliu@cs.sfu.ca).

Minghua Chen is with the Department of Information Engineering, The Chinese University of Hong Kong, Shatin, NT, Hong Kong (e-mail: minghua@ie.cuhk.edu.hk).

Haiyang Wang is with the Department of Computer Science, the University of Minnesota Duluth, MN, 55812 (e-mail: wang4905@d.umn.edu).

A preliminary version of this paper appeared in IEEE INFOCOM’12.

Manuscript received xxx xx, 2013; revised xxx xx, 2013.




Fig. 1: An illustration of the user demand distributions and variations of a popular live media streaming system (PPTV) on its two typical channels (CCTV3 and DragonBall) during one day.

having an all-round scalable, reliable, responsive, and cost-effectiveness solution remains an illusive goal.

To make it even worse, today’s live streaming applications have become highly globalized, with subscribers from all over the world. Such a globalization makes user behaviors and demands even more diverse and dynamic. For illustration, we examine the user demand distribution of PPTV¹, one of the most popular live media streaming systems in China with multi-million users from a trace analysis [22][27]. Fig. 1 shows the results of two representative channels (CCTV3 and DragonBall) during one day². It is easy to see that although PPTV is from China, it has attracted users from all over the world, and the peak time therefore shifts from region to region, depending on the timezone. For example, on the CCTV3 channel, the peak time of North America is around 20:00, while for Asian users, it is around 8:00. During the period of 12:00 to 20:00, the Asian users have very low demands while the European users generate most of their demands and the North American users also have moderate demands. Similar observations can also be found from the DragonBall channel, despite that the streaming contents delivered on the two channels are completely different. The impact of such globalized demand turbulence has yet to be addressed in existing systems that largely focus on regional services only.

The emergence of *cloud computing* however sheds new lights into this dilemma. A cloud platform offers reliable, elastic and cost-effective resource provisioning, which has been dramatically changing the way of enabling scalable and dynamic network services [11][3][20][5]. There have been studies on demand-driven resource provision [24][13][26][28]; there have been also initial attempts leveraging cloud service to support VoD (Video-on-Demand) applications, from both industry (e.g. Netflix) [16] and academia [25][9][12]. Live

¹www.pptv.com – formerly known as PPLive.

²For ease of comparison, the user demands are normalized by the corresponding maximum demand of each day.

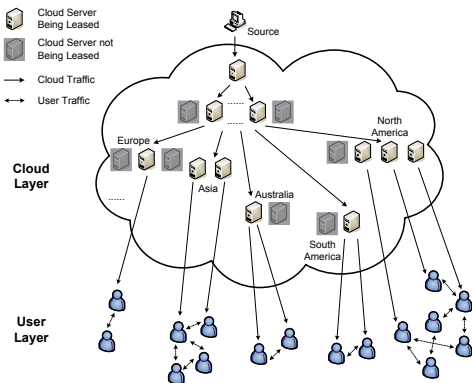


Fig. 2: An overview of CALMS.

media streaming however has more stringent playback delay constraints with content being updated in realtime. The larger, dynamic, and non-uniform client population further aggravates the problem, calling for new solutions toward a successful migration to the cloud.

In this paper, we present CALMS (Cloud-Assisted Live Media Streaming), a generic framework that facilitates a cost-effective migration to the cloud. CALMS adaptively leases and adjusts cloud servers in a fine granularity to accommodate temporal and spatial dynamics of user demands. We present optimal solutions to deal with cloud servers with diverse capacities and lease prices, as well as the potential latencies in initiating and terminating leases in real world cloud platforms. Our solution well accommodates location diversity, mitigating the impact from user globalization. It also enables seamless migration for existing streaming systems, e.g., peer-to-peer, and fully explores their potentials. We further develop a set of practical solutions for dynamic adjustment of lease schedules, smart user redirection, and cloud server organization. To understand the performance of CALMS, extensive simulations have been carried out with real data traces from both cloud service providers (Amazon EC2 and SpotCloud) and live media streaming service provider (PPTV). The results demonstrate that the proposed CALMS effectively mitigates the overall system deployment costs and yet provides users with satisfactory streaming latency and rate.

The remainder of this paper proceeds as follows: Section II presents an overview of the framework. In Section III, we first investigate the basic problem of leasing cloud service and provide an optimal solution, which is then extended by integrating locality-awareness and user assistance. The implementation issues and further optimization are discussed in Section IV. We evaluate CALMS by trace-driven simulations in Section V. Section VI further discusses some open issues. Finally, Section VII concludes the paper and discusses potential future directions.

II. CLOUD-ASSISTED LIVE MEDIA STREAMING (CALMS): AN OVERVIEW

Our CALMS intends to provide a generic framework that facilitates the migration of existing live media streaming services to a cloud-assisted solution. Fig. 2 shows an illustration of CALMS, which is divided into two layers, namely, *Cloud Layer* and *User Layer*. The Cloud Layer consists of the live

media source and dynamically leased cloud servers. Upon receiving a user's subscription request, the Cloud Layer redirects this user to a properly selected cloud server. Such a re-direction is transparent to the user, i.e., the Cloud Layer is deemed as a single source server from a user's perspective. Since the user demands change over time, which are also location-dependent, the Cloud Layer accordingly dynamically adjusts the amount and location distribution of the leased servers. Intuitively, it leases more server resources upon demand increase during peak times, and terminates leases upon decrease.

There are however a number of critical theoretical and practical issues to be addressed in this generic CALMS framework. First, the cloud servers have diverse capacities and lease prices; the lease duration is not infinitesimally short, either, e.g. 1 hour for Amazon EC2. As such, when being leased, the server and the pricing cannot be simply terminated at anytime. In addition, for a newly leased server, the configuration and boot up takes time, too, e.g., about 1-10 minutes for EC2 mainly depending on the used Operating System. Though the cloud services are improving, given the hardware, software, and network limits, such latencies can hardly be avoided in the near future. Therefore, CALMS must well predict when to lease a new servers to meet the ever-changing demands and when to terminate a server to minimize the lease costs. These problems are further complicated given the global heterogeneous distributions of the cloud servers and that of the user demands.

Note that we do not assume any particular implementation of the User Layer in this study. They can be individual users purely relying on the Cloud Layer, or served by peer-to-peer or a CDN infrastructure, but seeking for extra assistance from the Cloud Layer during load surges. In other words, our CALMS framework can smoothly migrate diverse existing live streaming systems. Also, we will explore the potential assistance from user peers in our study by investigating general solutions that well complement existing system designs, taking into account different issues on user dynamics such as user churn, user mobility and identifying good potential user helpers.

III. FRAMEWORK DESIGN AND SOLUTION

In this section, we discuss the detailed design issues and their solutions for migrating the live media streaming application to the cloud service. Our discussions first start from modeling the basic form of the leasing cloud service problem, and then extend to integrate with locality-awareness and user-assistance, respectively. We will also present centralized algorithms that yield optimized solutions for addressing these issues, which will further motivate the practical solutions for implementation and optimization in next section.

A. Basic Problem: Leasing Cloud Service

Denote the set of cloud servers that can be leased from the cloud service providers as $C = \{c_1, c_2, \dots, c_m\}$. In practice, most cloud providers have a minimum unit time for the duration of leasing a server (e.g. 1 hour for Amazon EC2) and when being leased, a cloud server must spend some time in setup and initialization before ready to use. We use

D_m to denote this required minimum unit duration and T_s as the latency for preparing the cloud server. For simplicity, we assume there is always a cloud server directly connecting to the live media source to act as the live media server and use c_0 to denote it. Let R be the rate of the live media streaming. We assume that there is a demand forecast algorithm (such as the algorithms proposed in [17]) to predict the demand in the next period T , where the demand may contain the estimated online population of users, their distributions at different areas or ISPs, and other type of information. At current stage, we only consider the online population of users and denote it as $\mathbb{P}(t)$ for a given time t ($t \leq T$). The discussions for utilizing other forecasted demand information will be deferred to later subsections. Define a cloud service lease schedule as $S = \{(x_1, t_1, d_1), (x_2, t_2, d_2), \dots, (x_k, t_k, d_k)\}$ ($t_1 \leq t_2 \leq \dots \leq t_k \leq T$), where a tuple (x_i, t_i, d_i) ($x_i \in C$, $d_i > 0$ and $d_i \bmod D_m = 0$ for $i = 1, 2, \dots, k$) means at time t_i , we start to lease cloud server x_i for the duration d_i . Our problem is thus to find a proper cloud lease schedule S , subjecting to the following constraints:

(1) Service Availability Constraint:

$$\begin{aligned} \forall (x_i, t_i, d_i) \in S, \text{ if } \exists (x_j, t_j, d_j) \in S \text{ and } x_i = x_j, \\ \text{then } [t_i, t_i + d_i] \cap [t_j, t_j + d_j] = \emptyset ; \end{aligned}$$

(2) Streaming Quality Constraint:

$$\forall T_s \leq t \leq T,$$

$$\mathbb{U}(c_0) + \sum_{(x_i, t_i, d_i) \in S} (\mathbb{U}(x_i) - R) \cdot I_{[t \in [t_i + T_s, t_i + d_i]]} \geq R \cdot \mathbb{P}(t) ;$$

where $\mathbb{U}(\cdot)$ is the upload capacity and $I_{[\cdot]}$ is the indicator function. The service availability constraint asks that at any given time, a cloud server can only appear in one schedule. And the streaming quality constraint asks that at any given time t , the streaming rate demands of $\mathbb{P}(t)$ users have to be satisfied. Our objective is thus to minimize the lease costs:

$$\mathbb{C}_l(c_0) + \sum_{(x_i, t_i, d_i) \in S} \mathbb{C}_l(x_i) \cdot d_i + \mathbb{C}_b(R \cdot \sum_{t \leq T} \mathbb{P}(t)) ,$$

where $\mathbb{C}_l(\cdot)$ and $\mathbb{C}_b(\cdot)$ are the costs for leasing a cloud server and out-cloud bandwidth usage, respectively³. As the first and last part can not be reduced, we focus on minimizing the middle part of the total costs, which we denote as $Cost_{lease}$:

$$Cost_{lease} = \sum_{(x_i, t_i, d_i) \in S} \mathbb{C}_l(x_i) \cdot d_i .$$

This problem is challenging as the cloud servers are scheduled both along the time dimension and at each time instance, along the user demand dimension. By exhaustively searching along both dimensions, the optimal solution can be achieved. However, simply using a naive searching algorithm can be quite inefficient as the solution space increases very fast with

³Amazon EC2 had no charges on the traffics into the cloud as well as within the cloud during this research was conducted. Yet its new policy starts to charge on the traffics among different AWS regions. Our model and solution here can be easily adapted to the new policy by adding an extra cost on the traffic of one streaming rate from outside of the region to inside of the region (except for the region which c_0 is in) for each used AWS region.

the combination of both dimensions.

To this end, we proposed an enhanced DFS algorithm, which can skip most parts of the solution space and find the optimal solution efficiently. The proposed algorithm is summarized in Fig. 3. To improve the search efficiency, we first sort the cloud servers in C by the ascendant order of the lease cost per unit upload bandwidth (line 1). This allows the servers with cheaper upload bandwidth being explored first and near-optimal solutions can then be quickly found. With such solutions, we can further cut other search branches with equal or higher costs (line 7) and greatly reduce the search space for the optimal solution. In addition, we use $leaseTime[c]$ to denote the remained lease time of server c and $leaseTime[c] > 0$ means that server c is leased. Since a newly leased server needs T_s for preparation, our algorithm makes decisions in advance of T_s , i.e., at $time$ we lease a new server c so that it can start to provide the service at $time + T_s$. Similarly, for a server c to be renewable at $time$, $leaseTime[c]$ must be equal to T_s , so that after we decide to renew it at $time$, it will continue providing the service at $time + T_s$.

We use variable $renew$ to distinguish whether the currently considered k -th server is renewed or newly leased (line 11-12). When all renewable servers have been considered, we also check if any new server needs to be leased (line 13). If so, we will reset k and $renew$ to further explore the branch (line 14). Every time a server c_k is selected (line 17-18), it will be leased for the time of D_m , i.e., $leaseTime[c_k]$ is increased by D_m . $Load[time + T_s]$ to $Load[time + T_s + D_m]$ (or $Load[time + D_m]$, depending on whether c_k is renewed or newly leased) will be reduced by $\mathbb{U}(c_k) - R$. After that, the algorithm checks if any other server needs to be leased or can be renewed (line 19). If not, $time$ will be increased (by $\Delta time$) to another position where a server can be renewed or new servers need to be leased. Then both k and $renew$ will be reset accordingly. $leastTime$ will also be updated by subtracting $\Delta time$ for leased servers.

When a search branch is cut off or fully explored, the search will revert to its previous status (line 25-27), where the previous $\{k, time, renew, leastTime\}$ will be popped up from *Stack* with *Load* and *Cost* being calculated oppositely to line 18. When the search finishes, the optimal cloud lease schedule S will be generated and returned (line 29-30), where a tuple (c_i, t_i, D_m) will be created and added into S for a server c_i newly leased at t_i , and another D_m will be added to the tuple's lease duration if it is renewed afterwards. We then have the following theorem.

Theorem 1: The algorithm proposed in Fig. 3 returns the optimal cloud lease schedule for the basic problem.

Proof: If without cutting branches, our solution will search exhaustively and return the optimal schedule. The proof thus can be done by showing that cutting branches (line 7) do not miss the optimal schedule, since $Cost$ is a lower bound of those schedules whose search paths contain the current search branch. A more detailed proof can be found in [21]. ■

B. Integrating with Locality-Awareness

As the users may be from various locations or time zones over the world, registering to different ISPs and in different

areas, a further optimization to the basic problem is thus to integrate with the locality-awareness, i.e., to maximize the number of users that connect to the local cloud servers. By achieving this, the delay between users and cloud servers can be effectively reduced, which is often a crucial factor for live media streaming. Another advantage is that such locality-awareness can also help to reduce the cross-boundary traffic (e.g. cross-ISP traffic), which is especially important when considering user assistance as discussed in the following subsections. To this end, we use an abstract notation “region” to represent the locality that we care about, which can be interpreted into different meanings in different contexts (e.g. an area with some extent of physical closeness or a group of ASes belonging to one ISP). Let $\mathbb{A} = \{A_1, A_2, \dots, A_n\}$ be the set of regions that the cloud servers and users may be in. For a cloud server c_i , we use $c_i \in A_j$ to denote that it is in the region A_j . In addition, we further extend $\mathbb{P}(t)$ to $\mathbb{P}(A, t)$ to denote the online population of users in region A at time t . The basic problem thus can be rewritten as to find a proper cloud lease schedule S , subjecting to the service availability constraint and the following new streaming quality constraint:

$$\begin{aligned} & \forall T_s \leq t \leq T, \\ & \mathbb{U}(c_0) + \sum_{(x_i, t_i, d_i) \in S} (\mathbb{U}(x_i) - R) \cdot I_{[t \in [t_i + T_s, t_i + d_i]]} \\ & \geq R \cdot \sum_{A \in \mathbb{A}} \mathbb{P}(A, t). \end{aligned}$$

Our objective is now to minimize the lease costs:

$$Cost_{lease} = \sum_{(x_i, t_i, d_i) \in S} \mathbb{C}_l(x_i) \cdot d_i,$$

as well as maximize the locality, which is defined as:

$$\begin{aligned} Locality = & \frac{1}{R \cdot \sum_{t \leq T} \sum_{A \in \mathbb{A}} \mathbb{P}(A, t)} \cdot \sum_{t \leq T} \sum_{A \in \mathbb{A}} \min \left(R \cdot \mathbb{P}(A, t), \right. \\ & \left. \sum_{(x_i, t_i, d_i) \in S} \mathbb{U}(x_i) \cdot I_{[t \in [t_i + T_s, t_i + d_i], x_i \in A]} \right). \end{aligned}$$

These two objectives may contradict with each other, as leasing more servers in each region improves the locality but also raises the lease cost. We adopt the following linear combination form to align them together with different weights:

$$p \cdot \frac{Cost_{lease}}{Cost_{max}} + q \cdot (1 - Locality),$$

where p and q are two parameters that can assign different weights to the two goals. Assuming that the demand estimation is upper bounded over time duration T , $Cost_{max}$ is thus the minimum lease costs of the case where the demand is constantly set to be the maximum demand within T . As $Locality$ is a ratio of the intra-region streaming traffic over the total streaming traffic, $(1 - Locality)$ is thus the ratio of the cross-region traffic over the total traffic, which should also be minimized like $Cost_{lease}$. To make the lease cost part also a ratio ranged between $[0, 1]$, we further divide $Cost_{lease}$ by $Cost_{max}$ and then use parameters p and q to linearly combine

Algorithm OptimalCloudLeaseSchedule()

```

1: Sort  $C$  by ascendant order of  $\mathbb{C}_l(\cdot)/\mathbb{U}(\cdot)$ ;
2: for  $\forall c \in C$ ,  $leaseTime[c] \leftarrow 0$ ; end for
3: for  $\forall t \leq T$ ,  $Load[t] \leftarrow R \cdot \mathbb{P}(t) - \mathbb{U}(c_0)$ ; end for
4: Set  $Stack$  empty;  $k \leftarrow 0$ ;  $time \leftarrow 0$ ;  $Cost \leftarrow 0$ ;
5:  $renew \leftarrow \text{false}$ ;  $Cost^* \leftarrow \infty$ ; Set  $Stack^*$  empty;
6: do
7:   if  $Cost \geq Cost^*$ , goto 25;
8:   else if  $time + T_s > T$ ,
9:      $Cost^* \leftarrow Cost$ ;  $Stack^* \leftarrow Stack$ ; goto 25;
10:  end if
11:   $k \leftarrow k + 1$ ;
12:  if  $k \leq |C|$  and  $(!renew \text{ and } leaseTime[c_k] > 0)$  or  $(renew \text{ and } leaseTime[c_k] \neq T_s)$ , continue;
13:  else if  $renew \text{ and } k > |C| + 1$  and  $Load[time + T_s] > 0$ ,
14:     $k \leftarrow 0$ ;  $renew \leftarrow \text{false}$ ; continue;
15:  else if  $(!renew \text{ and } k > |C|)$  or  $(renew \text{ and } k > |C| + 1)$ , goto 25;
16:  end if
17:  Push  $\{k, time, renew, leaseTime\}$  in  $Stack$ ;
18:  Update  $Load$ ,  $leaseTime$  and  $Cost$ ;
19:  if  $Load[time + T_s] \leq 0$  and  $(!renew \text{ or } k = |C| + 1)$ ,
20:    Increase  $time$  until  $time + T_s > T$  or
     $\exists c \in C$ ,  $leaseTime[c] - \Delta time = T_s$  or
     $Load[time + T_s] > 0$ ;
21:     $k \leftarrow 0$ ; Update  $leaseTime$  by  $\Delta time$ ;
22:     $renew \leftarrow \text{if } \exists c \in C, leaseTime[c] = T_s$ ;
23:  end if
24:  continue;
25:  if  $Stack$  is not empty,
26:    Pop  $Stack$ ; Update  $Load$  and  $Cost$ ;
27:  end if
28:  while  $Stack$  is not empty;
29:  Generate optimal cloud lease schedule  $S$  from  $Stack^*$ ;
30:  return  $S$ ;
    
```

Fig. 3: Algorithm to compute the optimal cloud lease schedule.

the two parts together. This new problem can also be solved by the algorithm proposed in Fig. 3, with some proper but simple modifications. We defer the detailed discussion to the next subsection with the consideration of exploring user resources.

C. Exploring User Resources

It is known that peer-to-peer streaming is highly scalable through exploring user contributed resources. However, in a pure peer-to-peer system, users may join or leave at their own wills, and their available upload bandwidth may vary significantly from time to time and from user to user. Even if the aggregate user contributed upload bandwidth is equal to or greater than the total demand, a pure peer-to-peer design may still suffer from content bottlenecks [15], where users have upload bandwidth but no expected streaming content available for sharing. An extreme is a flashcrowd; that is, during a peak time, many fresh users join the system with a vast amount of ready-to-share upload bandwidth but no content available.

Our CALMS could also benefit from such readily available resources in the User Layer, and yet it can elegantly mitigate the aforementioned problems. To this end, we introduce a group of parameters (α, β) , which we call as *user assistance index*. Both of the two parameters range from 0 to 1. α determines the ratio of the user contributed upload bandwidth that

can be effectively used to assist the live streaming. β denotes the minimum ratio of cloud servers to be reserved to deal with the content bottleneck (since the cloud servers always have the updated streaming content to share). The full version of our problem thus can be written as to find a proper cloud lease schedule S , subjecting to the service availability constraint and the following updated streaming quality constraint that considers both locality and user resources:

$$\begin{aligned} & \forall T_s \leq t \leq T, \\ & \mathbb{U}(c_0) + \sum_{(x_i, t_i, d_i) \in S} (\mathbb{U}(x_i) - R) \cdot I_{[t \in [t_i + T_s, t_i + d_i]]} \\ & + \min \left(\alpha \cdot \sum_{A \in \mathbb{A}} \mathbb{B}(A, t), (1 - \beta) \cdot R \cdot \sum_{A \in \mathbb{A}} \mathbb{P}(A, t) \right) \\ & \geq R \cdot \sum_{A \in \mathbb{A}} \mathbb{P}(A, t); \end{aligned}$$

where $\mathbb{B}(A, t)$ is the aggregate user upload bandwidth in region A at time t . Our objective is still to minimize

$$p \cdot \frac{Cost_{lease}}{Cost_{max}} + q \cdot (1 - Locality),$$

while with the user-assistance taken into account, the locality measure is now calculated as

$$\begin{aligned} Locality = & \frac{1}{R \cdot \sum_{t \leq T} \sum_{A \in \mathbb{A}} \mathbb{P}(A, t)} \cdot \sum_{t \leq T} \sum_{A \in \mathbb{A}} \min \left(R \cdot \mathbb{P}(A, t), \right. \\ & \left. \sum_{(x_i, t_i, d_i) \in S} \mathbb{U}(x_i) \cdot I_{[t \in [t_i + T_s, t_i + d_i], x_i \in A]} \right. \\ & \left. + \min (\alpha \cdot \mathbb{B}(A, t), (1 - \beta) \cdot R \cdot \mathbb{P}(A, t)) \right), \end{aligned}$$

so as to maximize the traffics from local users/servers, thus reducing the latency and improving the overall performance.

To address this problem, some modifications need to be applied to the algorithm proposed in Fig. 3. First, *Load* now needs to have a second dimension to distinguish the demands from different regions and when initialized (line 3), it needs to further subtract the user contributed upload bandwidth. In addition, the *Cost* computation (line 18 and 26) now needs to integrate with the locality, where for a time instance $t \geq time$ (which means that new cloud servers still may be leased for this time instance), we use $R \cdot \mathbb{P}(A, t)$ to approximate the locality in a region A , so that the computed *Cost* still keeps to be a lower bound and the optimality of the cloud lease schedule returned by the algorithm can thus be achieved. We omit more details here due to the space limitation, which can be found in [21]. We then have the following corollary.

Corollary 1: The modified algorithm returns the optimal cloud lease schedule with the consideration of both locality and user resources.

IV. MIGRATION IMPLEMENTATION AND OPTIMIZATION

Previous section addresses the major design problems for CALMS. However, in practice, there still remain some issues that need to be considered carefully. First, the user demand and




Fig. 4: Illustration of Cloud Layer organization.

capacity may not be forecasted precisely, where the accuracy often degrades as the predication time gets more forward than the current time. This renders that a statically computed optimal cloud lease schedule for a long period may become less useful due to the prediction error. Another issue is how to organize the leased servers in the Cloud Layer. Since they may be from different regions, a careless organization may introduce unnecessary cross-boundary and cross-region traffics and also degrade the performance. Similar situation also happens to the User Layer, where users need to be carefully organized to enable assistances among each other as well as enforce locality and good QoS. In this section, we present our solutions to address these issues.

A. Cloud Layer Organization and Evolution

As the Cloud Layer is the core part of the CALMS framework, its organization is thus very important to the performance of the migrated live media streaming application. As the tree structure is well known for its efficiency for organization, in our implementation, we adopt a tree structure for the Cloud Layer. In particular, we let server c_0 be the root of the tree and in charge of the whole Cloud Layer. The servers in the interior part of the tree relay the live streaming content to other servers to amplify the upload capacity. And the remaining servers in outskirts then transmit the live streaming content to the User Layer.

Since the leased cloud servers may be from different regions, naively organizing servers into a tree may still cause poor performance. Fig. 4 shows an example, where each circle denotes a cloud server and the number inside denotes the region that the server belongs to. In this figure, there are 3 cloud servers of 2 unit upload capacities leased from each of the regions of 1 to 3, which are expected to provide similar upload capacities to each region. By a naive organization shown on the left part, the upload capacities provided by cloud servers to each region is 0, 4 and 6, which is greatly different from the expectation; while by handling carefully as shown on the right part, the upload capacities for each region becomes 3, 3 and 4. In addition, by the naive organization, to arrive at the User Layer, the live media streaming traffics must cross at least one region boundary and may cross as many as 3 region boundaries; while the carefully handled organization can effectively reduce such cross-boundary traffics.

To deal with these issues, we let server c_0 select and keep tracking a root server for each region from the leased servers. When a cloud server is newly leased, it will first be redirected to the root server in the same region, and the root server will be responsible to help the newly leased server join the subtree rooted at itself. If a region currently has no root server, server c_0 will temporarily take the role until a server from that region




Fig. 5: Measured bandwidth distribution (CDF) between cloud servers and users.




Fig. 6: Region distribution of PPTV users in one day's traces.




Fig. 7: Impacts of different p/q values on lease cost and cross-region traffic.




Fig. 8: Impacts of different α or β values on playback discontinuity.

is leased. And if a root server is stopped due to the decreased demand in its region, the root server will select another server in the same region from its descendants to take its role. If no such server exists, server c_0 will take the role temporarily. In addition, the root server of each region will also help to optimize the server organization within its region, such as allocating the servers with higher upload capacities closer to itself as well as moving a server to a new parent closer to the root server than the current parent, so that the height of the subtree rooted at the root server can be minimized.

Through our simulation results in Section V, we observe this design, combined with the locality-aware scheduling proposed in the previous section, can substantially reduce the cross-boundary traffic as compared to a straightforward approach without locality-awareness.

B. User Layer Organization and Evolution

The main task of the User Layer organization is to enforce locality and good QoS as well as enable user-assistance. When a user joins the live media streaming session, it will first be redirected to the root server at the same region if there is available upload capacity within that region, otherwise the user will be redirected to the root server of other region with available upload capacity. The root server then decides where the user should go next. If there is available upload capacity directly from the servers in the region, the user will then be redirected to the server that can provide the upload capacity. If not, the user will be redirected to a server that contains information about users that can provide the upload capacity. This server will then randomly choose some users with enough aggregate upload capacities and send their IPs to the newly joined user. The user will then add these users as neighbors and exchange the live media streaming content with them by a peer-to-peer protocol.

To enforce good QoS and handle the content bottleneck problem that may exist, when a server randomly picks users for the newly joined user, only those with high playback buffer levels will be selected and sent out by the server. In addition, we also let a user with high upload capacity to preempt a user that directly downloads from a server but with low upload capacity. And when the playback buffer level of a user becomes low, either due to user dynamics or network bandwidth fluctuation, the user will request more neighbors from its server to maintain good live streaming quality.

To track such user information at the Cloud Layer while keeping good scalability, we adopt a hierarchical structure

that is naturally provided by the tree organization. We let the last server that the user has been redirected to keep the full information of that user. Such information will then be aggregated and reported to this server's parent. This process will continue on level by level until reaching the root server at that region. The root server of each region will then aggregate and report the user information directly to server c_0 . By this means, when a user is being redirected, the current server that issues the redirection can then choose the next stop for the user based on the collected user information.

To deal with user dynamics, each user will periodically report its updated information to its server. The server will also check whether the user has left if a report has not been received recently. And when a server needs to be stopped due to the decreased demands from its region, it will first redirect all its users one by one to server c_0 to redo the join process, then stop and leave the Cloud Layer.

C. Dynamic Lease Scheduling

In practice, the user demand and capacity forecast may be inaccurate, and such inaccurate may cause a statically computed cloud lease schedule for a long future period becomes less useful or even invalid. To overcome this problem, we use dynamic lease scheduling instead for the migration implementation. In particular, with the collected aggregation information (such as current total user demand and total upload capacity) from each region, server c_0 will dynamically recompute the cloud lease schedule in the following two cases:

Case 1: If current user demands are greater than the prediction or current user upload capacities are less than the prediction, server c_0 will dynamically recalculate the cloud lease schedule with the updated information and then lease additional cloud servers by the new schedule.

Case 2: If a cloud server has been leased for the multiple times of D_m (i.e., if necessary, the cloud server can now be stopped without introducing further delay and cost), server c_0 will check if current user demands in that region are less than the half of current upload capacities even with this server stopped. If so, it will stop this server to reduce costs and recompute the cloud lease schedule accordingly.

In addition, due to the content bottleneck issue, sometimes the QoS perceived at users may degrade even the total upload capacity is still greater than the total user demand. To make the Cloud Layer responsive to such situations, we also collect the playback buffer level of each user when tracking other information. Server c_0 will then check the minimum buffer




Fig. 9: Lease costs of different approaches.




Fig. 10: Total costs of different approaches.




Fig. 11: Cross-region traffics of different approaches.




Fig. 12: Playback discontinuity of different approaches.

level of the users who have already started playing the media streaming. If the minimum level is below a threshold, which indicates potential content bottleneck may happen, server c_0 will lease a new server to increase the content availability and recompute a new lease schedule accordingly.

V. PERFORMANCE EVALUATION

We run extensive simulations to evaluate CALMS. The simulations are conducted on the block level and driven by real traces and data sets collected from Amazon EC2 and a popular live media streaming system (PPTV). We first briefly introduce the collected traces and data sets, and then continue on to present the methodology as well as the evaluation results.

A. Data Sets and Traces

1) *Amazon EC2 Measurement*: Our measurement on Amazon EC2 is mainly on the bandwidth distribution between cloud servers and users. In particular, we first send DNS requests to the EC2 domains to find the IP addresses of EC2 servers. The DNS server replies with a list of unique IP prefixes which are reserved by Amazon for their EC2 instances. Based on this IP list, we further probe these IP prefixes with ICMP, TCP and UDP packets from different locations, identifying active instances and then measuring their bandwidth accordingly. Fig. 5 shows the measured bandwidth distributions between different cloud servers and users. Other settings are adopted from previous measurement and evaluation works [4][11][18] and the Amazon EC2 official web site [2]. A more detailed description can be found in [21].

2) *PPTV Traces*: PPTV is one of the largest commercial peer-to-peer live streaming systems to date, attracting over 100,000 online users during peak times, and is also one of the mostly examined systems in academia [6][8][27]. Our simulations are based on traces from two of its popular channels, namely, CCTV3 and DragonBall. These traces are gathered by an online crawler that continuously collects information from each channel [22]. Fig. 6 shows the region distribution of PPTV users in one day's traces and Fig. 1 shows the user demands distributions and variations. Other settings about the live media streaming users are adopted from [6][27].

B. Methodology

With the data sets and traces from Amazon EC2 and PPTV, we then conduct extensive block level simulations to evaluate our migration framework. We adopt a typical live media streaming setting as used in PPTV [6] and CoolStreaming [14].

In particular, we set each data block is of 1-second video and assume TCP is used for the transmission. The playback buffer is set to the size enough to hold 60 data blocks. The playback will start when at least 30 continuous data blocks are received at the buffer. For comparison, we implement other three approaches: *Max-Central* statically provisions servers all from one region by the maximum user demand in the corresponding trace, which emulates the solution that uses centralized server clusters to provide the live media streaming service. *Max-CDN* also provisions servers by the maximum user demand, but the provisioned servers may be selected from different regions based on the average user demands from each region. This approach emulates the solution that uses CDNs to provide the streaming service. *P2P-Locality* only uses the server c_0 as the streaming server and delivers streaming content by peer-to-peer technology and with locality-awareness, which emulates the solution for peer-to-peer live media streaming.

In addition, we use the following four metrics in our simulation: *Lease costs* is the costs for leasing cloud servers, which represents the major concern from the live media streaming service providers; *Cross-region traffic* is the amount of traffic that crosses different regions. This metric shows the locality-awareness of an approach and is also an indicator to the general performance as lower cross-region traffic means that users are closer to their servers and the Cloud Layer are well organized; *Playback discontinuity* is the average duration that a user may experience discontinued playback per hour, which is mainly caused by that the streaming data packets fail to arrive at a user before its playback deadline; *Startup latency* is the latency taken by a user between its requesting to join the session and receiving enough data to start playback.

C. Impacts of Different Parameter Settings

We first conduct simulations to investigate the impacts of different parameter settings. Fig. 7 demonstrates how the lease cost and cross-region traffic change with different p/q values. For ease of comparison, the results are normalized by the corresponding minimum values. When p/q is small ($\leq 10^{-2}$), the cross-region traffic is minimized while introducing the highest lease costs. On the other hand, when p/q grows large ($\geq 10^2$), it results in the minimum lease cost but at the expenses of excessive cross-region traffic. Moreover, within the region near 10^{-1} , both the lease cost and cross-region traffic stay relatively low. We thus pick up this value $p/q = 0.1$ as the default for the remaining evaluations.

We next investigate the impacts of different α and β values on the playback discontinuity. The results are shown in Fig. 8.




Fig. 13: CDF of playback discontinuity of CALMS and P2P-Locality.




Fig. 14: Startup latency of different approaches.




Fig. 15: Location distribution of SpotCloud servers.




Fig. 16: Measured bandwidth distribution (CDF) between SpotCloud servers and users.

It is easy to see that the impact of β is more significant than that of α , while the results are different. When α becomes larger, the playback discontinuity slightly increases, which matches the definition of α since more user-assistances are involved, resulting in that more content bottlenecks may happen and degrade the playback quality. On the other hand, the playback discontinuity changes inversely with the value of β . This also matches the definition of β as reserving more cloud servers to compensate the content bottlenecks will improve the playback quality. More interestingly, the playback discontinuity will change more dramatically as either of the two parameters changes within the region of (0.2, 0.7). We thus choose $\alpha = 0.2$ and $\beta = 0.7$ as the default setting.

D. Performance Observed at Service Providers

With the default parameter setting, we then conduct simulations to see how CALMS performs with both CCTV3 and DragonBall traces. Fig. 9 shows the lease costs of different approaches. For ease of comparison, the lease costs in each trace are normalized by the corresponding cost of the Max-Central approach. It is not surprising that P2P-Locality has the lowest costs. At the same time, our CALMS also has much lower lease costs, achieving 33.5%-45.1% of the Max-Central approach and 30.5%-39.6% of the Max-CDN approach, respectively. Another observation is that the lease costs in the DragonBall trace are generally higher than those in the CCTV3 trace. This is because the content provided in the DragonBall channel attracts much more Asian user demands than those from the other regions and the lease prices of Amazon EC2 in Asia are relatively higher, which further verifies the locality-awareness of both the Max-CDN and CALMS approaches. Since Amazon EC2 also charges for the data traffic transferred out of the cloud boundary, we also examine the total costs of different approaches, which is shown in Fig. 10. It is easy to see that even with the data transfer costs, CALMS can still reduce a great amount of the total costs by roughly 28%-30%, which further demonstrates the effectiveness of migrating the live media streaming application to the cloud service.

Fig. 11 shows the cross-region traffics generated by different approaches, which are also normalized by the corresponding cross-region traffic of the Max-Central approach. As the Max-CDN, CALMS and P2P-Locality approaches are all locality-aware, it is thus not surprising that their cross-region traffics are much lower than that of the Max-Central approach. Yet an interesting thing is that the cross-region traffics of the three approaches in the DragonBall trace are much lower

than those in the CCTV3 trace. This is also because the DragonBall channel attracts much more Asian user demands than from the other regions, which allows the server/peer selections to be more dedicated in Asia and results in more intra-region traffics instead of cross-region traffics. Also, by comparing among the three approaches, we can find that by both dynamically provisioning enough cloud servers and exploiting user-assistance, CALMS can actually achieve a balance of the locality provided by either mechanism, always avoid the worst case – having the highest cross-region traffic, and stays close to the best one.

E. QoS Perceived by Users

Previous subsection has shown that migrating the live media streaming application to the cloud service can achieve good performance as observed at the service providers. Next we go on to explore possible benefits that may be brought to users.

We first investigate the playback discontinuity, which is shown in Fig. 12. We can see that CALMS performs similar to both Max- approaches and achieves very low playback discontinuity. On the other hand, due to the peer dynamics and content bottlenecks, P2P-Locality suffers relatively high playback discontinuity with the average as many as over 120 seconds per hour. To better understand the QoS perceived by each user, we also draw the CDF of the playback discontinuity for both CALMS and P2P-Locality in Fig. 13. It is clear to see that for P2P-Locality, some users may suffer playback discontinuity up to near 800 seconds per hour, which means that these users cannot watch the live media for more than 20% of the time. On the other hand, by provisioning enough cloud servers to provide both upload contents and capacities, CALMS significantly reduces the worst case playback discontinuity to around 1 minute and affords more than 95% of the users to enjoy zero playback discontinuity.

Besides watching live media fluently, a user also prefers to short waiting time before the live media can start to play. Fig. 14 shows the startup latency of different approaches. Due to the delay of being redirected to other servers and filling up the playback buffers, both Max- approaches have the startup latency of about 2 seconds. CALMS, by provisioning enough cloud servers with upload contents and capacities, can also achieve a similar startup latency. On the other hand, the startup latency of P2P-Locality is almost of doubled length of the other three approaches due to the long latency for identifying enough peers with both the live media streaming contents and available upload capacities.




Fig. 17: Lease costs of different approaches by SpotCloud.




Fig. 18: Cross-region traffics of different approaches by SpotCloud.




Fig. 19: Playback discontinuity of different approaches by SpotCloud.




Fig. 20: Startup latency of different approaches by SpotCloud.

VI. FURTHER DISCUSSIONS

As this paper focuses on providing fundamental understandings on migrating the live media streaming to the cloud, we only provide basic schemes in Section IV to handle the inaccuracies in the user demand/capacity forecast. It is yet interesting to explore the impact of the forecast accuracy on the performance. Based on our preliminary simulations and analysis, it can be briefly summarized in several folds. First, if the user demand is overestimated, the cross-region traffic and the QoS perceived by users are not affected, while the lease cost increases. If the demand is underestimated, both the lease cost and QoS may drop and the cross-region traffic may increase. On the other hand, if the user capacity is underestimated, both the cross-region traffic and QoS are not affected, while the lease cost may increase. And if overestimated, it may degrade the QoS, increase the cross-region traffic, and reduce the lease cost. In CALMS, the capacity forecast inaccuracy is partially overcome by carefully setting α and β , where low α value or high β value can reduce the framework's dependence on the user capacities but may introduce extra lease cost. For the demand forecast error, we believe that to the overall performance, overestimation is better than underestimation. This suggests that one may map the demand forecast problem to the classic ski-rental problem [10] and apply the simple break-even online algorithm. We omit more details here due to the space limitation, which can be found in [21]. Our preliminary results show as long as the decisions of different approaches are made by the same demand forecast algorithm with only overestimation errors, CALMS can still achieve the similar performance as shown in the last section.

Another open issue is that recently increasing attentions from both industries and academia have been attracted by a new type of cloud platform featured as allowing the users to contribute/sell their own idle computing resources and build their own cloud services. It is thus interesting to investigate the potential of this type of platform to migrate the live streaming application. To this end, we select a typical example, the Enomaly's SpotCloud [19][23], to run simulations with the measured information shown in Figs. 15 and 16. The results are shown in Figs. 17-20, where the lease cost and cross-region traffic are normalized by those values of the Max-Central approach on Amazon EC2. Besides the superiority of our CALMS solution, we find two interesting observations comparing to the results on Amazon EC2: First, the lease cost of this new type of platform is very low and generally less than 10% of Amazon EC2. This is because this new platform dis-

tributes the maintenance cost to the server's contributor/seller, greatly reducing the lease cost of each server. The other observation is that in this new type of platform, although the server contributor/sellers are highly geo-distributed (e.g. 8 countries for SpotCloud in Fig. 15), the aggregate server capacities may not be as strong as the datacenter clouds (if there is one) in the same region. This explains why the cross-region traffic by SpotCloud is much higher than Amazon EC2. These observations motivate a possible hybrid design to use both types of cloud platforms to provide the live streaming application, where the datacenter clouds as the backbone to enforce enough server capacities and good performance, with the servers contributed/sold by others from this new type of platform as edge servers to further reduce the lease costs.

VII. CONCLUSION AND FUTURE WORK

This paper presented CALMS, a generic framework that facilitates migrating live media streaming to a cloud platform. Being the very first paper towards this direction, we strived to offer fundamental understandings on the practical feasibility and theoretical constraints in the migration. We first modeled the basic problem of leasing cloud services to support time-varying user demands and developed an optimal algorithm. We then extended our solution to integrate locality-awareness and user-assistance, alleviating the impact from the service globalization. We further designed a series of practical solutions for both cloud and user layer organization and optimization, as well as dynamic lease scheduling that accommodates inaccuracy in demand and capacity forecast. Extensive simulations driven by traces from both cloud service providers (Amazon EC2 and SpotCloud) as well as live media streaming service provider (PPTV) demonstrated the cost-effectiveness and superior streaming quality of CALMS, even with highly dynamic and globalized demands.

We are currently conducting more simulations to evaluate and improve CALMS with data sets and traces from other cloud providers and live media streaming providers. We expect to develop a prototype for further verification and evaluation. We are also interested in exploring other open issues such as designing better user demand/capacity forecast mechanisms and extending CALMS to other type of cloud platform, e.g., Amazon CloudFront [1], a cloud-based CDN platform.


ACKNOWLEDGMENT

Feng Wang's research is supported by a Start-up Grant from the University of Mississippi, MS. Jiangchuan Liu's

research is supported by a Canada NSERC Discovery Grant. Minghua Chen's research is supported by National Basic Research Program of China (Project No. 2013CB336700) and the University Grants Committee of the Hong Kong Special Administrative Region, China (Area of Excellence Grant Project No. AoE/E-02/08). The authors thank Dr. Xiaojun Hei for providing PPTV traces.

REFERENCES

- [1] Amazon CloudFront, <http://aws.amazon.com/cloudfront/>.
- [2] Amazon Elastic Compute Cloud, <http://aws.amazon.com/ec2/>.
- [3] M. Armbrust, R. G. A. Fox, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "Above the Clouds: A Berkeley View of Cloud Computing," University of California, Berkeley, Tech. Rep., 2009.
- [4] S. Garfinkel, "An Evaluation of Amazon's Grid Computing Services : EC2 , S3 and SQS," Harvard University, Tech. Rep., 2008.
- [5] M. Hajjat, X. Sun, Y.-W. E. Sung, D. Maltz, S. Rao, K. Sripanidkulchai, and M. Tawarmalani, "Cloudward Bound: Planning for Beneficial Migration of Enterprise Applications to the Cloud," in *ACM SIGCOMM*, 2010.
- [6] X. Hei, C. Liang, J. Liang, Y. Liu, and K. Ross, "A Measurement Study of a Large-Scale P2P IPTV System," *IEEE Trans. Multimedia*, vol. 9, no. 8, pp. 1672–1687, December 2007.
- [7] C. Huang, J. Li, and K. W. Ross, "Peer-Assisted VoD: Making Internet Video Distribution Cheap," in *USENIX IPTPS*, 2007.
- [8] Y. Huang, T. Z. J. Fu, D.-M. Chiu, J. C. S. Lui, and C. Huang, "Challenges, design and analysis of a large-scale P2P-VoD system," in *ACM SIGCOMM*, 2008.
- [9] Z. Huang, C. Mei, L. E. Li, and T. Woo, "CloudStream: Delivering High-Quality Streaming Videos through a Cloud-Based SVC Proxy," in *IEEE INFOCOM Mini-Conference*, 2011.
- [10] A. Karlin, M. Manasse, L. Rudolph, and D. Sleator, "Competitive Snoopy Caching," *Algorithmica*, vol. 3, no. 1, pp. 79–119, 1988.
- [11] A. Li, X. Yang, S. Kandula, and M. Zhang, "CloudCmp: Comparing Public Cloud Providers," in *ACM IMC*, 2010.
- [12] H. Li, X. Cheng, and J. Liu, "Understanding Video Sharing Propagation in Social Networks: Measurement and Analysis," to appear in *ACM Transactions on Multimedia Computing, Communications and Applications*, 2014.
- [13] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, "Dynamic Right-Sizing for Power-Proportional Data Centers," *IEEE/ACM Transactions on Networking*, vol. 21, no. 5, pp. 1378–1391, October 2013.
- [14] J. Liu, S. G. Rao, B. Li, and H. Zhang, "Opportunities and Challenges of Peer-to-Peer Internet Video Broadcast," *Proceedings of the IEEE*, vol. 96, no. 1, pp. 11–24, January 2008.
- [15] N. Magharei and R. Rejaie, "PRIME: Peer-to-Peer Receiver-driven MESH-Based Streaming," in *IEEE INFOCOM*, 2007.
- [16] Netflix, <http://www.netflix.com>.
- [17] D. Niu, Z. Liu, B. Li, and S. Zhao, "Demand Forecast and Performance Prediction in Peer-Assisted On-Demand Streaming Systems," in *IEEE INFOCOM Mini-Conference*, 2011.
- [18] R. Shea, F. Wang, H. Wang, and J. Liu, "A Deep Investigation Into Network Performance in Virtual Machine Based Cloud Environment," in *IEEE INFOCOM*, 2014.
- [19] SpotCloud, <http://www.spotcloud.com/>.
- [20] K. Sripanidkulchai, S. Sahu, Y. Ruan, A. Shaikh, and C. Dorai, "Are Clouds Ready for Large Distributed Applications?" in *SOSP LADIS Workshop*, 2009.
- [21] F. Wang, J. Liu, M. Chen, and H. Wang, "On Cloud-Assisted Live Media Streaming," University of Mississippi, Tech. Rep., 2014.
- [22] F. Wang, Y. Xiong, and J. Liu, "mTreebone: A Collaborative Tree-Mesh Overlay Network for Multicast Video Streaming," *IEEE Trans. Parallel and Distributed Systems*, vol. 21, no. 3, pp. 379–392, March 2010.
- [23] H. Wang, F. Wang, J. Liu, and J. Groen, "Measurement and Utilization of Customer-Provided Resources for Cloud Computing," in *IEEE INFOCOM*, 2012.
- [24] C. Wu, B. Li, and S. Zhao, "On Dynamic Server Provisioning in Multichannel P2P Live Streaming," *IEEE/ACM Transactions on Networking*, vol. 19, no. 5, pp. 1317–1330, October 2011.
- [25] Y. Wu, C. Wu, B. Li, X. Qiu, and F. C. Lau, "CloudMedia: When Cloud on Demand Meets Video on Demand," in *IEEE ICDCS*, 2011.
- [26] H. Xu and B. Li, "Joint request mapping and response routing for geodistributed cloud services," in *IEEE INFOCOM*, 2013.
- [27] K. Xu, H. Li, J. Liu, W. Zhu, and W. Wang, "PPVA: A Universal and Transparent Peer-to-Peer Accelerator for Interactive Online Video Sharing," in *IWQoS*, 2010.
- [28] Q. Zhang, M. F. Zhani, R. Boutaba, and J. L. Hellerstein, "HARMONY: Dynamic Heterogeneity-Aware Resource Provisioning in the Cloud," in *IEEE ICDCS*, 2013.



Feng Wang (S'07-M'13) received both the Bachelor's degree and Master's degree in Computer Science and Technology from Tsinghua University, Beijing, China in 2002 and 2005, respectively. He received the PhD degree in Computing Science from Simon Fraser University, Burnaby, British Columbia, Canada in 2012. He is currently an assistant professor in the Department of Computer and Information Science at the University of Mississippi, MS, USA. His research interests include cloud computing, peer-to-peer networks, socialized content sharing, wireless mesh/sensor networks and cyber-physical systems. He is a TPC co-chair in IEEE ICC'14 for Symposium on Wireless Networking and Multimedia.



Jiangchuan Liu (S'01-M'03-SM'08) received the BEng degree (cum laude) from Tsinghua University, Beijing, China, in 1999, and the PhD degree from The Hong Kong University of Science and Technology in 2003, both in computer science. He is a co-recipient of ACM TOMCCAP Nicolas D. Georganas Best Paper Award 2013, ACM Multimedia Best Paper Award 2012, IEEE Globecom 2011 Best Paper Award, and IEEE Communications Society Best Paper Award on Multimedia Communications 2009.

He is a Full Professor in the School of Computing Science, Simon Fraser University, British Columbia, Canada, and an EMC-Endowed Visiting Chair Professor of Tsinghua University, Beijing, China. He was an Assistant Professor in the Department of Computer Science and Engineering at The Chinese University of Hong Kong from 2003 to 2004.

His research interests include multimedia systems and networks, cloud computing, social networking, wireless ad hoc and sensor networks, and peer-to-peer and overlay networks. He serves on the editorial boards of IEEE Transactions on Multimedia, IEEE Communications Surveys and Tutorials, IEEE Access, and IEEE Internet of Things Journal. He is a TPC co-chair of IEEE/ACM IWQoS'14 and an Area Chair of ACM Multimedia'14.



Minghua Chen (S'04-M'07) Minghua Chen received his B.Eng. and M.S. degrees from the Department of Electronic Engineering at Tsinghua University in 1999 and 2001, respectively. He received his Ph.D. degree from the Department of Electrical Engineering and Computer Sciences at University of California at Berkeley in 2006. He spent one year visiting Microsoft Research Redmond as a Postdoc Researcher. He joined the Department of Information Engineering, the Chinese University of Hong Kong, in 2007, where he currently is an

Associate Professor. He is also an Adjunct Associate Professor in Peking University Shenzhen Graduate School in 2011-2014. He received the Eli Jury award from UC Berkeley in 2007 (presented to a graduate student or recent alumnus for outstanding achievement in the area of Systems, Communications, Control, or Signal Processing) and The Chinese University of Hong Kong Young Researcher Award in 2013. He also received several best paper awards, including the IEEE ICME Best Paper Award in 2009, the IEEE Transactions on Multimedia Prize Paper Award in 2009, and the ACM Multimedia Best Paper Award in 2012. He is currently an Associate Editor of the IEEE/ACM Transactions on Networking. His recent research interests include energy systems (e.g., microgrids and energy-efficient data centers), distributed optimization, multimedia networking, wireless networking, network coding, and secure network communications.

Haiyang Wang (S'08-M'13) is currently an Assistant Professor in the Department of Computer Science at the University of Minnesota Duluth, MN, USA. He received the PhD degree in Computing Science from Simon Fraser University, Burnaby, British Columbia, Canada in 2013. His research interests include cloud computing, big data, socialized content sharing, multimedia communications and peer-to-peer networks.

