

Best Papers of the ACM Multimedia Systems (MMSys) Conference 2017 and the ACM Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV) 2017

The 8th ACM International Conference on Multimedia Systems (MMSys'17) was successfully held in Taipei, Taiwan, between June 20 and 23, 2017, with the co-located workshops: the 27th ACM SIGMM Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'17), the 9th International Workshop on Massively Multiuser Virtual Environments (MMVE'17), and the 15th Annual Workshop on Network and Systems Support for Games (NetGames'17).

There were several initiatives at ACM MMSys'17 for: (i) ensuring reproducibility, (ii) broadening the topics, and (iii) increasing interests from scientists and practitioners. In particular, ACM MMSys'17 was the first ACM SIGMM conference to implement the new ACM reproducibility badge system. The scientific reproducibility committee of the conference has contacted all accepted papers and invited the authors to make their dataset and code available, in order to obtain the ACM badge, which is visible at the ACM Digital Library. Moreover, to further grow the multimedia systems community, ACM MMSys'17 added four special sessions: (i) Cloud-based Multimedia, (ii) Multi-sensory Experiences, (iii) Immersive Experiences in AR (Augmented Reality) and VR (Virtual Reality), and (iv) IoT (Internet-of-Things) Connectivity and Infrastructure. We are happy to see that these special sessions attract submissions from authors who are new to the MMSys community, resulting in high diversity in the accepted papers and conference sessions.

ACM MMSys'17 accepted 13 Research Track papers out of 47 submissions, leading to an acceptance rate of 28%. In the special sessions, ACM MMSys'17 accepted 11 papers out of the 23 submissions to create more diversity and attract new researchers to the MMSys community. Several high-quality ACM MMSys'17 papers have been shortlisted as candidates to Best Paper Awards and DASH (Dynamic Adaptive Streaming over HTTP) Industry Forum Excellence in DASH Award. As a good tradition, the ACM MMSys'17 Technical Program Committee (TPC) co-chairs invited the authors of these high-quality papers to submit their extended manuscripts to this TOMM special section. The ACM NOSSDAV'17 TPC co-chairs also invited the authors of the NOSSDAV'17 papers with high review scores to submit. After a couple of rounds of thorough reviews, mainly on the new materials beyond the conference publications required by ACM TOMM, we accepted seven submissions in this special section.

The first two articles are on video streaming. In "ASAP: Adaptive Stall-Aware Pacing for Improved DASH Video Experience in Cellular Network" (<https://doi.org/10.1145/3219750>), Zahran et al. present novel DASH traffic management solutions, which reduce the playback stalls while maintaining good Quality of Experience (QoE). Their solutions jointly employ information from networks and clients, so as to pace the video flows under diverse channel conditions. With trace-driven simulations over LTE networks, the authors show their solutions indeed reduce the number of stalls and total stall durations. In "On the Effectiveness of Offset Projections for 360-degree Video Streaming" (<https://doi.org/10.1145/3209660>), Zhou et al. consider the emerging 360-degree video

2018 Copyright is held by the owner/author(s).

1551-6857/2018/06-ART60

<https://doi.org/10.1145/3214700>

streaming, where the spherical surface has to be projected into a rectangular video before being compressed. The authors tackle the problem of understanding the implication of offset projection. With offset projection, the spherical surface is distorted by an offset, so that more pixels are encoded in a pixel-concentrated direction. Through real experiments, the authors report various interesting findings, e.g., offset projections give better visual quality as long as the Head-Mounted Display (HMD) viewer's orientation is within 40 or 50 degrees of the pixel-concentrated direction.

The third article is about mesh simplification algorithms for VR, which is titled "Designing and Evaluating a Mesh Simplification Algorithm for Virtual Reality" (<https://doi.org/10.1145/3209661>). In this article, Bahirat et al. study the problem of developing a mesh simplification algorithm designed for VR applications, because existing algorithms are vulnerable to boundaries and non-manifold meshes, which are rather common in 3D models. The authors also show that their solution comes with two advantages: (i) smaller geometric errors due to curvature-based boundary preservation and (ii) more preserved surface properties for each vertex. The resulting open-source implementation of their algorithm produces low-complexity and high-fidelity 3D models, which may be useful to VR practitioners.

The following two articles aim for higher scalability in two different applications: video analytics and media retrieval. In the fourth article, titled "Enabling Live Video Analytics with a Scalable and Privacy-Aware Framework" (<https://doi.org/10.1145/3209659>), Wang et al. consider privacy-preserving real-time video analytics. They leverage edge computing for higher scalability, i.e., they have IoT devices with cameras send videos to nearby small multi-tenant data centers, called cloudlets, for video analytics. Two open-source projects: a face recognition software system, called OpenFace, and a face tracking system, called RTFace, are presented. These two projects are then used to build a privacy-aware IoT ecosystem with many cameras for real-time video analytics. In the fifth article, titled "Prototyping a Web-scale Multimedia Retrieval Service Using Spark" (<https://doi.org/10.1145/3209662>), Gudmundsson et al. experiment with automatically distributed cloud computing frameworks to support scalable media retrieval. Particularly, the authors evaluate the effectiveness of implementing a high-throughput media retrieval service using Spark. Their main contribution is to develop, implement, and evaluate detailed Spark pipelines for index construction and query processing algorithms. Regarding large-scale service, the authors adopt 43 billion Scale-Invariant Feature Transform (SIFT) feature vectors for experiments, which is one of the largest high-dimensional feature vectors in the literature. They also plan to make their Spark pipeline code publicly available.

The special section is closed with two articles related to user behaviors. One article is titled "Characterizing User Behaviors in Mobile Personal Livecast: Towards an Edge Computing-assisted Paradigm" (<https://doi.org/10.1145/3219751>), in which Ma et al. study an emerging multimedia application called mobile personal livecast. In mobile personal livecast, a large number of broadcasters distribute video content of their daily life or special events to many worldwide viewers, and viewers and broadcasters exchange real-time messages during the broadcasts. The authors first analyze the behavior of the mobile personal livecast users using a dataset with 11 million users and report interesting insights like that most livecasts are locally popular. Next, the authors propose to jointly leverage cloud and edge computing to optimize the overall system performance in terms of efficiency and scalability. The other and last article is titled "User Behavior Analysis and Video Popularity Prediction on a Large-Scale VoD System" (<https://doi.org/10.1145/3226035>). In that work, Huang et al. analyze the individuals' viewing behaviors in a large dataset from a VoD (Video-on-Demand) provider in China. Specifically, the daily request patterns (such as frequency of watching videos) and viewing interests (such as eagerness of new content) are carefully analyzed. Based on their analysis, the authors propose a video popularity prediction algorithm based on the Auto-Regressive-Moving-Average (ARMA) model, where video-dependent training

parameters are used. The algorithm is shown to outperform other prediction algorithms, because the proposed algorithm is finer-grained, and trend-, periodicity-, and autocorrelation-aware. Last, the authors demonstrate that their prediction algorithm indeed improves the performance of an Least Frequently Used (LFU) based cache, which is a concrete sample application.

We hope you, as a reader, will enjoy the special section. We would like to sincerely thank the authors of all the articles in this special section for their tremendous efforts in preparing and polishing the manuscripts. We also want to thank the reviewers, award committee members at ACM MMSys'17 and NOSSDAV'17, for carefully selecting the candidates of the extended best paper submissions. Furthermore, we are also grateful to the reviewers of the special section for their valuable and constructive comments throughout the rigorous review process. Without their great help, we wouldn't be able to get the special section ready. Last, we are grateful to ACM TOMM editorial members: Professor Alberto Del Bimbo and Professor Stefano Berretti for all their generous support. Hopefully, this special section will provide some more momentum on further growing the ACM Multimedia Systems community.

Pablo Cesar
Centrum Wiskunde & Informatica

Cheng-Hsin Hsu
National Tsing Hua University

Chun-Ying Huang
National Chiao Tung University

Pan Hui
University of Helsinki and Hong Kong University of Science and Technology

Guest Editors