

SIGMM RECORDS

Table of Contents

- 1 **Volume 5, Issue 2, June 2013 (ISSN 1947-4598)**
- 1 **Editorial**
- 1 **Introducing the new Board of ACM SIG Multimedia**
- 3 **Open Source Column: OpenIMAJ – Intelligent Multimedia Analysis in Java**
- 8 Papers
- 8 **FXPAL hires Dick Bulterman to replace Larry Rowe as President**
- 9 **Call for Bids: ACM Multimedia 2016**
- 10 **MediaEval Multimedia Benchmark: Highlights from the Ongoing 2013 Season**
- 11 **MPEG Column: Press release for the 104th MPEG meeting**
- 13 **PhD Thesis Summaries**
- 13 Johannes Schels
- 13 Tiia Ojanperä
- 14 Tomas Kupka
- 15 Ulrich Lampe
- 16 **Recently published**
- 16 MMSJ Volume 19, Issue 3
- 16 MMSJ Volume 19, Issue 4
- 16 TOMCCAP, Volume 9, Issue 2
- 17 TOMCCAP, Volume 9, Issue 3
- 17 **Job Opportunities**
- 17 PhD position in HDR imaging
- 17 Research Assistants/PhD Positions at University of Vienna
- 18 **Calls for Contribution**
- 18 CFPs: Sponsored by ACM SIGMM
- 18 CFPs: Sponsored by ACM
- 18 CFPs: Not ACM-sponsored
- 19 **Back Matter**
- 19 Notice to Contributing Authors to SIG Newsletters
- 19 Impressum



SIGMM Records

Volume 5, Issue 2, June 2013 (ISSN 1947-4598)

Editorial

Dear Member of the SIGMM Community, welcome to the second issue of the SIGMM Records in 2013.

SIGMM has elected a new board, to guide the SIG through the next couple of years and develop it further. The new board, under the chairmanship of Professor Shih-Fu Chang introduces itself in this issue of the Records.

Among the first acts of the new board was the call for bids for ACM Multimedia 2016, announced also in this issue.

Of course, we have also several other contributions: the OpenSource column introduces OpenIMAJ, while the MPEG column brings the press release for the 104th MPEG meeting. We can also reveal a change of leadership in FXPal, a research company with many SIGMM members in its ranks and a former SIGMM chair as departing president.

We put also a spotlight on the ongoing season for MediaEval, the multimedia benchmarking initiative, and we include four PhD thesis summaries in this issue.

Of course, we include also a variety of calls for contribution. Please give attention to two particular ones: TOMCCAP has chosen its special issue topic for 2014 and includes a call for papers in this issue of the Records. And also MTAP has issued a special issue call for paper.

Last but most certainly not least, you find pointers to the latest issues of TOMCCAP and MMSJ, and several job announcements.

We hope that you enjoy this issue of the Records.

The Editors

Stephan Kopf, Viktor Wendel, Lei Zhang, Pradeep Atrey, Christian Timmerer, Pablo Cesar, Mathias Lux, Carsten Griwodz

Introducing the new Board of ACM SIG Multimedia

Statement of the new Board

As we celebrate the 20th anniversary of the ACM Multimedia conference, we are thrilled to see the rapidly expanding momentum in developing innovative multimedia technologies in both the commercial and academic worlds. These can be easily seen by the continuously growing attendance size at the SIGMM flagship conference (ACMMM) and the action-packed technical demo session contributed by many researchers and practitioners from around the world each year.

However, the community is also faced with a set of new challenges in the coming years:

- how to solidify the intellectual foundation of the SIGMM community and gain broader recognition as a field;
- how to rebuild the strong relations with other communities closely related to multimedia;
- how to broaden the scope of the multimedia conferences and journals and encourage participation of researchers from other disciplines;
- how to develop a simulative environment for nurturing next-generation leaders including a stronger female participation in the community.

To respond to these challenges and explore new opportunities, we will actively explore following initiatives:

- create opportunities for stronger collaboration with relevant fields through joint workshops, tutorials, and special issues;
- recruit active contribution from broad technical, geographical, and organizational areas;

- strengthen the SIGMM brand by establishing community-wide resources such as SIGMM distinguished lecture series, open source software, and online courses;
- expand the mentoring and educational activities for students and minority members.

Chair: Shih-Fu Chang



Shih-Fu Chang

shih.fu.chang@columbia.edu

(**<shih.fu.chang@columbia.edu>**)

Shih-Fu Chang is the Richard Dicker Professor and Director of the Digital Video and Multimedia Lab at Columbia University. He is an active researcher leading development of innovative technologies for multimedia information extraction and retrieval, while contributing to fundamental advances of the fields of machine learning, computer vision, and signal processing. Recognized by many paper awards and citation impacts, his scholarly work set trends in several important areas, such as content-based visual search, compressed-domain video manipulation, image authentication, large-scale high-dimensional data indexing, and semantic video search. He co-led the ADVENT university-industry research consortium with participation of more than 25 industry sponsors. He has received IEEE Signal Processing Society Technical Achievement Award, ACM SIG Multimedia Technical Achievement Award, IEEE Kiyo Tomiyasu Award, Service Recognition Awards from IEEE and ACM, and the Great Teacher Award from the Society of Columbia Graduates. He served as the Editor-in-Chief of the IEEE Signal Processing Magazine (2006-8), Chairman of Columbia Electrical Engineering Department (2007-2010), Senior Vice Dean of Columbia Engineering School (2012-date), and advisor for several companies and research institutes. His research has

been broadly supported by government agencies as well as many industry sponsors. He is a Fellow of IEEE and the American Association for the Advancement of Science.

Vice Chair: Rainer Lienhart



Rainer Lienhart

rainer.lienhart@informatik.uni-augsburg.de

(**<rainer.lienhart@informatik.uni-augsburg.de>**)

Rainer Lienhart is a full professor in the computer science department of the University of Augsburg heading the Multimedia Computing & Computer Vision Lab (MMC Lab). His group is focusing on all aspects of very large-scale image, video, and audio mining algorithms including feature extraction, image/video retrieval, object detection, and humans pose estimation. Rainer Lienhart has been an active contributor to OpenCV. From August 1998 to July 2004 he worked as a Staff Researcher at Intel's Microprocessor Research Lab in Santa Clara, CA, on transforming a network of heterogeneous, distributed computing platforms into an array of audio/video sensors and actuators capable of performing complex DSP tasks such as distributed beamforming, audio rendering, audio/visual tracking, and camera array processing. At the same time, he was also continuing his work on media mining. He is well known for his work in image/video text detection/recognition, commercial detection, face detection, shot and scene detection, automatic video abstraction, and large-scale image retrieval. The scientific work of Rainer Lienhart covers more than 80 refereed publications and more than 20+ patents. He was a general co-chair of ACM Multimedia 2007 and SPIE Storage and Retrieval of Media Databases 2004 & 2005. He serves in the editorial boards of 3 international journals. For more than a decade he is a committee member of ACM Multimedia. Since July 2009 he is the vice chair of SIGMM. He is

also the executive director of the Institute for Computer Science at the University of Augsburg since April 2010.

Director of Conferences: Nicu Sebe



Nicu Sebe

sebe@disi.unitn.it (<sebe@disi.unitn.it>)

Nicu Sebe is Associate Professor with the University of Trento, Italy, where he is leading the research in the areas of multimedia information retrieval and human-computer interaction in computer vision applications. He was involved in the organization of the major conferences and workshops addressing the computer vision and human-centered aspects of multimedia information retrieval, among which as a General Co-Chair of the IEEE Automatic Face and Gesture Recognition Conference, FG 2008, ACM International Conference on Image and Video Retrieval (CIVR) 2007 and 2010. He is the general chair of ACM Multimedia 2013 and a program chair of ECCV 2016 and was a program chair of ACM Multimedia 2011 and 2007. He has been a visiting professor in Beckman Institute, University of Illinois at Urbana-Champaign and in the Electrical Engineering Department, Darmstadt University of Technology, Germany. He is a co-chair of the IEEE Computer Society Task Force on Human-centered Computing and is an associate editor of IEEE Transactions on Multimedia, Computer Vision and Image Understanding, Machine Vision and Applications, Image and Vision Computing, and of Journal of Multimedia.

Open Source Column: OpenIMAJ – Intelligent Multimedia Analysis in Java

Introduction

Multimedia analysis is an exciting and fast-moving research area. Unfortunately, historically there has been a lack of software solutions in a common programming language for performing scalable integrated analysis of all modalities of media (images, videos, audio, text, web-pages, etc). For example, in the image analysis world, OpenCV and Matlab are commonly used by researchers, whilst many common Natural Language Processing tools are built using Java. The lack of coherency between these tools and languages means that it is often difficult to research and develop rational, comprehensible and repeatable software implementations of algorithms for performing multimodal multimedia analysis. These problems are also exacerbated by the lack of any principled software engineering (separation of concerns, minimised code repetition, maintainability, understandability, premature optimisation and over optimisation) often found in research code. OpenIMAJ is a set of libraries and tools for multimedia content analysis and content generation that aims to fill this gap and address the concerns. OpenIMAJ provides a coherent interface to a very broad range of techniques, and contains everything from state-of-the-art computer vision (e.g. SIFT descriptors, salient region detection, face detection and description, etc.) and advanced data clustering and hashing, through to software that performs analysis on the content, layout and structure of webpages. A full list of the all the modules and an overview of their functionalities for the latest OpenIMAJ release can be found here. OpenIMAJ is primarily written in Java and, as such is completely platform independent. The video-capture and hardware libraries contain some native code but Linux, OSX and Windows are supported out of the box (under both 32 and 64 bit JVMs; ARM processors are also supported under Linux). It is possible to write programs that use the libraries in any JVM language that supports Java interoperability, for example Groovy and Scala. OpenIMAJ can even be run on Android phones and tablets. As it's written using Java, you can run any application built using OpenIMAJ on any of the supported platforms without even having to recompile the code.

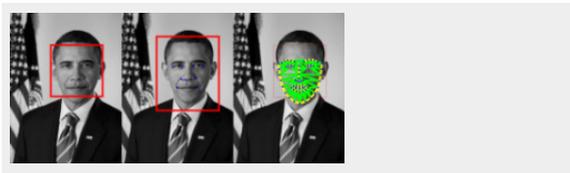
Some simple programming examples

The following code snippets and illustrations aim to give you an idea of what programming with OpenIMAJ is like, whilst showing some of the powerful features.

```
...
// A simple Haar-Cascade face detector
HaarCascadeDetector det1 =
    new HaarCascadeDetector();
DetectedFace face1 =
    det1.detectFaces(img).get(0);
new SimpleDetectedFaceRenderer()
    .drawDetectedFace(mbf, 10, face1);

// Get the facial keypoints
FKEDetector det2 =
    new FKEDetector();
KEDetectedFace face2 =
    det2.detectFaces(img).get(0);
new KEDetectedFaceRenderer()
    .drawDetectedFace(mbf, 10, face2);

// With the CLM Face Model
CLMFaceDetector det3 =
    new CLMFaceDetector();
CLMDetectedFace face3 =
    det3.detectFaces(img).get(0);
new CLMDetectedFaceRenderer()
    .drawDetectedFace(mbf, 10, face3);
...
```



Face detection, keypoint localisation and model fitting

```
...
// Find the features
DoGSIFTEngine eng = new DoGSIFTEngine();
List sourceFeats =
    eng.findFeatures(sourceImage);
List targetFeats =
    eng.findFeatures(targetImage);

// Prepare the matcher
final HomographyModel model =
    new HomographyModel(5f);
final RANSAC ransac = new RANSAC(model,
    1500,
    new RANSAC.BestFitStoppingCondition(),
    true);
ConsistentLocalFeatureMatcher2d matcher
```

```
= new ConsistentLocalFeatureMatcher2d
    (new FastBasicKeypointMatcher(8));

// Match the features
matcher.setFittingModel(ransac);
matcher.setModelFeatures(sourceFeats);
matcher.findMatches(targetFeats);
...
```

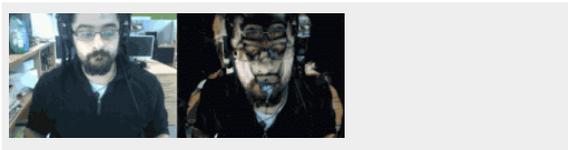


Finding and matching SIFT keypoints

```
...
// Access First Webcam
VideoCapture cap = new VideoCapture(
    640, 480);

//grab a frame
MBFImage last = cap.nextFrame().clone();

// Process Video
VideoDisplay
    .createOffscreenVideoDisplay(cap)
    .addVideoListener(
        new VideoDisplayListener() {
            public void beforeUpdate(MBFImage
                frame) {
                frame.subtractInplace(last).abs();
                last = frame.clone();
            }
        }
    );
...
```



Webcam access and video processing

The OpenIMAJ design philosophy

One of the main goals in the design and implementation of OpenIMAJ was to keep all components as modular as possible, providing a clear separation of concerns whilst maximising code reusability, maintainability and understandability. At the same time, this makes the code easy to use and extend. For example, the OpenIMAJ difference-of-Gaussian SIFT implementation allows different parts of the algorithm to be replaced or modified at will without having to modify the source-code

of the existing components; an example of this is our min-max SIFT implementation [1], which allows more efficient clustering of SIFT features by exploiting the symmetry of features detected at minima and maxima of the scale-space. Implementations of commonly used algorithms are also made as generic as possible; for example, the OpenIMAJ RANSAC implementation works with generic *Model* objects and doesn't care whether the specific model implementation is attempting to fit a homography to a set of point-pair matches or a straight line to samples in a space. Primitive media types in OpenIMAJ are also kept as simple as possible: Images are just an encapsulation of a 2D-arrays of pixels; Videos are just encapsulated iterable collections/streams of images; Audio is just an encapsulated array of samples. The speed of individual algorithms in OpenIMAJ has not been a major development focus, however OpenIMAJ can not be called slow. For example, most of the algorithms implemented in both OpenIMAJ and OpenCV run at similar rates, and things such as SIFT detection and face detection can be run in real-time. Whilst the actual algorithm speed has not been a particular design focus, scalability of the algorithms to massive datasets has. Because OpenIMAJ is written in Java, it is trivial to integrate it with tools for distributed data processing, such as Apache Hadoop. Using the OpenIMAJ Hadoop tools [3] on our small Hadoop cluster, we have extracted and indexed visual term features from datasets with sizes in excess of 50 million images. The OpenIMAJ clustering implementations are able to cluster larger-than-memory datasets by reading data from disk as necessary.

A history of OpenIMAJ

OpenIMAJ was first made public in May 2011, just in time to be entered into the 2011 ACM Multimedia Open-Source Software Competition [2] which it went on to win. OpenIMAJ was not written overnight however. As shown in the following picture, parts of the original codebase came from projects as long ago as 2005. Initially, the features were focused around image analysis, with a concentration on image features used for CBIR (i.e. global histogram features), features for image matching (i.e. SIFT) and simple image classification (i.e. cityscape versus landscape classification).



A visual history of OpenIMAJ

As time went on, the list of features began to grow; firstly with more implementations of image analysis techniques (i.e. connected components, shape analysis, scalable bags-of-visual-words, face detection, etc). This was followed by support for analysing more types of media (video, audio, text, and web-pages), as well as implementations of more general techniques for machine learning and clustering. In addition, support for various hardware devices and video capture was added. Since its initial public release, the community of people and organisations using OpenIMAJ has continued to grow, and includes a number of internationally recognised companies. We also have an active community of people reporting (and helping to fix) any bugs or issues they find, and suggesting new features and improvements. Last summer, we had a single intern working with us, using and developing new features (in particular with respect to text analysis and mining functionality). This summer we're expecting two or three interns who will help us leverage OpenIMAJ in the 2013 MediaEval campaign. From the point-of-view of the software itself, the number of features in OpenIMAJ continues to grow on an almost daily basis. Since the initial release, the core codebase has become much more mature and we've added new features and implementations of algorithms throughout. We've picked a couple of the highlights from the latest release version and the current development version below:

Reference Annotations

As academics we are quite used to the idea of thoroughly referencing the ideas and work of others when we write a paper. Unfortunately, this is not often carried forward to other forms of writing, such as the writing of the code for computer software. Within OpenIMAJ, we implement and expand upon much of our own published work, but also the published work of others. For the 1.1 release of OpenIMAJ we decided that we wanted to make it explicit where the idea for an implementation of each algorithm and technique came from. Rather than haphazardly adding references and citations in the Javadoc comments, we decided that the process of referencing should be more formal, and that the references should be machine readable. These machine-readable references are automatically inserted into the generated documentation, and can also be accessed programmatically. It's even possible to automatically generate a bibliography of all the techniques used by any program built on top of OpenIMAJ. For more information, take a look at this blog post. The reference annotations are part of a bigger framework currently under development that aims to encourage better code development for experimentation purposes. The overall aim of this is to provide the basis for repeatable software implementations of experiments and evaluations, with automatic gathering of the basic

statistics that all experiments should have, together with more specific statistics based on the type of evaluation (i.e. ROC statistics for classification experiments; TREC-style Precision-Recall for information retrieval experiments, etc).

Stream Processing Framework

Processing streaming data is a hot topic currently. We wanted to provide a way in OpenIMAJ to experiment with the analysis of streaming multimedia data (see the description of the “Twitter’s visual pulse” application below for example). The OpenIMAJ `Stream` classes in the development trunk of OpenIMAJ provide a way to effectively gather, consume, process and analyse streams of data. For example, in just a few lines of code it is possible to get and display all the images from the live Twitter sample stream:

```
//construct your twitter api key
TwitterAPIToken token = ...

// Create a twitter dataset instance
// connected to the live twitter sample
// stream
StreamingDataset<Status> dataset =
    new TwitterStreamingDataset(token, 1);

// use the Stream#map() method to
// transform the stream so we get images
dataset
    // process tweet statuses to produce
    // a stream of URLs
    .map(new TwitterLinkExtractor())
    // filter URLs to just get those that
    // are URLs of images
    .map(new ImageURLExtractor())
    // consume the stream and display
    // images
    .forEach(new Operation<URL>() {
        public void perform(URL url) {
            DisplayUtilities.display(
                ImageUtilities.readMBF(url));
        }
    });
```

The stream processing framework handles a lot of the hard-work for you. For example it can optionally drop incoming items if you are unable to consume the stream at a fast enough rate (in this case it will gather statistics about what it’s dropped). In addition to the Twitter live stream, we’ve provided a number of other stream source implementations, including the one based on the Twitter search API and one based on IRC chat. The latter was used to produce a simple visualisation of a world map that shows where current Wikipedia edits are currently happening.

Improved face pipeline

The initial OpenIMAJ release contained some support for face detection and analysis, however, this has been and continues to be improved. The key advantage OpenIMAJ has over other libraries such as OpenCV in this area is that it implements a complete pipeline with the following components:

1. Face Detection
2. Face Alignment
3. Facial Feature Extraction
4. Face Recognition/Classification

Each stage of the pipeline is configurable, and OpenIMAJ contains a number of different algorithm implementations for each stage as well as offering the possibility to easily implement more. The pipeline is designed to allow researchers to focus on a specific area of the pipeline without having to worry about the other components. At the same time, it is fairly easy to modify and evaluate a complete pipeline. In addition to the parts of the recognition pipeline, OpenIMAJ also includes code for tracking faces in videos and comparing the similarity of faces.

Improved audio processing & analysis functionality

When OpenIMAJ was first made public, there was little support for audio processing and analysis beyond playback, resampling and mixing. As OpenIMAJ has matured, the audio analysis components have grown, and now include standard audio feature extractors for things such as Mel-Frequency Cepstrum Coefficients (MFCCs), and higher level analysers for performing tasks such as beat detection, and determining if an audio sample is human speech. In addition, we’ve added a large number of generation, processing and filtering classes for audio signals, and also provided an interface between OpenIMAJ audio objects and the CMU Sphinx speech recognition engine.

Example applications

Every year our research group holds a 2-3 day Hackathon where we stop normal work and form groups to do a mini-project. For the last two years we’ve built applications using OpenIMAJ as the base. We’ve provided a short description together with some links so that you can get an idea of the varied kinds of application OpenIMAJ can be used to rapidly create.

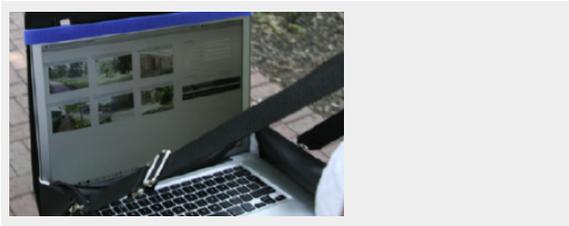
Southampton Goggles

In 2011 we built “Southampton Goggles”. The ultimate aim was to build a geo-localisation/geo-information

system based on content-based matching of images of buildings on the campus taken with a mobile device; the idea was that one could take a photo of a building as a query, and be returned relevant information about that building as a response (i.e. which faculty/school is located in it, whether there are vending machines/ cafe's in the building, the opening times of the building, etc). The project had two parts: the first part was data collection in order to collect and annotate the database of images which we would match against. The second part involved indexing the images, and making the client and server software for the search engine. In order to rapidly collect images of the campus, we built a hand-portable *streetview* like camera device with 6 webcams, a GPS and compass. The software for controlling this used OpenIMAJ to interface with all the hardware and record images, location and direction at regular time intervals. The camera rig and software are shown below:



The Southampton Goggles Capture Rig



The Southampton Goggles Capture Software, built using OpenIMAJ

For the second part of the project, we used the SIFT feature extraction, clustering and quantisation abilities of OpenIMAJ to build visual-term representations of each image, and used our ImageTerrier software [3,4] to build an inverted index which could be efficiently queried. For more information on the project, see this blog post.

Twitter's visual pulse

Last year, we decided that for our mini-project we'd explore the wealth of visual information on Twitter. Specifically we wanted to look at which images were trending based not on counts of repeated URLs, but on the detection of near-duplicate images hosted at different URLs. In order to do this, we used what has now

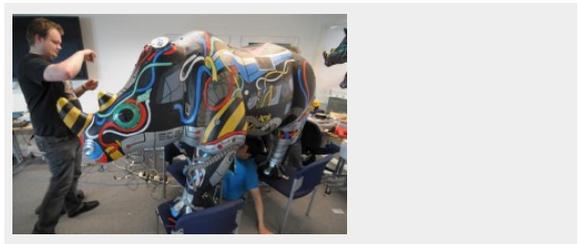
become the OpenIMAJ stream processing framework, described above, to:

1. ingest the Twitter sample stream,
2. process the tweet text to find links,
3. filter out links that weren't images (based on a set of patterns for common image hosting sites),
4. download and resample the images,
5. extract sift features,
6. use locality sensitive hashing to sketch each SIFT feature and store in an ensemble of temporal hash-tables.

This process happens continuously in real-time. At regular intervals, the hash-tables are used to build a duplicates graph, which is then filtered and analysed to find the largest clusters of duplicate images, which are then visualised. OpenIMAJ was used for all the constituent parts of the software: stream processing, feature extraction and LSH. The graph construction and filtering uses the excellent JGraphT library that is integrated into the OpenIMAJ core-math module. For more information on the "Twitter's visual pulse" application, see the paper [5] and this video.

Erica the Rhino

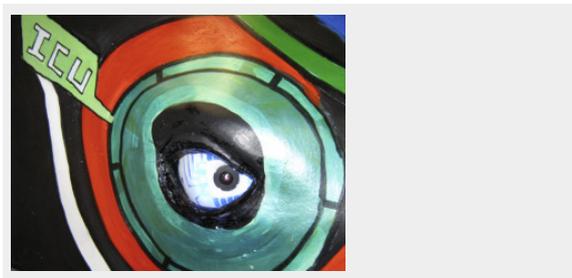
This year, we're involved in a longer-running hackathon activity to build an interactive artwork for a mass public art exhibition called Go! Rhinos that will be held throughout Southampton city centre over the summer. The Go! Rhinos exhibition features a large number of rhino sculptures that will inhabit the streets and shopping centres of Southampton. Our school has sponsored a rhino sculpture called Erica which we've loaded with Raspberry Pi computers, sensors and physical actuators. Erica is still under construction, as shown in the picture below:



Erica, the OpenIMAJ-powered interactive rhino sculpture

OpenIMAJ is being used to provide visual analysis from the webcams that we've installed as eyes in the rhino

sculpture (shown below). Specifically, we're using a Java program built on top of the OpenIMAJ libraries to perform motion analysis, face detection and QR-code recognition. The *rhino-eye* program runs directly on a Raspberry Pi mounted inside the sculpture.



Erica's eye is a servo-mounted webcam, powered by software written using OpenIMAJ and running on a Raspberry Pi

For more information, check out Erica's website and YouTube channel, where you can see a prototype of the OpenIMAJ-powered eye in action.

Conclusions

For software developers, the OpenIMAJ library facilitates the rapid creation of multimedia analysis, indexing, visualisation and content generation tools using state-of-the-art techniques in a coherent programming model. The OpenIMAJ architecture enables scientists and researchers to easily experiment with different techniques, and provides a platform for innovating new solutions to multimedia analysis problems. The OpenIMAJ design philosophy means that building new techniques and algorithms, combining different approaches, and extending and developing existing techniques, are all achievable. We welcome you to come and try OpenIMAJ for your multimedia analysis needs. To get started watch the introductory videos, try the tutorial, and look through some of the examples. If you have any questions, suggestions or comments, then don't hesitate to get in contact.

Acknowledgements

Early work on the software that formed the nucleus of OpenIMAJ was funded by the European Unions 6th Framework Programme, the Engineering and Physical Sciences Research Council, the Arts and Humanities Research Council, Ordnance Survey and the BBC. Current development of the OpenIMAJ software is primarily funded by the European Union Seventh Framework Programme under the ARCOMEM and TrendMiner projects. The initial public releases were also funded by the European Union Seventh Framework Programme under the LivingKnowledge together with

the LiveMemories project, funded by the Autonomous Province of Trento.

Papers

- [1] Hare, Jonathon, Samangoeei, Sina and Lewis, Paul (2011) Efficient clustering and quantisation of SIFT features: Exploiting characteristics of the SIFT descriptor and interest region detectors under image inversion. At The ACM International Conference on Multimedia Retrieval (ICMR 2011), Trento, Italy, 17 - 20 Apr 2011. ACM Press.
- [2] Hare, Jonathon, Samangoeei, Sina and Dupplaw, David (2011) OpenIMAJ and ImageTerrier: Java Libraries and Tools for Scalable Multimedia Analysis and Indexing of Images. At ACM Multimedia 2011, Scottsdale, Arizona, USA, 28 Nov - 01 Dec 2011. ACM, 691-694.
- [3] Hare, Jonathon, Samangoeei, Sina, Dupplaw, David and Lewis, Paul H. (2012) ImageTerrier: an extensible platform for scalable high-performance image retrieval. At ACM International Conference on Multimedia Retrieval (ICMR'12), Hong Kong, HK, 05 - 08 Jun 2012. 8pp.
- [4] Hare, Jonathon, Samangoeei, Sina and Lewis, Paul (2012) Practical scalable image analysis and indexing using Hadoop. Multimedia Tools and Applications, 1-34. (doi:10.1007/s11042-012-1256-0).
- [5] Hare, Jonathon, Samangoeei, Sina, Dupplaw, David and Lewis, Paul. (2013) Twitter's visual pulse. In, the 3rd ACM conference on multimedia retrieval, Dallas, US, 2pp, 297-298. (doi:10.1145/2461466.2461514).

FXPAL hires Dick Bulterman to replace Larry Rowe as President

FX Palo Alto Laboratory announced on July 1, 2013, the hiring of Prof.dr. D.C.A. (Dick) Bulterman to be President and COO. He will replace Dr. Lawrence A. Rowe on October 1, 2013. Dick is currently a Research Group Head of Distributed and Interactive Systems at Centrum Wiskunde & Informatica (CWI) in Amsterdam, The Netherlands. He is also a Full Professor of Computer Science at Vrije Universiteit, Amsterdam. His research interests are multimedia authoring and document processing. His recent research concerns socially-aware multimedia, interactive television, and media analysis. Together with his CWI colleagues, he has won a series of best paper and research awards

within the multimedia community, and has managed a series of successful large-scale European Union projects. He was also the founder and first managing director of Oratrix Development BV.

Dick commented on his appointment: "I am very pleased to be able to join an organization with FXPAL's reputation and scope. Becoming President is a privilege and a challenge. I look forward to working with the rest of the FXPAL team to continue to build a first-class research organization that addresses significant problems in information capture, sharing, archiving and visualization. I especially hope to strengthen ties with academic research teams around the world via joint projects and an active visitor's program in Palo Alto."

Larry Rowe has led FXPAL for the past six years. During his tenure, he significantly upgraded the physical facilities and computing infrastructure and hired six PhD researchers. Working with VP Research Dr. Lynn Wilcox, he led the lab in pursuing a variety of research projects involving distributed collaboration (e.g., virtual and mixed reality systems), image and multimedia processing (e.g., Embedded Media Markers), information access (e.g., TalkMiner lecture video search system and Querium exploratory search system), video authoring and playback, awareness and presence systems (e.g., MyUnity), and mobile/cloud applications. He worked with VP Marketing & Business Development Ram Sriram to develop the Fuji Xerox SkyDesk cloud applications product offering. Larry commented, "it has been an honor to lead FXPAL during this exciting time, particularly the changes brought about by the internet, mobile devices, and cloud computing."

Larry will spend the next several months working on applications using media wall technology to simplify presentations and demonstrations that incorporate multiple windows displaying content from mobile devices and cloud applications.

FXPAL is a leading multimedia and human-computer interface research laboratory established in 1995 by Fuji Xerox Co., Ltd. FXPAL researchers invent technologies that address key issues affecting businesses and society. For more information, visit <http://www.fxpal.com>.

Call for Bids: ACM Multimedia 2016

The ACM Special Interest Group on Multimedia is inviting for bids for holding its flagship "ACM International Conference on Multimedia". The bids are invited for holding the 2016 conference in EUROPE.

The details of the two required bid documents, evaluation procedure and deadlines are explained below (find updates and corrections at http://disi.unitn.it/~sebe/2016_bid.txt). The bids are due on September 2, 2013.

Required Bid Documents

Two documents are required:

1. Bid Proposal: This document outlines all of the details except the budget. The proposal should contain:
 - a. The organizing team: Names and brief bios of General Chairs, Program Chairs and Local Arrangements Chairs. Names and brief bios of at least one chair (out of the two) each for Workshops, Panels, Video, Brave New Ideas, Interactive Arts, Open Source Software Competition, Multimedia Grand Challenge, Tutorials, Doctoral Symposium, Preservation and Technical Demos. It is the responsibility of the General Chairs to obtain consent of all of the proposed team members. Please note that the SIGMM Executive Committee may suggest changes in the team composition for the winning bids. Please make sure that everyone who has been initially contacted understands this.
 - b. The Venue: the details of the proposed conference venue including the location, layout and facilities. The layout should facilitate maximum interaction between the participants. It should provide for the normal required facilities for multimedia presentations including internet access. Please note that the 2016 ACM Multimedia Conference will be held in Europe.
 - c. Accommodation: the bids should indicate a range of accommodations, catering for student, academic and industry attendees with easy as well as quick access to the conference venue. Indicative costs should be provided. Indicative figures for lunches/dinners and local transport costs for the location must be provided.
 - d. Accessibility: the venue should be easily accessible to participants from Americas, Europe and Asia (the primary sources of attendees). Indicative cost of travel from these major destinations should be provided.
 - e. Other aspects:
 - i. commitments from the local government and organizations
 - ii. committed financial and in-kind sponsorships

- iii. institutional support for local arrangement chairs
- iv. conference date in September/October/November which does not clash with any major holidays or other major related conferences
- v. social events to be held with the conference
- vi. possible venue(s) for the TPC Meeting
- vii.any innovations to be brought into the conference
- viii.cultural/scenic/industrial attractions

2. Tentative Budget: The entire cost of holding the conference with realistic estimated figures should be provided. This template budget sheet should be used for this purpose:http://disi.unitn.it/~sebe/ACMMM_Budget_Template.xlsPlease note that the sheet is quite detailed and you may not have all of the information. Please try to fill it as much as possible. All committed sponsorships for conference organization, meals, student subsidy and awards must be highlighted. Please note that estimated registration costs for ACM members, non-members and students will be required for preparing the budget. Estimates of the number of attendees will also be required.

Feedback from ACM Multimedia Steering Committee:

The bid documents will also be submitted to the ACM Multimedia Steering Committee. The feedback of this committee will have to be incorporated in the final submission of the proposal.

Bid Evaluation Procedure

Bids will be evaluated on the basis of:

1. Quality of the Organizing Team (both technical strengths and conference organization experience)
2. Quality of the Venue (facilities and accessibility)
3. Affordability of the Venue (travel, stay and registration)to the participants
4. Viability of the Budget: Since SIGMM fully sponsors this conference and it does not have reserves, the aim is to minimize the probability of making a loss and maximize the chances of making a small surplus.

The winning bid will be decided by the SIGMM Executive Committee by vote.

Bid Submission Procedure

Please up-load the two required documents and any other supplementary material to a web-site. The general chairs then should email the formal intent to host along with the bid documents web-site URL to the SIGMM Chair (sfchang@ee.columbia.edu) and the Director of Conferences (sebe@disi.unitn.it) by Sep 02, 2013.

Time-line

Sep 02, 2013:	Bid URL to be submitted to SIGMM Chair and Director of Conferences
Sep 2013:	Bids open for viewing by SIGMM Executive Committee and ACM Multimedia Steering Committee
Oct 01, 2013:	Feedback from SIGMM Executive Committee and ACM Multimedia Steering Committee made available
Oct 15, 2013:	Bid Documents to be finalized
Oct 15, 2013:	Bids open for viewing by all SIGMM Members
Oct 24, 2013:	10-min Presentation of each Bid at ACM Multimedia 2013
Oct 25, 2013:	Decision by the SIGMM Executive Committee

Please note that there is a separate conference organization procedure which kicks in for the winning bids whose details can be seen at: http://www.acm.org/sigs/volunteer_resources/conference_manual

MediaEval Multimedia Benchmark: Highlights from the Ongoing 2013 Season

MediaEval is an international multimedia benchmarking initiative that offers tasks to the multimedia community that are related to the human and social aspects of multimedia. The focus is on addressing new challenges in the area of multimedia search and indexing that allow researchers to make full use of multimedia techniques that simultaneously exploit multiple modalities.

A series of interesting tasks is currently underway in MediaEval 2013. As every year, the selection of tasks is made using a community-wide survey that gauges what multimedia researchers would find most interesting and useful. A new task to watch closely this year is Search and Hyperlinking of Television Content, which follows on the heels of a very successful pilot last year. The other main tasks running this year are:

- Social Event Detection for Social Multimedia
- Placing: Geo-coordinate Prediction for Social Multimedia
- Violent Scenes Detection in Film (Affect Task)
- Visual Privacy: Preserving Privacy in Surveillance Videos
- Spoken Web Search: Spoken Term Detection for Low Resource Languages

The tagline of the MediaEval Multimedia Benchmark is: "The 'multi' in multimedia: speech, audio, visual content, tags, users, context". This tagline explains the inspiration behind the choice of the Brave New Tasks, which are running for the first time this year. Here, we would like to highlight Question Answering for the Spoken Web, which builds on the Spoken Web Search tasks mentioned above. This task is a joint effort between MediaEval and the Forum for Information Retrieval Evaluation, a India-based information retrieval benchmark. MediaEval believes strongly in collaboration and complementarity between benchmarks and we hope that this task will help us to better understand how joint-tasks should be best designed and coordinated.

The other Brave New Tasks at MediaEval this year are:

- Soundtrack Selection for Commercials a task from MusiClef
- Similar Segments of Social Speech
- Retrieving Diverse Social Images
- Emotion in Music
- Crowdsourcing for Social Multimedia

The MediaEval 2013 season culminates with the MediaEval 2013 workshop, which will take place in Barcelona, Catalunya, Spain, on Friday-Saturday 18-19 October 2013. Note that this is just before ACM Multimedia 2013, which will be held Monday-Friday 21-25 October 2013, also in Barcelona. We are currently working on finalizing the registration site for the workshop and it will open very soon and will be announced on the MediaEval website.

In order to further foster our understanding and appreciation of user-generated multimedia, each year we designate a MediaEval filmmaker to make a YouTube video about the workshop. The MediaEval

2012 workshop video was made by John N.A. Brown and has recently appeared online.

John's decided to focus on the sense of community and the laughter than he observed at the workshop. Interestingly, his focus recalls the work done at MediaEval 2010 on the role of laughter in social video, see: <http://www.youtube.com/watch?v=z1bjXwxkgBs&feature=youtu.be&t=1m29s>

We hope that this video inspires you to join us in Barcelona.

MPEG Column: Press release for the 104th MPEG meeting

Multimedia ecosystem event focuses on a broader scope of MPEG standards

The 104th MPEG meeting was held in Incheon, Korea, from 22 January to 26 April 2013.

MPEG hosts Multimedia Ecosystem 2013 Event

During its 104th meeting, MPEG has hosted the MPEG Multimedia Ecosystem event to raise awareness of MPEG's activities in areas not directly related to compression. In addition to world class standards for compression technologies, MPEG has developed media-related standards that enrich the use of multimedia such as MPEG-M for Multimedia Service Platform Technologies, MPEG-U for Rich Media User Interfaces, and MPEG-V for interfaces between real and virtual worlds. Also, new activities such as MPEG Augmented Reality Application Format, Compact Descriptors for Visual Search, Green MPEG for energy efficient media coding, and MPEG User Description are currently in progress. The event was organized with two sessions including a workshop and demonstrations. The workshop session introduced the seven standards described above while the demonstration session showed 17 products based on these standards.

MPEG issues CfP for Energy-Efficient Media Consumption (Green MPEG)

At the 104th MPEG meeting, MPEG has issued a Call for Proposals (CfP) on energy-efficient media consumption

(Green MPEG) which is available in the public documents section at <http://mpeg.chiariglione.org/>. Green MPEG is envisaged to provide interoperable solutions for energy-efficient media decoding and presentation as well as energy-efficient media encoding based on encoder resources or receiver feedback. The CfP solicits responses that use compact signaling to facilitate reduced consumption from the encoding, decoding and presentation of media content without any degradation in the Quality of Experience (QoE). When power levels are critically low, consumers may prefer to sacrifice their QoE for reduced energy consumption. Green MPEG will provide this capability by allowing energy consumption to be traded off with the QoE. Responses to the call are due at the 105th MPEG meeting in July 2013.

APIs enable access to other MPEG technologies via MXM

The MPEG eXtensible Middleware (MXM) API technology specifications (ISO/IEC 23006-2) have reached the status of International Standard at the 104th MPEG meeting. MXM specifies the means to access individual MPEG tools through standardized APIs and is expected to help the creation of a global market of MXM applications that can run on devices supporting MXM APIs in addition to the other MPEG technologies. The MXM standard should also help the deployment of innovative business models because it will enable the easy design and implementation of media-handling value chains. The standard also provides reference software as open source with a business friendly license. The introductory part of the MXM family of specifications, 23006-1 MXM architecture and technologies, will soon be also freely available on the ISO web site.

MPEG introduces MPEG 101 with multimedia

MPEG has taken a further step toward communicating information about its standards in an easy and user-friendly manner; i.e. MPEG 101 with multimedia. MPEG 101 with multimedia will provide video clips containing overviews of individual standards along with explanations of the benefits that can be achieved by each standard, and will be available from the MPEG web site (<http://mpeg.chiariglione.org/>). During this 104th MPEG meeting, the first video clip on the Unified Speech and Audio Coding (USAC) standard has been prepared. USAC is the newest MPEG Audio standard, which was issued in 2012. It provides performance as good as or better than state-of-the-art codecs that are designed specifically for a single class of content,

such as just speech or just music, and it does so for any content type, such as speech, music or a mix of speech and music. Over its target operating bit rate, 12 kb/s for mono signals through 32 kb/s for stereo signals, USAC provides significantly better performance than the benchmark codecs, and continues to provide better performance as the bitrate is increased to higher rates. MPEG will employ the MPEG 101 with multimedia communication tool to other MPEG standards in near future.

Digging Deeper – How to Contact MPEG

Communicating the large and sometimes complex array of technology that the MPEG Committee has developed is not a simple task. Experts, past and present, have contributed a series of tutorials and vision documents that explain each of these standards individually. The repository is growing with each meeting, so if something you are interested in is not yet there, it may appear shortly – but you should also not hesitate to request it. You can start your MPEG adventure at <http://mpeg.chiariglione.org/>

Further Information

Future MPEG meetings are planned as follows:

- No. 105, Vienna, AT, 29 July – 2 August 2013
- No. 106, Geneva, CH, 28 October – 1 November 2013
- No. 107, San Jose, CA, USA, 13 – 17 January 2014
- No. 108, Valencia, ES, 31 March – 04 April 2014

For further information about MPEG, please contact:
Dr. Leonardo Chiariglione (Convenor of MPEG, Italy)
Via Borgionera, 103
10040 Villar Dora (TO), Italy
Tel: +39 011 935 04 61
leonardo@chiariglione.org

or

Dr. Arianne T. Hinds
Cable Television Laboratories 858
Coal Creek Circle Louisville, Colorado 80027, USA
Tel: +1 303 661 3419
a.hinds@cablelabs.com.

The MPEG homepage also has links to other MPEG pages that are maintained by the MPEG subgroups. It also contains links to public documents that are freely available for download by those who are not MPEG members. Journalists that wish to receive MPEG Press Releases by email should contact Dr. Arianne T. Hinds at a.hinds@cablelabs.com.

PhD Thesis Summaries

Johannes Schels

Object Class Detection Using Part-Based Models Trained from Synthetically Generated Images

Supervisor(s) and Committee member(s): Rainer Lienhart (advisor), Bernhard Möller (committee member), Eckehard Steinbach (committee member)
 URL: <http://opus.bibliothek.uni-augsburg.de/opus4/frontdoor/index/index/docId/2269>



This thesis presents part-based approaches to object class detection in single 2D images, relying on pre-built CAD models as a source of synthetic training data. Part-based models, representing an object class as a deformable constellation of object parts, have demonstrated state-of-the-art results with respect to object class detection. Typically, the majority of part-based approaches rely on real training images of publicly available image data sets and consequently, the positive output of those detectors is restricted to the viewpoints which are represented by those real training images. However, progress in the domain of computer graphics enables the generation of photo-realistic renderings on demand from a database of CAD models, which can serve as training source for learning an object class detection approach. In this thesis, we present part-based object class detection methods which are based on synthetically generated positive training images and real negative training images, thereby combining the advantages of the two domains described above. More specifically, photo-realistic object parts, representing the object class being trained, are learnt in an unsupervised way without

requiring any manual bounding box, object part, or viewpoint annotations during the training process. The established object parts are efficiently combined into an object class detection framework relying on two part-based models with different learning paradigms. In addition, we outline an extension of our detection framework which is able to cope with multiple object classes. The approaches to object class detection are evaluated on standard benchmark data sets and achieve state-of-the-art results with respect to object class detection in single 2D images.

Multimedia Computing & Computer Vision Lab (MMC)
 URL: <http://www.multimedia-computing.de/>

Tiia Ojanperä

Cross-layer Optimized Video Streaming in Heterogeneous Wireless Networks

Supervisor(s) and Committee member(s): Prof. Timo Ojala (supervisor), Prof. Jörg Ott (reviewer), Prof. Christer Åhlund (reviewer), Prof. Gerald C. Maguire Jr. (opponent)
 URL: <http://urn.fi/urn:isbn:9789526201511>
 ISBN: 978-952-62-0151-1



The volume of data transmitted across wireless and mobile networks continues to grow at a rapid rate. Videos already account for most of this data traffic, and their share is expected to grow even bigger in the near future. The thesis presents a novel video service architecture designed to optimise video streaming according to wireless network capacity in order to achieve improved quality and efficient network resource usage. The solution essentially enables Quality of Service (QoS) sensitive video streaming services to take full advantage of the access diversity of heterogeneous wireless networks and to adapt efficiently to available

network capacity. This nevertheless requires support beyond the layered communication architecture of today's Internet as proposed in the thesis.

The proposed architecture relies on end-to-end cross-layer signaling and control for video streaming optimization. The architecture supports extensive context information collection and transfer in heterogeneous networks; thus, enabling the efficient management of video stream adaptation and user terminal mobility in a diverse and dynamic network environment. The thesis also studies and proposes cross-layer enhancements for adaptive video streaming and mobility management functions enabled by the cross-layer architecture. These include cross-layer video adaptation, congestion-triggered handovers, and concurrent utilization of multiple access networks in the video stream transport. For the video adaptation and multipath transmission, the flexible adaptation and transmission capabilities of the novel scalable video coding technology are used. Regarding the mobility management, the proposed solutions essentially enhance the handover decision-making of the Mobile IP protocol to better support QoS-sensitive video streaming. Finally, the thesis takes a holistic view on the application adaptation and mobility management, and proposes a solution for coordinated control of these two operations in order to achieve end-to-end optimization.

The resulting mobile video streaming system architecture and the cross-layer control algorithms are evaluated using network simulations and real prototypes. Based on the results, the proposed mechanisms can be seen to be viable solutions for improving video streaming performance in heterogeneous wireless networks. They require changes in the communication end-points and the access network but support gradual deployment. This allows service providers and operators to select only a subset of the proposed mechanisms for implementation. The results also support the need for cross-layer signaling and control in facilitating efficient management and utilization of heterogeneous wireless networks.

VTT Technical Research Centre of Finland

URL: <http://www.vtt.fi/>

Tomas Kupka

On the HTTP segment streaming potentials and performance improvements

Supervisor(s) and Committee member(s): Pål Halvorsen (advisor), Carsten Griwodz (advisor), Niklas Carlsson

(opponent), Jörg Ott (opponent), Stein Gjessing (opponent)

URL: <http://heim.ifi.uio.no/~paalh/students/TomasKupka-phd.pdf>



Video streaming has gone a long way from its early years in the 90's. Today, the prevailing technique to stream live and video on demand (VoD) content is adaptive HTTP segment streaming as used by the solutions from for example Apple, Microsoft, and Adobe. The reasons are its simple deployment and management. The HTTP infrastructure, including HTTP proxies, caches and in general Content Delivery Networks (CDNs), is already deployed. Furthermore, HTTP is the de facto standard protocol of the Internet and is therefore allowed to pass through most firewalls and Network Address Translation (NAT) devices. The goal of this thesis is to investigate the possible uses of adaptive HTTP segment streaming beyond the classical linear streaming and to look at ways to make HTTP servers dealing with HTTP segment streaming traffic more efficient.

In addition to the deployment and management benefits, the segmentation of video opens new application possibilities. In this thesis, we investigate those first. For example, we demonstrate on the fly creation of custom video playlists containing only content relevant to a user query. Using user surveys, we show, that it not only saves time to automatically get playlists created from relevant video excerpts, but the user experience increases significantly as well.

However, already the basic capabilities of HTTP segment streaming, i.e., streaming of live and on demand video, are very popular and are creating a huge amount of network traffic. Our analysis of logs provided by a Norwegian streaming provider Comoyo indicates that a substantial amount of the traffic data must be served from places other than the origin server. Since

a substantial part of the traffic comes from places other than the origin server, it is important that effective and efficient use of resources not only takes place on the origin server, but also on other, possibly HTTP segment streaming unaware servers.

The HTTP segment streaming unaware servers handle segment streaming data as any other type of web data (HTML pages, images, CSS files, javascript files etc.). It is important to look at how the effectiveness of data delivery from this kind of servers can be improved, because there might be potentially many “off the shelf” servers serving video segments (be it a homemade solution or an HTTP streaming-unaware CDN server). In general, there are three possible places to improve the situation: on the server, in the network and on the client. To improve the situation in the network between the server and the client is generally impossible for a streaming provider. Improving things on the server is possible, but difficult because the serving server might be out of the control of the streaming provider. Best chances are to improve things on the client. Therefore, the major part of this thesis deals with the proposal and evaluation of different modifications to the client-side and only some light modifications to the server-side. In particular, the thesis looks at two types of bottlenecks that can occur. The thesis shows how to deal with a client-side bottleneck using multiple links. In this context, we propose and evaluate a scheduler for partial segment requests. After that, we discuss different techniques on how to deal with a server-side bottleneck with for example different modifications on the transport layer (TCP congestion control variant, TCP Congestion Window (CWND) limitation) and the application layer (the encoding of segments, segment request strategy).

The driving force behind many of these modifications is the on-off traffic pattern that HTTP segment streaming traffic exhibits. The on-off traffic leads in many cases of live streaming to request synchronization as explained in this thesis. The synchronization in turn leads to increased packet loss and hence to a downgrade of throughput, which exhibits itself by decreased segment bitrate, i.e., lower quality of experience. We find that distributing client requests over time by means of a different client request strategy yields good results in terms of quality and the number of clients a server can handle. Other modifications like the limiting of the CWND or using a different congestion control algorithm can also help in many cases.

All in all, this thesis explores the potentials of adaptive HTTP segment streaming beyond the linear video streaming and it explores the possibilities to increase the performance of HTTP segment streaming servers.

Media Performance Group

URL: <http://simula.no/research/mpg>

Ulrich Lampe

Monetary Efficiency in Infrastructure Clouds - Solution Strategies for Workload Distribution and Auction-based Capacity Allocation

Supervisor(s) and Committee member(s): Ralf Steinmetz (Supervisor), Schahram Dustdar (Rapporteur), Klaus Hofmann (Examiner), Helmut Schlaak (Examiner)
URL: <http://d-nb.info/1037291751>
ISBN: 978-3843911016



Since the early days of computing, a vision has been to provide Information Technology services in the form of a utility, just like water, electricity, or telephony. With the advancement of the cloud computing paradigm since the mid-2000s, this vision has been put into realization. Cloud computing builds on and combines multiple existing technologies and paradigms, such as virtualization and Service-Oriented Architecture, to deliver various forms of Information Technology services over the Internet. In this thesis, our focus is on the most elementary class of Information Technology: computing infrastructure. In this context, we examine two important research problems and propose solution strategies, based on the conjoint objective of monetary efficiency.

As the first major contribution, we introduce the so-called Cloud-oriented Workload Distribution Problem (CWDP). This problem concerns the distribution of a workload,

which comprises multiple computational jobs, across leased infrastructure. We assume the position of a cloud user, who aims at cost-minimal deployment under consideration of resource constraints. On the basis of a mathematical optimization model, we propose the exact solution approach CWDP-EXA.KOM. Given its high time complexity, we further propose the heuristic optimization approach CWDP-HEU.KOM, which is complemented by the improvement procedure CWDP-IMP.KOM. The practical applicability and performance of these optimization approaches is demonstrated using a quantitative evaluation, based on realistic data from the cloud computing market.

As the second key contribution, we examine the Equilibrium Price Auction Allocation Problem (EPAAP). This problem refers to the allocation of Virtual Machine instances based on an equilibrium price auction scheme. For that matter, we focus on the role of a cloud provider, who pursues the aim of profit maximization. We formalize the problem as an optimization model, which permits to deduce the exact optimization approach EPAAP-EXA.KOM. We further propose a heuristic optimization approach, named EPAAP-HEU.KOM, and improvement procedure EPAAP-IMP.KOM. All three approaches are thoroughly analyzed through a quantitative evaluation.

Service-oriented Computing, KOM, Tu Darmstadt
URL: <http://www.kom.tu-darmstadt.de/>

Recently published

MMSJ Volume 19, Issue 3

Editor-in-Chief: Thomas Plagemann
URL: <http://www.springer.de/>
Published: June 2013

Special Issue on Network and Systems Support For Games

Guest Editors: Shervin Shirmohammadi, Carsten Griwodz, Grenville Armitage

- Shervin Shirmohammadi, Carsten Griwodz, Grenville Armitage: **Guest editorial for special issue on network and systems support for games**
- Meng Zhu, Alf Inge Wang, Hong Guo: **From 101 to nnn: a review and a classification of computer game architectures**
- Mirko Suznjevic, Maja Matijasevic: **Player behavior and traffic characterization for MMORPGs: a survey**

- A. L. Cricenti, P. A. Branch: **The Ex-Gaussian distribution as a model of first-person shooter game traffic**
- Mirko Suznjevic, Ivana Stupar, Maja Matijasevic: **A model and software architecture for MMORPG traffic generation based on player behavior**
- Amir Yahyavi, Kévin Huguenin, Bettina Kemme: **Interest modeling in games: the case of dead reckoning**
- Cheryl Savery, T. C. Nicholas Graham: **Timelines: simplifying the programming of lag compensation for the next generation of networked games**
- Juan M. Silva, Abdulmotaleb El Saddik: **Exertion interfaces for computer videogames using smartphones as input controllers**
- Daniel Pittman, Chris GauthierDickey: **Match +Guardian: a secure peer-to-peer trading card game protocol**

MMSJ Volume 19, Issue 4

Editor-in-Chief: Thomas Plagemann
URL: <http://www.springer.de/>
Published: July 2013

- Chung-Hua Chu: **Super-resolution image reconstruction for mobile devices**
- Chandan Singh, Pooja Sharma: **Performance analysis of various local and global shape descriptors for image retrieval**
- Zhang Liu, Chaokun Wang, Jianmin Wang, Hao Wang, Yiyuan Bai: **Adaptive music resizing with stretching, cropping and insertion**
- Jun-Bin Yeh, Chung-Hsien Wu, Shi-Xin Mai: **Multiple visual concept discovery using concept-based visual word clustering**

TOMCCAP, Volume 9, Issue 2

Editor-in-Chief: Ralf Steinmetz
URL: <http://tomccap.acm.org/>
Published: May 2013
sponsored by ACM SIGMM

The Transactions on Multimedia Computing, Communication and Applications are the SIGMM's own Transactions. As a service to Records readers, we provide direct links to ACM Digital Library for the papers of the latest TOMCCAP issue.

- Juan M. Silva, Mauricio Orozco, Jongeun Cha, Abdulmotaleb El Saddik, Emil M. Petriu: **Human**

perception of haptic-to-video and haptic-to-audio skew in multimedia applications

- Chidansh A. Bhatt, Pradeep K. Atrey, Mohan S. Kankanhalli: **A reward-and-punishment-based approach for concept detection using adaptive ontology rules**
- Fawaz A. Alsulaiman, Nizar Sakr, Julio J. Valdés, Abdulmotaleb El Saddik: **Identity verification based on handwritten signatures with haptic information using genetic programming**
- Qianni Zhang, Ebroul Izquierdo: **Multifeature analysis and semantic context learning for image classification**
- Zhen Wei Zhao, Sameer Samarth, Wei Tsang Ooi: **Modeling the effect of user interactions on mesh-based P2P VoD streaming systems**
- Yang Yang, Yi Yang, Heng Tao Shen: **Effective transfer tagging from image to video**
- Zhen Wei Zhao, Wei Tsang Ooi: **APRICOD: An access-pattern-driven distributed caching middleware for fast content discovery of noncontinuous media access**

TOMCCAP, Volume 9, Issue 3

Editor-in-Chief: Ralf Steinmetz
 URL: <http://tomccap.acm.org/>
 Published: June 2013
 sponsored by ACM SIGMM

The Transactions on Multimedia Computing, Communication and Applications are the SIGMM's own Transactions. As a service to Records readers, we provide direct links to ACM Digital Library for the papers of the latest TOMCCAP issue.

- Tao Mei, Lin-Xie Tang, Jinhui Tang, Xian-Sheng Hua: **Near-lossless semantic video summarization and its applications to video analysis**
- Oluwakemi A. Ademoye, Gheorghita Ghinea: **Information recall task impact in olfaction-enhanced multimedia**
- Lo-Yao Yeh, Jiun-Long Huang: **A conditional access system with efficient key distribution and revocation for mobile pay-TV systems**
- Ruchira Naskar, Rajat Subhra Chakraborty: **A generalized tamper localization approach for reversible watermarking algorithms**
- Jonathan Doherty, Kevin Curran, Paul Mckeivitt: **A self-similarity approach to repairing large dropouts of streamed music**
- Edmond S. L. Ho, Jacky C. P. Chan, Taku Komura, Howard Leung: **Interactive partner control in close interactions for real-time applications**

Job Opportunities**PhD position in HDR imaging**

High Dynamic Range (HDR) imaging is believed to be the next frontier in imaging similar to transition from gray level to color or 2D to 3D. However, despite many recent developments there is still no standard approach for compression of HDR images and video and no standard way, in which quality is assessed in such imaging modality. The successful PhD candidate will address these shortcomings by determining, via proposed new subjective evaluations and objective metrics, the influence of context and environmental parameters on HDR image and video content. The outcome of these evaluations will be used to design new HDR compression algorithms for images and video sequences. The compression algorithms will be later extended to 3D image and video content.

Applicants should have experience in image/video signal processing and hold a MSc in related field (electrical engineering, mathematics, or computer science). Experience in scientific programming (Matlab/C/C++) is essential. Strong knowledge of video and image compression, HDR imaging, 3D video, and computer vision would be a plus. Applicants should be very well versed in both oral and written English.

Starting date is September 1, 2013.

Employer: Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
 Expiration date: Thursday, August 1, 2013
 More information date: <http://mmspg.epfl.ch/cms/op/edit/lang/en/pid/58408>

Research Assistants/PhD Positions at University of Vienna

The University of Vienna (15 faculties, 4 centers, about 180 fields of study, approx. 9.400 members of staff, approx. 90.000 students) is inviting applications for three (3) positions of Research Assistants (PhD Positions) at the Research Group Multimedia Information Systems (headed by Prof. Dr. Wolfgang Klas) at the Faculty of Computer Science. Details for the job announcement are available at: http://cs.univie.ac.at/uploads/media/OpenPosition_PraeDoc__4198_-_2013-07.pdf

Employer: University of Vienna
 Expiration date: Sunday, September 1, 2013
 More information date: http://cs.univie.ac.at/uploads/media/OpenPosition_PraeDoc__4198_-_2013-07.pdf

Calls for Contribution

CFPs: Sponsored by ACM SIGMM

ACM MMSys 2014

ACM Multimedia Systems Conference

Submission deadline: 27. September 2013
Location: Singapore
Dates: 19. March 2014 -21. March 2014
More information: <http://www.mmsys.org/>
Sponsored by ACM SIGMM

Call for Papers The ACM Multimedia Systems Conference provides a forum for researchers, engineers, and scientist to present and share their latest research findings in multimedia systems. While research about specific aspects of multimedia systems is regularly published in the various proceedings and transactions of the networking, operating system, real-time system, and database communities, MMSys aims ... Read more →

ACM MMSys 2014 Dataset Track

MMSys'14 Call for Papers: Dataset Track

Submission deadline: 20. December 2013
Location: Singapore
Dates: 19. March 2014 -21. March 2014
More information: <http://www.mmsys.org/>
Sponsored by ACM SIGMM

As an integral part of the conference since 2012, the Dataset Track provides an opportunity for researchers and practitioners to make their work available to the multimedia community.

TOMCCAP Special Issue on “Multiple Sensorial (MulSeMedia) Multi-modal Media: Advances and Applications”

Transactions on Multimedia Computing, Communications and Applications

Multiple Sensorial (MulSeMedia) Multi-modal Media: Advances and Applications

Submission deadline: 14. October 2013

Special issue
Sponsored by ACM SIGMM

Multimedia applications have primarily engaged two of the human senses – sight and hearing. With recent advances in computational technology, however it is possible to develop applications that also consider, integrate and synchronize inputs across all senses, including tactile, olfaction, and gustatory. This integration of multiple senses leads to a ... Read more →

CFPs: Sponsored by ACM

ICMR 2014

ACM International Conference on Multimedia Retrieval (ICMR)

Submission deadline: 02. December 2013
Location: Glasgow, UK
Dates: 01. April 2014 -05. April 2014
More information: <http://www.icmr2014.org/>
Sponsored by ACM

The annual ACM International Conference on Multimedia Retrieval offers a great opportunity for exchanging leading-edge multimedia retrieval ideas among researchers, practitioners and other potential users of multimedia retrieval systems.

MMM 2014

The 20th Anniversary International Conference on MultiMedia Modeling (MMM)

Submission deadline: 01. December 2013
Location: Dublin, Ireland
Dates: 08. January 2014 -10. January 2014
More information: <http://mmm2014.org/>
Sponsored by ACM

MMM is a leading international conference for researchers and industry practitioners for sharing new ideas, original research results and practical development experiences from all MMM related areas.

CFPs: Not ACM-sponsored

CVIU Special Issue on “Large Scale Multimedia Semantic Indexing”

Elsevier Computer Vision and Image Understanding

Large Scale Multimedia Semantic Indexing

Submission deadline: 15. August 2013
Special issue
More information: <http://ees.elsevier.com/cviu>

IEEE ISM 2013

IEEE International Symposium on Multimedia (ISM2013)

Submission deadline: 01. August 2013
Location: Anaheim, CA, USA
Dates: 09. December 2013 -11. December 2013
More information: <http://ism.eecs.uci.edu/ISM2013/>
Sponsored by IEEE

MTAP Special Issue on “Content Based Multimedia Indexing”

Multimedia Tools and Applications

Content Based Multimedia Indexing

Submission deadline: 01. September 2013
Special issue
More information: <http://cbmi2013.mik.uni-pannon.hu/index.php/ctp>

MUM 2013

The 12th International Conference on Mobile and Ubiquitous Multimedia (MUM 2013)

Submission deadline: 17. August 2013
Location: Lulea, Sweden
Dates: 02. December 2013 -05. December 2013
More information: <http://www.mum2013.org/>
In cooperation with ACM

NetGames 2013

The 12th Annual Workshop on Network and Systems Support for Games (NetGames)

Submission deadline: 30. August 2013
Location: Denver, CO, USA
Dates: 09. December 2013 -10. December 2013
More information: <http://netgames2013.cs.du.edu/>
In cooperation with ACM SIGMM

Back Matter

Notice to Contributing Authors to SIG Newsletters

By submitting your article for distribution in this Special Interest Group publication, you hereby grant to ACM the following non-exclusive, perpetual, worldwide rights:

- to publish in print on condition of acceptance by the editor
- to digitize and post your article in the electronic version of this publication
- to include the article in the ACM Digital Library and in any Digital Library related services
- to allow users to copy and distribute the article for noncommercial, educational or research purposes

However, as a contributing author, you retain copyright to your article and ACM will refer requests for republication directly to you.

Impressum

Editor-in-Chief

Carsten Griwodz, Simula Research Laboratory

Editors

Stephan Kopf, University of Mannheim
Viktor Wendel, Darmstadt University of Technology
Lei Zhang, Microsoft Research Asia
Pradeep Atrey, University of Winnipeg
Christian Timmerer, Klagenfurt University
Pablo Cesar, CWI
Mathias Lux, Klagenfurt University

