

Opencast Matterhorn 1.1: Reaching New Heights

Christopher A. Brooks^{*}
University of Saskatchewan
Saskatoon, SK, Canada
cab938@mail.usask.ca

Markus Ketterl^{*}
Universität Osnabrück
Osnabrück, Germany
mketterl@uni-
osnabrueck.de

Adam Hochman
University of California at
Berkeley
Berkeley, CA, USA
adam@media.berkeley.edu

Josh Holtzman
University of California at
Berkeley
Berkeley, CA, USA
jholtzman@berkeley.edu

Judy Stern
University of California at
Berkeley
Berkeley, CA, USA
judy@media.berkeley.edu

Tobias Wunden
Eidgenössische Technische
Hochschule Zürich
Zurich, Switzerland
tobias.wunden@id.ethz.ch

ABSTRACT

This paper gives a short overview of the Opencast Matterhorn system. Built by an open community of individuals and institutions, Matterhorn provides a lecture capture platform for both research and production environments. Matterhorn is comprehensive and scalable, and includes components for the acquisition, processing, and playback of content. Matterhorn is licensed under the liberal Educational Community License (ECL 2.0), a flexible OSI approved open source license, and the Opencast community is free for all institutions, corporations, or individuals to join.

Keywords

Human, Social, and Educational Aspects of Multimedia, Multimedia Systems and Middleware

1. INTRODUCTION

The Opencast community is a collaboration of higher education institutions, commercial partners, and individuals working together to explore, create, and document best practices and technologies with respect to rich media in higher education. The principal activity this community has undergone is the development of the Matterhorn product: an open source and enterprise strength platform for media acquisition, processing, and playback for lecture capture.

In 2009 a group of thirteen higher education institutions formed a team of designers, developers, and researchers charged with creating a reliable open source solution for lecture cap-

^{*}Corresponding authors. Other authors listed alphabetically. Corresponding authors are Ph.D. students at their institutions.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM MM 2011 Scottsdale, Arizona, USA
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

ture and content usage. Consisting of members from computer science research groups as well as institutional education and technology infrastructure support groups, the aim of this team was to take lessons learned from deployed home grown systems and build a scalable, accessible, open solution for higher education. Notably, the system was intended to be both a platform for research in multimedia, pedagogy, and education, as well as a production system intended to scale to the media processing needs of the institution, federations of institutions, and tens of thousands of students accessing and using lecture-created content.

Supported by a Mellon and Hewlett Foundation grant totalling \$1.5M, and institutional matching funds of \$2M, development has continued for two years and is currently maintained by over 17 committers from institutions and corporations from around the world. The Opencast community has over 650 members with more than 200 of them interested in the Matterhorn build project¹. More information about Opencast Matterhorn can be found in [5] and by visiting the project webpage at www.opencastproject.org.

2. ARCHITECTURE & SYSTEM

Matterhorn is made up of three subsystems corresponding to content acquisition, processing, and playback, referred to within the project as capture agents, workflows, and engage respectively. In addition to these distinct systems, a variety of user interfaces and API's for administration and choreographing these subsystems are provided.

Matterhorn is built as a J2EE OSGI system, allowing for services to be swapped out dynamically at runtime as needed. API's are provided as REST endpoints, encouraging reuse outside of the Java domain. Matterhorn relies on other open source solutions (*gststreamer*, *ffmpeg*, and *apache felix* to name a few of the more significant frameworks used), and all contributed code is licensed under the OSI approved Educational Community License (ECL) 2.0. No proprietary code is part of Matterhorn, however, several hardware and software solutions (some open source, some proprietary) external to the project have been connected by interested third parties.

¹As there are no membership fees or official membership process these numbers reflect subscriptions to our public mailing lists.

2.1 Content Acquisition (Capture Agents)

The principal method of content acquisition in Matterhorn is the *capture agent*; a hardware device connected to classroom audio/video infrastructure. The reference platform for this agent is based on commodity laptop-style x86 hardware and works with a number of different video capture devices including NTSC/PAL-based cards (such as the Hauppauge and bt787 chipsets), usb webcams, and specialized VGA capture devices (such as the Epiphan line of products). A fully capable agent with both VGA (5 FPS) capture and NTSC/PAL capture (30/25 FPS) costs roughly \$1,500, and several vendors are beginning to offer embedded or value-added solutions to integrate with Matterhorn capture at a similar price point.

The agent polls a central service which provide iCalendar encoded schedules. Fault tolerance is provided through caching and retry timeouts, allowing capture agents to be deployed in environments with intermittent networking such as wireless or mobile cart situations. Schedules identify both the metadata associated with a capture (e.g. course call number, instructor name, etc.) as well as the inputs to capture from. Underlying the scheduling system the agent uses the *gstreamer* framework to acquire, transform, and encode video. The system allows for multiple specialized encoding pipelines to be run at once, and initial support for transforming video to still frames for confidence monitoring during capture is available.

Once captured, video and other artifacts are aggregated and packaged into an extensible structure called a *MediaPackage* and ingested into the Matterhorn workflow system. Integration with the video processing services of Matterhorn is purposefully limited to the *MediaPackage* declarative format; we wish to encourage maximum flexibility when choosing capture agents for different rooms or purposes (HD, research systems, live streaming, etc.), and using a media aggregation format instead of strong codec-specific contracts makes this more tractable.

2.2 Video Processing (Workflows)

The core of Matterhorn is the video processing workflow system. This system uses a service oriented architecture to dispatch jobs to worker nodes which may be specialized for particular services or provide a variety of different general purpose services. Workflows are linear compositions of these services (which may be synchronous or asynchronous) meant to be run in parallel depending upon the infrastructure topology. In this way Matterhorn can be scaled with more computing resources both to accommodate new services as well as to accommodate an increased load or responsiveness. It also holds out the opportunity to leverage the specialization of certain services to certain pieces of hardware; for instance, encoding services could be run across GPU-accelerated infrastructure while metadata services could be run on more general purpose machines.

Content is ingested into the system in one of three ways; through REST APIs as described previously, through an “inbox” metaphor to accommodate legacy applications (e.g. those that may use file share copying, FTP, etc.), or through end-user web interfaces. Once ingested, a workflow is started by the ingestion system which automatically distributes jobs across the platform depending on worker loads. Worker loads are calculated automatically, and new nodes can be enabled or disabled depending on demand .

Services included in the Matterhorn 1.1 distribution include video encoding (using *ffmpeg* or *gstreamer* as desired), metadata generation, scene detection, preview image generation, trimming and captioning (requiring user input from a web-based tools), and text analysis (using *tesseract*). A command line operation handler is also included allowing for end-user integration of other tools as services in a declarative manner.

Matterhorn publishing services for both progressive download and the open source *red5* streaming server are included, though out of the box Matterhorn will work with a variety of third party streaming services. A local search index service using *solr* is also available. Services for distribution to other popular portals such as YouTube.edu and iTunesU are highly requested and are under active development.

2.3 Media Playback (Engage)

Flexible media distribution is a built-in goal of the system, and standards-based aggregation solutions such as *rss* and *atom* feeds are generated by the system to group events. These feeds can further be filtered by search attributes, and enclose a variety of different media types and qualities specified by the workflow definition. This search feeds an ideal solution for integration to external platforms such as campus information systems and learning content management systems.

Accessibility has been a driving factor throughout the development of Matterhorn user interfaces. The media player, called *engage* was designed and developed to be accessible by supporting assistive technology across multiple platforms. The player can be used in different situations; as a standalone application, or embedded inside of other applications such as blogs, wikis or content management systems. The player enables features based on underlying Matterhorn services such as slide segmentation for previews and faster navigation based on video OCR and text analysis. The technology stack for *engage* is HTML and Javascript-based for navigation, and an AJAX bridge that forwards user interaction to an Adobe Flex-based video component is included. This hybrid approach allows to use the strength of plain standard based HTML regarding accessibility (e.g. tab order or keyboard navigation) extended with ARIA roles that describe functionality according to accessibility specifications while maintaining a quality user experience.

The video component provided with Matterhorn uses the *open source media framework (OSMF)* and supports a powerful, reliable, and platform independent video solution that can handle different distribution mechanisms such as progressive or streaming video. It allows for synchronization between different content sources, allowing the display of multiple videos (e.g. both the presentation and presenter video) in the browser at the same time. The video component can be easily replaced depending on client requirements (e.g. HTML 5 for mobile clients).

3. PRODUCTION SYSTEM

Matterhorn is used at institutions across the globe as a lecture capture solution for thousands of students. The open source nature of the product makes measurement of adoption difficult, however both corresponding authors of this document have overseen widespread local deployments of Matterhorn and offer a glimpse here of the effect it has had at their institution. More user testimonials can be found on

the Opencast Matterhorn project website.

3.1 University of Saskatchewan

The University of Saskatchewan piloted Matterhorn 1.0 capture agents in 11 classrooms between September 2010 and April 2011. Custom Matterhorn workflows were created to connect to a legacy media player which had already been integrated into campus infrastructure.

Over this time period 944 events were captured with the majority of them being regular undergraduate lectures (between 1 hour and 4 hours long) along with some graduate lectures, special events, and meetings. A breadth first approach was taken to lecture capture deployment, and a variety of disciplines were recorded including the natural sciences, social sciences, humanities, and professional colleges.

During this two term deployment, 2,639 unique students watched over 1,543 days worth of video, an average of roughly 41 hours per student. Students were restricted to viewing videos from their enrolled classes only, though many students were enrolled in multiple classes at once.

In addition to this, several students at the university have been involved in the customization of Matterhorn through volunteering, summer research positions, professional internships, and graduate theses. The University of Saskatchewan was a mentoring institution in the Undergraduate Capstone Open Source projects (UCOSP)² where five senior undergraduate students from institutions across the country worked on various aspects of Matterhorn for completion of their bachelors degree. Some of these students became and committers on the project, and several of them continue to contribute to the system despite the conclusion of the program.

The University of Saskatchewan is increase its deployment with Matterhorn 1.1 over the summer of 2011 in preparation for the fall academic term.

3.2 Universität Osnabrück

The Universität Osnabrück migrated to Opencast Matterhorn 1.0 for the winter semester 2010/2011. The system was used to record 16 weekly lectures with Matterhorn capture agent hardware and a couple of additional events (e.g. workshops, conferences, talks) with traditional A/V equipment, subsequently ingested for further processing and distribution. In total, more than 300 (typically 1.5 hour long) recordings were processed and distributed using the Matterhorn ecosystem. Several of the recorded lectures were watched by more than 1,000 unique students over a small period of time on modest hardware infrastructure.

Running in full production with the first public 1.0 release in 2010 using a large scale distributed system unveiled the kinds of bugs that only longitudinal testing can identify. By being involved both in development and deployment of Matterhorn, the Universität Osnabrück acted as a sort of real-world litmus test as to the direction Matterhorn needed to be put in to achieve its scalability goals. While not a primary purpose, local experiences and knowledge have been leveraged to aid in undergraduate student learning as well. Several of the committers involved in the project have been instructing senior level software engineering courses focusing on the technologies, techniques, and solutions found in Matterhorn. Students in these courses are encouraged to get involved with the project and contribute to the open source world.

²See <http://ucosp.ca/>

The Universität Osnabrück has since migrated to Matterhorn 1.1 for its current deployment.

4. RESEARCH PLATFORM

One of the powerful aspects of a common platform for lecture capture in higher education is the ability to share not only research results but research innovations. Several institutions are actively developing, deploying, and validating systems based on Matterhorn, and the ease at which these institutions can share those results back with the broader community, allows for greater impact while reducing the costs of applying research in practice. Further, a collaborative production environment encourages a collaborative research program; many of the individuals involved with Matterhorn now have collaborative research agendas adding to the quality and impact of their results.

A sample of some of the research endeavours using Matterhorn that we are familiar with includes:

- **Student interaction modelling through clustering.** Using Matterhorn recorded lectures, clusters of students who share similar viewing patterns were discovered in [1]. Tying these clusters to pedagogical goals and educational strategies can allow for content recommendation or early warning of student problems. This research ties aspects of education and the social sciences to the use of rich media in the classroom.
- **Scene detection using Supervised Machine Learning.** The Matterhorn scene detection service was analysed in [2] against several other scene detection techniques. The workflow operation handler features of Matterhorn make both multiple detection techniques as well as supervised machine learning techniques feasible in large scale implementations if computational resources are available. This technique is being extended to add novel methods of trim point detection and pause detection to Matterhorn.
- **Social navigation support in web lectures.** User awareness has become one of the major features of today's Web 2.0 experience. Understanding how different people are using multimedia content opens up new possibilities for navigation and analysis [4]. Social footprints allow pedagogy and content researchers to track which parts of long video recordings have been watched more frequently by students. This information can then be used for machine learning or also for semantic based recommender systems [3].
- **Automated camera tracking.** A novel research spinoff of the Matterhorn project has been an undergraduate honours thesis work focusing on automated camera tracking. Written primarily by one of the Matterhorn committers, it will soon be released to the Opencast community as an open source solution for tracking lecturers as they move about the classroom.

5. CONCLUSIONS

The Opencast Matterhorn platform is an open source and enterprise quality system for capturing, processing, and engaging with lecture content in higher education. It is a platform that marries the scalability and reliability demands of

production environments with the possibilities made available by cutting edge research. Principle amongst its features is the service oriented approach which allows new image and audio recognition techniques to be plugged in dynamically as needed. Experimentation of new codecs, metadata analysis, and playback tools can all be added without the need for developers to understand (and potentially destabilize) other aspects of the system.

In addition to these services, the platform is built around three coarse grained components; capture agents (for content acquisition), workflows (for content processing), and engage tools (for content playback). Each of these components maintains low coupling with one another, allowing them to be swapped out as institutional or research requirements may dictate. This has been leveraged to quickly add new commercial capture devices to the system, to integrate with legacy processing engines, and to explore device-specific (e.g. mobile) playback systems.

Matterhorn began as an open source solution that was funded by 13 higher education institutions and the Mellon and Hewlett foundations. It is no longer grant funded, and is instead maintained by 17 committers representing various higher education institutions across the globe. Further, an ecosystem of solutions providers around the Matterhorn codebase is starting to emerge, including both hardware and service vendors. Matterhorn offers reduced costs and increased flexibility for service units in the academy, and provides a platform for rapid roll-out of novel research related to audio and video processing, recognition, and pedagogy.

6. ACKNOWLEDGMENTS

We would like to thank the Hewlett and Mellon foundations for their critical support that allowed us to coordinate, plan, execute, and sustain the vision we have for media in higher education.

In addition to the many committers who made this work possible, we would like to thank the following committers emeritus and contributors who played a pivotal role in bringing the product to fruition (in alphabetical order): Stefan Altevogt, Allison Bloodworth, Nils Birnbaum, Johannes Emden, Sebastian Ewert, Mara Hancock, Jamie Hodge, Andre Klassen, Bostjan Pajntar, Olaf Schulte, and Aaron Zeckoski.

7. REFERENCES

- [1] C. A. Brooks, C. Demmans Epp, G. Logan, and J. Greer. The Who, What, When, and Why of Lecture Capture. In *1st International Conference on Learning Analytics and Knowledge 2011 (LAK'11)*, 2010.
- [2] G. S. Johnston, C. A. Brooks, C. Thompson, and J. Greer. Quantifying Accuracy of Algorithms Trained on Multiple Raters in Disagreement. In *In submission to the Conference on Educational Data Mining 2011*, 2011.
- [3] M. Ketterl, J. Emden, and O. Vornberger. Using social navigation for multimedia content suggestion. In *Proceedings of the 2010 IEEE Fourth International Conference on Semantic Computing, ICSC '10*, pages 448–449, Washington, DC, USA, 2010. IEEE Computer Society.
- [4] M. Ketterl, R. Mertens, and O. Vornberger. Bringing web 2.0 to web lectures. *International Journal of Interactive Technology and Smart Education*, 6(2):82–96, 2009.
- [5] M. Ketterl, O. A. Schulte, and A. Hochman. Opencast matterhorn: A community-driven open source software project for producing, managing, and distributing academic video. *Interact. Techn. Smart Edu.*, 7(3):168–180, 2010.

8. ADDITIONAL AUTHORS

The cover page lists authors who are members of the product owner team and executed the vision for Matterhorn 1.0. Additional authors of this work include our numerous committers and designers (listed here alphabetically): Kristofor Amundson, Greg Logan, Kenneth Lui, Adam McKenzie, Denis Meyer, Markus Moormann, Matjaz Rihtar, Ruediger Rolf, Nejc Skofic, Micah Sutton, Rubén Pérez Vázquez, and Benjamin Wulff. Full citation of this work should acknowledge these persons as extended authors.