

MIR WHEN ALL RECORDINGS ARE GONE: RECOMMENDING LIVE MUSIC IN REAL-TIME

Marco Lüthy

Temple University, Japan Campus
marco.luethy@temple.edu

Jean-Julien Aucouturier

Temple University, Japan Campus
aucouturier@gmail.com



Figure 1. Incredulous Londoners catch a glimpse of the Beatles’s unannounced rooftop concert atop the Apple headquarters near Picadilly Circus, on January 30, 1969. Extract from the 1970 film *Let It Be* (dir. M. Lindsay-Hogg).

A context

Since year 2000 when it saw its revenues decrease for the first time in 15 years, the Music Industry has entered crisis. CD shipments in the U.S. were down 20% in 2006, and again down 17.5% in 2007 - something trade groups such as the RIAA readily attribute to illegal file-sharing on the internet, although academic research disagrees[1]. Beyond a “Napster” generation effect, what is under question here is the value of music fixed on a medium. In May 2009, British alternative rock band Coldplay announced that it would give away copies of their new recording to those attending its upcoming tour as “a thank you to our fans”[2]. In July 2007 already, American Funk musician Prince’s *Planet Earth* LP was distributed as a freebie with UK’s tabloid Mail on Sunday - shipping 3 million copies in one day[3]. Increasingly, the *recording* is slipping out of the business model for artists and their labels. That transition is not in favor of video, either: in January 2009, American Industrial Rock band Nine Inch Nails shared with their fans more than 400 Gb of unedited raw concert videos in HD format via their website. A few years ago, such costly material would be edited into a highly profitable DVD, selling several 100,000 units globally. So now where do the revenues come from?

First, live. After his Sunday Mail campaign, Prince went

on to sell out a series of 21 shows in London - in minutes. Live performance cannot easily be pirated: either because it involves tons of lighting and stage rigging, or simply because of the feeling of *being there*, interacting with the artist - not something that can easily be triggered from a mp3 file. Second, beyond ticket and merchandising revenues commonly associated with concert production, artists increasingly build wealth on the more general idea of personal “connection” with their fans [4]. Finding USB keys with free mp3s in some random spots of concert venues, deciphering the text of the band’s T-shirts to find secret website URLs, seeing your own geographical position displayed in real-time on the band’s website as you, and thousands of others, download their songs: being a fan of a band like Nine Inch Nails¹ increasingly resembles the playing of a video game - and not one that consists in buying a mp3 file from a database for \$0.99.

At the same time as the music industry could be forecast to let go of static recordings and transition to revenue models based on live performance, the world of information retrieval also shows signs of transition toward “real-time” information. In a noted interview, Internet Search company Google’s CEO recently regretted his company had no good solution yet to search the millions of real-time updates of social networking sites like Twitter or Facebook[5]. Several start-ups engulfed in the breach in the first months of 2009 - these include search engines like Collecta, Crowd-Eye, Topsy, OneRiot. Real-time information has value to monitor fast phenomena like epidemics, political crisis of public opinion [6], as well as more mundane aspects like customer satisfaction for a product or traffic jams on the way to work.

Beyond the need for faster retrieval, this reveals that media and information itself is becoming as short-lived as a Twitter update: it outdates by the minute. The musical experience of playing Tap Tap Revenge, the most downloaded free game on Apple iPhone in 2008, to the sound of (again) Coldplay or Nine Inch Nails has the timespan of a few minutes, and can never be experienced again identically. So is the surprise revelation of TV show Britain’s Got Talent, Susan Boyle, on 11 April 2009. The value of the Beatles’ 1969

¹ see <http://www.nin.com>

rooftop concert (Figure 1) was to be there, on the instant, spotting it from the street - before the police's arrival to stop it down. Recommending its recording 40 years later is one thing a state-of-art MIR system would probably excel at doing. Recommending it on the instant, so that *I* can make it in time to Piccadilly Circus, is a different feat.

A story

Komuso Tokugawa is a live powerhouse. Since his acclaimed appearance at the nation's premier blues festival, he has been playing constantly more than ten shows per week, selling out every one of them. His fan base is uniquely devoted, often coming to attend shows from all over the world, several times in the same week. Many fans have followed him since the early days, when he could still be found singing in the small pub round the corner. Between shows, fans gather on online forums and blogs to review and comment each of the artist's appearances with astonishing enthusiasm. During his concerts, money tips fly onto the stage constantly, as much as cheers and shouts from a crowd in perfect rapture. Komuso is what every performing pop musician would like to be.

Incidentally, Komuso doesn't exist.

In this world, that is. Komuso is the 3D avatar in the virtual world of Second-Life of Tokyo-based musician Paul Cohen. Introduced in 2003 by California-based Linden Labs Inc., Second-Life (hereafter SL) is a computer-based recreation of the real-world, entirely built by its users - now a thriving community of over 1,400,000 people. As in a video game, real-life users of SL appear in virtual reality as 3D avatars and they can interact with other users' avatars and their environment. Like a small but increasing number of real-life musician, Paul aka Komuso has made performing in Second-Life his quasi-exclusive source of revenue [7].

"Performing virtually has a number of advantages over RL", says Paul Cohen, "instant setup, no travel from home studio to gig, and less wear and tear on gear"². Over the past 3 years, Paul/Komuso has played over 800 concerts in SL, topping at more than 40 a month. None of this music is recorded, but rather streamed live from Paul's studio. For him, going from one concert venue to the next is as simple as clicking on a "teleport" link (a feature of SL, where avatars can fly, duplicate and break numerous other physical laws). The audience can be recruited by a computer-based instant message, and they too can teleport to the concert in seconds. In a mixed-reality event we organized in April 2009 in Tokyo, Paul/Komuso performed live from SL, projected onto a screen in a real-world venue. With no advertising,

² all SL interviews reported hereafter were conducted as part of a student project at Temple University, see <http://sl-music.info>



Figure 2. Second-Life musician Paul Cohen performing live as his avatar Komuso Tokugawa, for a mixed audience both in real-life and virtual reality. Virtual worlds like Second-Life prefigure a music industry based on immediate, real-time performance rather than archived recordings.

his SL studio filled with 50 avatar concert-goers in seconds - word of mouth at the millisecond timescale (Figure 2).

Transposing this to the real-world music industry is vertiginous. Leonie Smith, a jazz singer in real-world and a SL TV host by the name of Paisley Beebe, puts it in a deceptively simple way: "In the real world, you might never get to see your favorite musician live. This is why you buy CDs. In SL however, you can teleport out anytime if you want to catch that musician maybe once a week, maybe 8 times a week". If seeing a given artist live becomes a commodity, who needs to acquire a recording?

One may be skeptical about Second-Life. Experiencing a virtual concert conjures visions of uninspiring 3D polygons, awkward navigation with a mouse and keyboard, and finding one's avatar mistakenly rendered half-way through a pillar or a wall. All this is true, but irrelevant. First, technology is constantly improving, and we'll soon be able to teleport at will in the middle of an ongoing music video of perfect HD quality³. Second, the question here is not whether virtual worlds like SL offer conditions for music that are better or worse than, say, the radio or the mp3 before them. The key question is whether the emerging micro-economy of music in SL is a trend so strong as to shape the future of the music industry, something technology will have to follow rather

³ See e.g. Blue Mars, a photorealistic virtual world under development by K. Hashimoto, the lead animator for the 2001 feature film "Final Fantasy: The Spirits Within"

than create: ubiquitous real-time live exclusive performance replacing the archived recording, at least for a segment of popular music.

This creates the following question for MIR research: How do we recommend music that (1) has never been played before, (2) is being played *now* and (3) will be irrelevant when it's over, in a few minutes?

A prototype

We are currently developing a technical platform to experiment with these new information retrieval needs. Our prototype is based on a communication device, embodied in a Second-Life object, which interfaces with a web server outside of SL (Figure 3). The object allows an artist to post an update (e.g. "I'm playing now") and all his/her followers to receive a real-time notification with a teleport link ("take me there"), not unlike a Second-Life Twitter. The communication with the server is based on a HTTP protocol, and triggered from SL using Linden Lab's custom scripting language. The server runs a Ruby on Rails application, processing requests in interaction with a database, and simulating a push mechanism from server (artist) to client (fan). While this update functionality is already available with the built-in instant messaging (IM) tools of SL, implementing it with a custom object and protocol has many advantages: first, all communications can be traced and logged, which allows us to study e.g. how many fans actually teleport after the notification; second, SL objects can be copied and shared, which allows artists to update an audience with which they don't have direct IM connection (friends of friends); finally, the system can grow to incorporate automatic recommendation of which updates a given user should receive.

We are currently collecting data in a pilot experiment in partnership with SL musician Paul Cohen/Komuso Tokugawa. One unique advantage of SL is that one can easily test a technology in real conditions, i.e. recommending real music to real people and measuring the impact by generating real revenue. Our presentation at the workshop will include our findings on such questions as: how does such a notification device diffuse among fans? How many copies of the object are shared by word-of-mouth, with which network dynamics? How many people choose to teleport to a concert after a notification? Are people who get recommended a live event in this way more likely to tip during the event?

In a second phase of our study, we plan to incorporate automatic recommendation techniques into the system. To this aim too, virtual worlds like SL provide unique opportunities. Information about an avatar's name, location, physical characteristics is readily available through a programming API. More details about e.g. an avatar's possessions or history can be mined with minor modi-




Figure 3. Prototype of a musical communication device in Second-Life, provided by artists to their fan to receive real-time notifications of concerts. Here, the artist unit (green) is activated, and notifications are received asynchronously at each of the fan units (red) - after transiting by an external web-server which records traffic and regulate recommendations.

fications of the client open-source code. This opens the way for recommendation based on e.g. who co-attended a concert, behaviors ("you like to stay at the back") or socio-cultural factors ("people who wear the same gothic outfit are attending this concert"). We will report on our technical experiments trying to gather such information, as well as some preliminary results on how "real-time" recommendation differs from traditional retrieval. And we promise: you'll never miss that rooftop concert again.

1 ACKNOWLEDGMENTS

M.L. is supported by a Temple University CARAS Fellowship.

2 REFERENCES

- [1] Oberholzer-Gee, Felix and Strumpf, Koleman. "The Effect of File Sharing on Record Sales: An Empirical Analysis", in *Journal of Political Economy* 115, no. 1 (February 2007): 1-42.
- [2] "Coldplay Reward Fans With Free Live LP", *Rolling Stone Magazine*, 2009-05-01.
- [3] "Why Prince's Free CD Ploy Worked", *TIME Magazine*, 2007-07-18.
- [4] Kelly, Kevin. "A Thousand True Fans", available: http://kk.org/thetechnium/archives/2008/03/1000_true_fans.php, posted March 2008, retrieved June 19th 2009.
- [5] "Google falling behind Twitter", *The Guardian*, available online: <http://www.guardian.co.uk/business/2009/may/19/google-twitter-partnership>, posted 19 May 2009, retrieved June 19th 2009.
- [6] Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S., and Brilliant, L. "Detecting influenza epidemics using search engine query data", in *Nature*, 2008.
- [7] Aucouturier, J.-J. "Real music for a virtual world", *Japan@Inc Magazine*, n.86, June 2009.