UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF FLORIDA

CASE NO. 11-20427-WILLIAMS/TURNOFF

DISNEY ENTERPRISES, INC.,
TWENTIETH CENTURY FOX FILM CORPORATION,
UNIVERSAL CITY STUDIOS PRODUCTIONS LLLP,
COLUMBIA PICTURES INDUSTRIES, INC., and
WARNER BROS. ENTERTAINMENT INC.,

Plaintiffs,

v.

HOTFILE CORP., ANTON TITOV, and
DOES 1-10.

Defendants.
____/

HOTFILE CORP.,

Counterclaimant,

v.

WARNER BROS. ENTERTAINMENT INC.,

Counterdefendant.
____/

DECLARATION OF DR. ERLING WOLD IN SUPPORT OF PLAINTIFFS' MOTION FOR SUMMARY JUDGMENT

1. My name is Erling Wold and I currently hold the position of Chief Scientist at Audible Magic, Inc. I have a Bachelor's degree in Electrical Engineering 1978 from the California Institute of Technology, and a Master's and PhD in Electrical Engineering and Computer Science from the University of California at Berkeley 1987, where my primary area of study was the development of computer algorithms and

architectures for the analysis of audio and music. My thesis was on the nonlinear parameter estimation of acoustic models, but during my time as a graduate student I also published papers in computer graphics ("Antialiasing Through Stochastic Sampling" SIGGRAPH 1985), applied mathematics ("Fast Fourier Transform Processors Using Gaussian Residue Arithmetic," J. Parallel and Distributed Computing 1985) and VLSI design ("Pipeline and Parallel-Pipeline FFT Processors for VLSI Implementations" IEEE Transactions on Computers 1984).

2. I was a Chief Engineer and member of the research staff at Yamaha Music Technologies from 1988 to 1992 where I primarily developed algorithms for the analysis and synthesis of music (e.g. "Method and apparatus for analyzing and synthesizing a sound by extracting and controlling a sound parameter" US 5536902) but also wrote patents with my colleagues on sensor technology ("Position-based controller for electronic musical instrument" US 5541358), software architecture ("Apparatus and method for linking software modules" US 5386568) and nanotechnology ("Musical tone generating apparatus employing microresonator array" US 5569871). After Yamaha, I was a partner in Muscle Fish LLC for eight years, a consulting group that specialized in the areas of computer music and audio, audio analysis, and signal processing but also general software programming and architecture. While there, my colleagues and I did some of the earliest work on automatic audio classification, similarity and retrieval. Our initial paper in this area ("Content-Based Classification, Search and Retrieval of Audio" IEEE Multimedia 1996) has been widely cited. At the time I also coauthored a number of other papers in audio classification and similarity matching and contributed to several books. For the last fifteen years, I have focused almost entirely on the problem of automatically identifying audio and video.

3. For the last decade, I have been an employee of Audible Magic, a company that offers media file identification services to a broad array of clients in a variety of application areas including advertising, royalty distribution, direct consumer-facing identification on phones and televisions, and copyright compliance. I am a named inventor on a number of granted patents assigned to Audible Magic in its service and product areas ("System for identifying content of digital data" US 8006314; "Method and apparatus for identifying new media content" US 7877438; "Method and apparatus for

cache promotion" EP 1485815B1; "Method and apparatus for creating a unique audio signature" US 7562012; "Method and apparatus for identifying an unknown work" US 7529659 and US 6968337; "Method and article of manufacture for content-based analysis, storage, retrieval and segmentation of audio information" US 5918223). At both Muscle Fish and Audible Magic I was the main developer and designer of their audio and video fingerprinting algorithms. I was also directly involved in the evaluation of a number of other audio and video fingerprinting technologies that were offered to us for licensing. A copy of my CV, including a list of my publications, is attached as Exhibit A. Other than in this case, I have not given testimony, whether by trial or by deposition, in the past four years.

4. I have been asked by the Plaintiffs in this case to provide general background to the Court regarding digital content recognition technology (often referred to as "fingerprinting" technology) and to describe the availability and effectiveness of such technology to websites that host content uploaded by users, such as Hotfile, from 2008 through the present. My opinions and bases therefore are expressed below. In preparing this report, I have relied upon my experience in the industry, including more than a decade's work at Audible Magic and Muscle Fish leading to the patents and papers detailed above, my personal knowledge and testing of the public work of our competitors, and my focused reading of their publications over the years. I have also reviewed a number of academic and engineering papers, which I list in Exhibit B.

5. I am being compensated by the Plaintiffs for my study and testimony in this case at a rate of $250 per hour. If called as an expert at trial, I would testify to the opinions and conclusions expressed in this declaration.

I. Summary.

6. Digital fingerprinting is a method used to identify the content contained within digital copies of media files. It is widely used today in a variety of different commercial settings to take an unknown digital file and identify whether the file contains a representation of a known media asset, such as an audiovisual work (e.g., a motion picture, television program, or music video) or an audio work (e.g., a sound recording or a book on tape). As I describe below, the technology is highly effective, has found commercial acceptance in many contexts, including and beyond copyright enforcement,

and was both commercially available for use by websites that host files uploaded by users and widely used by such websites prior to 2009, when I understand Hotfile began operations.

II. Background on Digital Fingerprinting.

A. Definition of Fingerprinting.

7. Digital fingerprinting is a technique that extracts a set of features from a digital file that serves to represent the file for the purposes of identification. The term is derived from the use of human fingerprints in identification. A "fingerprinting technique" is a combination of a fingerprint and an algorithm that can compare two fingerprints to determine if they match or if they do not. A "fingerprinting system" describes the entire business operation surrounding the technique that makes it viable in the marketplace.

8. A media asset – for example a movie or an audio recording – may have many different digital representations. For example, the original master of *Citizen Kane*, originally on film, may be digitized and then released on a DVD or a Blu-ray disc in digital form, and these discs may be ripped by a computer user to any one of many available digital file formats and encoded using any one of many available audio and video encoders. It may also be broadcast over a TV channel that is imperfect, resulting in added errors and noise, and may be altered intentionally, letterboxed or cropped, sped up or slowed down, and this version might be rerecorded by a viewer. Alternately, a theatrical presentation of the film might be surreptitiously recorded by a theatergoer using a camcorder and then uploaded to her or his computer. All of the resulting digital files will be different in size and content, but an individual viewing them all, for example by using a movie player on a computer, would immediately agree that they were all representations of *Citizen Kane*.

9. An effective fingerprinting technique is one that, like a human viewer, can reasonably identify the digital representations above as representations of the movie *Citizen Kane*. At the same time, it must be able to distinguish this collection of representations from the collection of representations of any other different media asset, e.g. *Gone with the Wind.*

B. Relationship to Other Engineering and Scientific Disciplines.

10. Although the use of digital fingerprints for identification of media files is a more recent application, it grows out of the established and broader discipline of classification and recognition of audio and video, and as such relies on many decades of work that have come before it. Media file fingerprinting technologies use many of the same techniques, algorithms, scientific knowledge and mathematical methods that have been developed for the broader subject. Fingerprinting is narrower in scope than other problems in this area that have also seen mature, robust and widely commercialized applications: e.g. speech recognition, now on mobile phones and answering machines; and computer vision, used in applications ranging from the detection of handwritten postal codes in letter sorters to space exploration.

11. Identification through digital fingerprinting is a special case of the broader field of pattern recognition, a general mathematical and technological framework for classifying and clustering data of all types, including text, DNA sequences, radar and sonar detection, credit scoring, and medical diagnosis. Again, the same techniques, algorithms and mathematics used in these applications are commonly used in fingerprinting. In many of these applications, the same issues arise, namely extracting a small and tractable set of pertinent features from the data, reducing the dimensions of that data, ignoring features in the original data that are unimportant, and deriving efficient and reliable algorithms for classification.

C. Features Used by Fingerprinting Systems.

12. Given a digital representation of a media asset, a fingerprinting technique reduces this file to a set of features, typically a set of numerical values. As with human fingerprints, which are reduced to minutiae as well as other features, the media file features chosen must capture those aspects of the media file that identify it as what it is, regardless of the actual bits contained in the media file. Some identification techniques use features that are derived from knowledge of human hearing and vision, both in terms of the physical sensors in the ear and eye and the way the brain emphasizes some aspects of the sound or image or motion over others. Some identification techniques model the mathematical transformations that occur when the original media asset is converted into a particular representation and optimize their features to be those that are present

throughout this process and that change the least when subjected to such transformations.

13. However they are derived, all possible feature sets must capture the aspects of the original media asset that are common to all representations. As an example, a black-and-white version of a color movie still has the same shapes in the same relative positions on the screen, and the scene changes occur in the same time locations. Thus a feature set that captured those particular qualities would be robust in the presence of alterations of color.

14. Many of the features used in fingerprinting systems have been used in computer speech, audio and vision systems for many years, in some cases many decades. During that period, their efficacy in pattern recognition and classification systems has been tested and peer-reviewed.

D. Matching Unknown Digital Files Against Known Media Assets.

15. To discover if an unknown media file is a representation of a particular media asset, fingerprinting systems first convert both the unknown file and the original master media asset to fingerprints using the feature extraction methods above. Once the two fingerprints are available, the system compares them to each other to see whether they match. This can be accomplished through a variety of techniques, but in practice, the exact method used is not that critical. If the two fingerprints are close enough to each other, a match is declared. If they are dissimilar, the system returns no match.

E. Reference Fingerprint Database.

16. For a fingerprinting system to be commercially viable for applications that require the identification of unknown files, the vendor of the fingerprinting system must also maintain a large database of reference fingerprints. If the fingerprint of a particular master media asset is not in the database, the fingerprinting system will fail to identify any digital representations of that original asset. Audible Magic, for instance, has agreements with many copyright holders to generate fingerprints for all their media assets as part of their production or maintenance processes. Other fingerprinting vendors have simply purchased CDs and DVDs and have computed fingerprints directly from those, or fingerprinted off live broadcast feeds. Current vendors in the marketplace have millions or tens of millions of fingerprints in their databases. As well as containing the fingerprints, these databases have other information about the original media asset,

including title, copyright holder, industry identifications like International Standard Recording Code ("ISRC") numbers, and so on.

F. Example Process.

17. The following is a description of the process followed by a media file fingerprint identification system. Say a user of a fingerprinting system had an unknown media file and wanted to know what it is. The user would first run the feature extraction software supplied by the identification service on this file to produce a fingerprint. This small package of information would then be transmitted, say over the Internet, to a computer maintained by the identification service. The computer would contain the reference database above and the software that implements the matching algorithm. The computer would compare the incoming fingerprint to all the fingerprints in the database. The system would then report back to us whether there was a match and, if there was, what was matched.

G. Low Incidence of Errors.

18. Fingerprinting applications are highly accurate, and potential sources of errors can be accounted for in designing and implementing a fingerprinting system. There are two types of errors that can occur in a fingerprint identification system, one or the other of which may matter more in a particular application. False negatives occur when the system fails to identify an unknown that should match something in the database. False positives occur when the system reports a match for something that actually should not have been matched.

19. Note that is possible for such errors to occur due to human error, for example, an incorrect title attached to the fingerprint. What is more important is the fundamental error rate of the low-level algorithms themselves. In modern fingerprinting solutions, both the false negative and false positive rates are extremely low. Based on my experience using fingerprinting algorithms to identify media files, one can achieve false negative rates that are much less than one in a hundred thousand files and false positive rates that are much less than one in a million files. One can trade off these two error rates against each other and against other performance requirements such as the cost of identification.

III. Scientific Background.

A. Basis for Choices of Features.

20. Fingerprinting methods – both audio and video – rely on features of a work that distinguish it from others. As discussed above, in audio fingerprinting, many approaches use features that are based on psychoacoustic models. The cochlea is to first approximation a filter bank, i.e. a spectrum analyzer. Most features used in fingerprinting are features derived from the spectrum of the signal, including spectral peaks (which are very robust in the presence of noise), quantities known as "mel-filtered cepstral coefficients" (which come from speech research and model the overall shape of the spectrum), and "filter bank coefficients" that quantify the amount of signal in each spectral range. As these parameters have psychoacoustic analogs, and as they are derived from nature, they are the natural fingerprint of a sound. In addition, many of the aspects of the sound that are ignored by such features are those that are known from audio and speech research to also be ignored by the ear-brain system, e.g. those that are psychoacoustically masked.

21. Similarly, for use in image and video identification, one can look to eye-brain models to find meaningful features. For example, the eye does not see all possible frequencies of light and has limitations on its ability to detect spatial and temporal changes. These considerations allow developers of a fingerprinting system to ignore much of the data in a video stream and to concentrate on those features that are most important.

22. However, it is not necessary for features to follow psychophysical parameters. Some methods simply use prominent features of the signal, features that have been experimentally or mathematically determined to remain more or less constant even in the presence of the modifications. This is again analogous to a human fingerprint, where identification relies on such features as dots and bifurcations and line terminations, prominent attributes that will survive the transfer of oils to a surface.

B. Basis for Matching Algorithms.

23. As discussed previously, the exact matching algorithm is typically not critical if one has a well-designed set of robust features. All that is important in determining whether two fingerprints match is to determine the similarity of the two sets of features. There are standard and statistically well-understood techniques from the field

of pattern matching that can be applied. For example, if a fingerprint is a set of numbers, one can measure the distance between the two sets using a standard Euclidean distance, just like a normal distance on a map or in space, but generalized to the case of many more dimensions. With such a distance measure, a match is declared if the distance is closer than a particular threshold and no match is returned if the distance is greater.

IV. History of Fingerprinting Techniques.

A. Early Use.

24. Automated media recognition systems have a history extending back at least to the early 1980s, when Broadcast Data Systems patented a solution for audio broadcast monitoring using filter banks. Given the state of technology at the time, only audio detection was possible, and the technique was simple and probably error-prone. However, even this early technique was based on a simple model of the cochlea as a bank of frequency-sensitive filters.

B. Copyright Enforcement in Peer-to-Peer Context.

25. When Napster, KaZaa, BitTorrent and other file sharing systems started to become popular for music sharing in 1999 and beyond, there was a great deal of research and interest in both the academy and industry to develop audio identification systems which could be used for copyright compliance. A number of systems were developed which worked very well, and iMesh, one of the peer-to-peer systems, began using Audible Magic's system in 2005 to identify unauthorized sound recordings. Due to the effectiveness of the systems developed during this period, they quickly found other application niches. One of the most famous of these is the Shazam music identification service, which allows cell phone users to identify a piece of music being played even in a noisy environment like a club or bar, and to purchase it if desired.

C. Use By Web 2.0 Companies.

26. Fingerprinting systems have also been used for several years to identify unknown media assets (including both audio and video) on websites that host content uploaded by users (often called "Web 2.0 companies"). This includes household names such as YouTube, Facebook, MySpace and others. Initially, such sites agreed to use fingerprint identification systems to block or notify users when copyrighted material was uploaded. However, what is most interesting is that the companies listed above, realizing

how well the systems worked, quickly began to use fingerprint-based identifications for their own business purposes, beyond copyright filtering, including ad placement and to sell media assets directly to their users.

27. Fingerprinting systems were commercially available to file-sharing sites well before 2009, when I understand that Hotfile commenced operations. For instance, by 2008, Audible Magic had a combined video and audio fingerprinting system that was utilized by all three of the companies named above, as well as many others,[1] and other companies were offering fingerprinting systems as well.

D. Further Commercial Applications Beyond Copyright.

28. Over the years, such systems have improved in many ways, not in the least due to the rapid increase in hardware performance, which has allowed standard pattern recognition algorithms to be applied to audio and video in real time. Automated broadcast monitoring has continued to be an application area worldwide, now used for both radio and television for royalty distribution and verification of ad placement.

29. It is important to note that the development of digital fingerprinting was not initially motivated primarily by copyright compliance. Media file fingerprinting systems are relied on to pay royalties, to identify what someone is watching or listening to for statistical tracking purposes, or to present metadata and other information that a user desires to know. There has been broad uptake of fingerprint-based identification solutions across the media application landscape.

30. More recently, the growth of the tablet market and the connected television market has spawned another burst of development in the area. A number of television producers have applications on tablets and phones that listen to what is playing on the television, identify it using fingerprinting techniques, and allow users to interact with the show and other fans. Television manufacturers are beginning to incorporate identification systems in the televisions themselves, allowing similar applications, plus commercial detection, coupon promotions, and more. All of this speaks to the maturity

---

[1] This list is by no means exhaustive. While Facebook, YouTube, and MySpace are household names, Audible Magic's technology was used by many websites, big and small, that hosted user-uploaded media files during this time period, including DailyMotion, Break.com, Veoh, Bebo, Crackle, Microsoft Soapbox, Dada, eSnips, Eyespot, and GoFish.

and reliability of fingerprint-based media identification technology.

31. The number of industries using digital fingerprinting is large, from those built around applications on cell phones and televisions and tablets to computer media playback devices and social networking sites. It is a mature industry, used by many companies to process billions of identifications per year, and the marketplace has already weeded out those vendors that do not provide working solutions. Digital fingerprinting has also been developed internally by companies who use it to improve their systems and to increase revenue. The cost / benefit tradeoffs of digital fingerprinting have been in the hands of the marketplace for many years, and the marketplace has chosen the viable techniques over others.

V. Fingerprinting in Practice and in the Marketplace.

32. In practice, digital fingerprinting systems work very well. There has been a great flowering of applications in response to their efficacy. Identification systems have many motivations, not just copyright compliance, but also by the desire for device manufacturers and device users to know what is playing or being stored on their systems. Audible Magic conducts roughly a billion identifications every year, and Shazam last year reported having one hundred million users of its cell-phone identification systems.

33. In practice, the techniques are straightforward. Digital fingerprinting systems have become the stuff of student projects. Digital fingerprinting identification has become a commodity service, with many competing vendors, and even open source projects given away for free and maintained by volunteers. The vendor companies are now competing primarily in application features, in cost, and in ease-of-use. The areas of continuing research are in pushing the limits of the algorithms, for example, how well they function in extremely noisy or visually distracting environments, or at identifying extremely short portions of the original, or at providing the service on very small and less powerful embedded devices. The basic task – identifying an unknown media asset contained in a digital file, hosted on a server, that has not been substantially distorted, and where the underlying work is contained within a vendor's reference database – is now, and has for several years been, a well-understood task that can be performed accurately and reliably.

34. Moreover, as stated above, such fingerprinting systems for identifying

unknown media assets uploaded to file-sharing websites have been available and in widespread use for several years. I understand that Hotfile claims that it is now using a fingerprinting system (from Vobile) to identify files uploaded to its service. I am not aware of any technical reason why Hotfile could not have utilized any of the commercially available fingerprinting systems to carry out this same function back in 2010 or 2009.[2]

[2] I understand that Hotfile has suggested that earlier use of a fingerprinting system would have been complicated by the fact that many files on Hotfile are archived or compressed into file formats that do not allow the media content to be accessed directly without further processing to decompress or unarchive the file. In my opinion, this is not a meaningful barrier to the use of a fingerprinting system. One of Audible Magic's products (its "CopySense Appliance," used by many universities to prevent unauthorized transfers of copyrighted material on their networks) can decompress and/or unarchive files prior to extracting a fingerprint from the underlying media asset, and has had this ability for years. This is not a technically challenging process.

I declare under penalty of perjury that the foregoing is true and correct.

Executed in San Francisco, California this 15th day of February, 2012.

Erling H. Wold, PhD

Exhibit A

ERLING H. WOLD, PhD

629 Wisconsin St
San Francisco, CA 94107
(415) 902 9653
erling@erlingwold.com

PROFESSIONAL

2000-present AUDIBLE MAGIC, INC.
Chief Scientist
Developed and patented algorithms and software for audio and video identification, identification of live television and internet audio and video broadcasts, antipiracy, reassembly and identification of media files on computer networks, audio classification for security and a variety of other projects for Audible Magic's product line as well as consulting projects for clients.

1992-2000 MUSCLE FISH MULTIMEDIA ENGINEERING.
Partner
Founding partner of consulting firm specializing in media related software design and implementation. Developed algorithms for audio analysis, classification, processing, identification and similarity search that were licensed to a number of media, computer and music software companies. Worked on audio and music software contracts for a wide variety of clients. Developed the SoundFisher™ sound effects browser. Was the technical lead of a large project at Sun Microsystems to develop a new audio infrastructure.

1988-1992 YAMAHA MUSIC TECHNOLOGIES USA, INC.
Chief Engineer
Led a three-person team on the development of a new music analysis/synthesis technique. Developed new synthesis techniques. Designed and Implemented a C++ class library for communication between software modules. Led all of the above through the patenting process. Designed and implemented a C++ class library for real-time scheduling of multi-media events. Gave demonstrations of these projects to Japanese management. Developed many projects and proposals for projects and which were presented to Japanese management.

1987-1988 UNIT PRODUCTIONS
Partner
Formed firm with Mark Dippé, later at Industrial Light and Magic. We developed computer video and graphics software and equipment, including 3-D rendering, paint systems, peripheral drivers and color correction of scanned slides.

1981-1987 UNIVERSITY OF CALIFORNIA, BERKELEY
Research Assistant
Work dealt with Ph.D. and M.S. topics described below, and on stochastic sampling algorithms for generating antialiased computer graphics and music.

1979-1981 NORTH STAR COMPUTERS, INC.
Design Engineer
Developed analog read/write electronics for floppy disk system. Developed test software and hardware. Designed alternate user interface devices.

1978 PERSCI, INC.
Design Engineer
Designed voice coil motor servo for an 8 inch floppy disk drive.

EDUCATION

1981-1987 UNIVERSITY OF CALIFORNIA, BERKELEY
Ph.D. in EECS. Ph.D. title: *Nonlinear Parameter Estimation of Acoustic Models.* M.S. in EECS. M.S. title: *FFT Structures for Integrated Circuit Implementation.*
Concentrated on signal processing, control, parameter estimation, nonlinear mathematics, residue arithmetic, computer architecture, VLSI design and computer music. Samuel Silver Award.

1985 UNIVERSITY OF CALIFORNIA, BERKELEY
Studied music composition with Gerard Grisey.

1983-1984 STANFORD UNIVERSITY
Studied computer music at CCRMA with John Chowning.

1978 UNIVERSITY OF CALIFORNIA, BERKELEY
Studied music composition with Andrew Imbrie.

1974-1978 CALIFORNIA INSTITUTE OF TECHNOLOGY, PASADENA, CA
B.S. in EECS. Emphasized applied mathematics and circuit design.

1976-1978 OCCIDENTAL COLLEGE, LOS ANGELES, CA
Studied music composition with Richard Grayson and Robert Gross.
Eleanor Remick Warren Award.

PUBLICATIONS

2004 Doug Keislar, Erling Wold and Thom Blum, *Audio Fingerprints: Technology and Applications*, Proceedings of the 117th Convention of the Audio Engineering Society, San Francisco, CA, USA.

1999 Doug Keislar, Thom Blum, James Wheaton and Erling Wold, *A Content-Aware Sound Browser*, Proceedings of the 1999 International Computer Music Conference.

1999 Erling Wold, Thom Blum, Douglas Keislar and James Wheaton, *Classification, Search and Retrieval of Audio*, in CRC Handbook of Multimedia Computing.

1996 Erling Wold, Thom Blum, Douglas Keislar and James Wheaton, *Content-Based Classification, Search and Retrieval of Audio*, IEEE Multimedia 3(3).

1995 Doug Keislar, Thom Blum, James Wheaton and Erling Wold, *Audio Analysis for Content-Based Retrieval*, Proceedings of the 1995 International Computer Music Conference.

1995 Thom Blum, Doug Keislar, James Wheaton and Erling Wold, *Audio Databases with Content-Based Retrieval*, 1995 International Joint Conference on Artificial Intelligence.

1992 Erling Wold, *Crash*, Leonardo Music Journal, Vol.1, No. 1, pp. 98-99.

1992 Mark Dippé, Erling Wold, *Stochastic Sampling, Theory and Application*, in *Progress in Computer Graphics,* Ablex Publishing, Norwood, N.J.

1987 Erling Wold, Kim Pépard, Comments on *Stochastic Sampling in Computer Graphics* (by Robert L. Cook), in *ACM Transactions on Graphics.*

1987 Erling Wold, *Nonlinear Parameter Estimation of Acoustic Models*, UCB/CSD Report No. 87/354, Computer Science Division, U.C. Berkeley, Berkeley, CA.

1986 Erling Wold, Al Despain, *Parameter Estimation of Acoustic Models: Audio Signal Separation,* Proceedings 1986 IEEE ASSP Workshop on Applications of Signal Processing to Audio and Acoustics, New Paltz, NY.

1985 Erling Wold, Mark Dippé, *Alias-Free Sound Synthesis by Stochastic Sampling,* Proceedings of the International Computer Music Conference, Vancouver, BC.

1985 Mark Dippé, Erling Wold, *Antialising Through Stochastic Sampling*, ACM SIGGRAPH Proceedings, San Francisco, CA.

1985 A. Despain, A. Peterson, O. Rothaus, E. Wold, *Fast Fourier Transform Processors Using Gaussian Residue Arithmetic*, Journal of Parallel and Distributed Computing, Vol. 2, pp. 219-237.

1984 E.H. Wold, A. M. Despain, *Pipeline and Parallel-Pipeline FFT Processors for VLSI Implementations*, IEEE Transactions on Computing, Vol. C-33, No. 5.

1982 A. Despain, C. Sequin, C. Thompson, E. Wold, D. Lioupis, *VLSI Implementation of Digital Fourier Transforms*, UCB/CSD Report No. 82/111, Computer Science Division, U.C. Berkeley, Berkeley, CA.

PATENTS

2011 US 8,006,314 System for identifying content of digital data

2011 US 7,877,438 Method and apparatus for identifying new media content

2009 US 7,562,012 Method and apparatus for creating a unique audio signature

2009 EP 1 485 815 Method and apparatus for cache promotion

2009	US 7,529,659 Method and apparatus for identifying an unknown work
2005	US 6,968,337 Method and apparatus for identifying an unknown work
1999	US 5,918,223 Method and article of manufacture for content-based analysis, storage, retrieval, and segmentation of audio information
1996	US 5,569,871 Musical tone generating apparatus employing microresonator array
1996	US 5,541,358 Position-based controller for electronic musical instrument
1996	US 5,536,902 Method of and apparatus for analyzing and synthesizing a sound by extracting and controlling a sound parameter
1995	US 5,386,568 Apparatus and method for linking software modules

AFFILIATIONS
Executive Director and cofounder, San Francisco Composers Chamber Orchestra.
Board member, Intersection for the Arts, 2004-2006.
Composer-in-residence, ODC Theater, from 2000 to 2003.
Steering Committee, American Composers Forum, SF Bay Area Chapter, 1997-2000

Exhibit B

Exhibit B

Materials Reviewed

"Distortion Discriminant Analysis for Audio Fingerprinting," published in the 2003 IEEE Transactions on Speech and Audio Processing, a respected, peer-reviewed journal. This paper shows the development of a set of audio fingerprint features using mathematical models for the expected audio distortions in the application area. It demonstrates experimental methods for evaluating the efficacy of the system, and develops a mathematical model for the probability of identification error.

"Extracting Noise-Robust Features from Audio Data," published in the 2002 Proceedings of the IEEE Conference on Acoustics, Speech and Signal Processing, a major research conference. Similar to the previous paper, it shows experimental methods applied to estimate error rates.

"Robust Video Fingerprinting for Content-Based Video Identification," published in the 2008 IEEE Transactions on Circuits and Systems for Video Technology, another peer-reviewed journal. This paper develops a mathematical model of the fingerprint features, describes an experimental setup to measure feature robustness and reports the results of these experiments. The authors experimentally compare their chosen feature set to a variety of other published video features. It also shows how false-positive and false-negative rates can be traded off against each other in a typical system.

"A Highly Robust Audio Fingerprinting System," published in the 2002 Proceedings of the International Society for Music Information Retrieval, an international group that has for more than 10 years specialized in problems in music and audio identification, retrieval and analysis. The paper describes experiments and results for testing the robustness of the features in the presence of different distortions.

"A robust image fingerprinting system using the Radon transform," published in Signal Processing: Image Communication 2004. It shows a mathematical model used to derive a set of features for image identification, a basic building block of video identification. The paper demonstrates the robustness of the features experimentally.

"Perceptual Audio Hashing Function," published in EURASIP Journal on Applied Signal Processing 2005, another peer-reviewed journal. The paper shows experimental results for both false negative and false positive error rates.

"Automatic Song Identification in Noisy Broadcast Audio," published in Signal and Image Processing 2002, another peer-reviewed journal. The paper shows experimental measurement of error rates of their set of audio identification features.

"Modulation-Scale Analysis for Content Identification," published in IEEE Transactions on Signal Processing 2004, a peer-reviewed journal. This paper shows experimental measurements of error rates and compares behavior of their feature set to other popular feature sets in audio fingerprint-based identification.