

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

FREE STREAM MEDIA CORP. D/B/A SAMBA TV,
Petitioner

v.

GRACENOTE, INC.,
Patent Owner

Inter Partes Review Case No. IPR2020-00216
U.S. Patent No. 9,066,114

DECLARATION OF DR. PIERRE MOULIN

TABLE OF CONTENTS

I. INTRODUCTION	3
A. BACKGROUND AND QUALIFICATIONS	5
II. LEGAL FRAMEWORK	9
A. OBVIOUSNESS.....	9
B. CLAIM CONSTRUCTION.....	16
III. OPINIONS	17
A. LEVEL OF A PERSON HAVING ORDINARY SKILL IN THE ART.....	17
B. SUMMARY OF THE CHALLENGED PATENTS.....	19
C. SUMMARY OF THE PRIOR ART REFERENCES RELIED ON IN THE PETITION.....	22
D. MURPHY DOES NOT TEACH “A REFERENCE TIME POINT IDENTIFIED BY THE REFERENCE FINGERPRINT,” AS REQUIRED BY CLAIMS 1, 8, AND 10 OF THE ’114 PATENT	28
E. MURPHY DOES NOT TEACH “REFERENCE FINGERPRINTS PREVIOUSLY DERIVED” FROM THE “CONTENT BEING PLAYED BACK,” AS REQUIRED BY CLAIMS 1, 8, AND 10 OF THE ’114 PATENT	29
F. MURPHY DOES NOT TEACH “AN ACTION TO BE TRIGGERED DURING PLAYBACK OF THE SEGMENT,” AS REQUIRED BY CLAIMS 11, 20, AND 21 OF THE ’114 PATENT	30
G. MURPHY DOES NOT TEACH “ACCESSING SEGMENTS OF CONTENT” DURING FINGERPRINT GENERATION, AS REQUIRED BY CLAIMS 11, 20, AND 21 OF THE ’114 PATENT	31
H. THE CLAIMED INVENTION OF THE CHALLENGED PATENTS DIFFERENTIATES OVER MURPHY IN VIEW OF BRUNK.....	33
I. MURPHY DOES NOT TEACH “DERIVING A SECOND FINGERPRINT BASED ON A SEGMENT OF THE MULTIMEDIA SIGNAL THAT CORRESPONDS TO THE SECOND TRIGGER TIME POINT” AS REQUIRED BY CLAIMS 1C, 8C, AND 15C OF THE ’962 PATENT	35
J. MURPHY DOES NOT TEACH “RECEIVING A SECOND TRIGGER TIME POINT,” AS REQUIRED BY CLAIMS 1, 8, AND 15 OF THE ’962 PATENT	38
K. MURPHY DOES NOT TEACH “RECEIVING ... A SECOND ACTION,” AS REQUIRED BY CLAIMS 1, 8, AND 15 OF THE ’962 PATENT	40
L. MURPHY DOES NOT TEACH “THE ACTION BEING ASSOCIATED WITH A TIME POINT INDICATING WHEN, IN THE MULTIMEDIA CONTENT, THE ACTION IS TO BE PERFORMED” AS REQUIRED BY CLAIMS 1, 11, 19, AND 24 OF THE ’831 PATENT..	41
M. MURPHY DOES NOT TEACH “A PLURALITY OF REFERENCE FINGERPRINTS HAVING BEEN PREVIOUSLY DERIVED” FROM THE MULTIMEDIA CONTENT BEING PLAYED BACK, AS REQUIRED BY CLAIM 1 OF THE ’831 PATENT	44

I, Pierre Moulin, hereby declare the following:

I. INTRODUCTION

1. My name is Dr. Pierre Moulin, D. Sc., and I am a Professor of Electrical and Computer Engineering (“ECE”) and Statistics at the University of Illinois. I have been retained by counsel for Gracenote, Inc. (“Patent Owner”) as a technical expert in the above-captioned case, as well as two related cases. Specifically, I have been asked to respond to certain issues raised by Petitioner, Free Stream Media Corp. d/b/a Samba TV (“Petitioner”) in IPR petition No. IPR2020-00216 regarding U.S. Patent No. 9,066,114 (“the ’114 Patent”), IPR petition No. IPR2020-00217 regarding U.S. Patent No. 9,407,962 (“the ’962 Patent”), and IPR petition No. IPR2020-00218 regarding U.S. Patent No. 9,479,831 (“the ’831 Patent”) (the ’114 Patent, the ’962 Patent, and the ’831 Patent are collectively referred to as “the Challenged Patents”). For the sake of convenience, I have been asked to submit a single declaration with regard to opinions I have reached for each of these three respective IPR petitions, and I understand this same declaration is being submitted separately for each of these three matters.

2. I am being compensated in this matter, and no part of my compensation depends on my opinions or the outcome of these proceedings. I note

that I have no financial interest in Petitioner or Patent Owner, and I have no other interest in the outcome of these matters.

3. In reaching my opinions in this matter, I have reviewed the following materials:

- Petition for *Inter Partes* Review of U.S. Patent No. 9,066,114 (“’114 Petition”);
- Patent Owner Preliminary Response to Petition for *Inter Partes* Review of U.S. Patent No. 9,066,114;
- Petition for *Inter Partes* Review of U.S. Patent No. 9,407,962 (“’962 Petition”);
- Patent Owner Preliminary Response to Petition for *Inter Partes* Review of U.S. Patent No. 9,407,962;
- Petition for *Inter Partes* Review of U.S. Patent No. 9,479,831 (“’831 Petition”);
- Patent Owner Preliminary Response to Petition for *Inter Partes* Review of U.S. Patent No. 9,479,831;
- Exhibit 1001 – U.S. Patent No. 9,066,114 (“the ‘114 Patent”);
- Exhibit 1002 – U.S. Patent No. 9,407,962 (“the ’962 Patent”);
- Exhibit 1003 – U.S. Patent No. 9,479,831 (“the ’831 Patent”);
- Exhibit 1004 - Select portions of prosecution history of the ’114 Patent
- Exhibit 1005 - Select portions of prosecution history of the ’962 Patent
- Exhibit 1006 - Select portions of prosecution history of the ’831 Patent
- Exhibit 1007 – Expert Declaration of Dr. Ahmed H. Tewfik (“’114 Tewfik Decl.”);
- Exhibit 1008 – Expert Declaration of Dr. Ahmed H. Tewfik (“’962 Tewfik Decl.”);
- Exhibit 1009 – Expert Declaration of Dr. Ahmed H. Tewfik (“’831 Tewfik Decl.”);

- Exhibit 1011 – U.K. Published Patent Application GB 2,375,907A to Murphy (“Murphy”);
- Exhibit 1012 – U.S. Patent Application Publication 2002/0126872 to Brunk, et al. (“Brunk”);
- Exhibit 1013 - W. Kate, et al. “trigg&link: A New Dimension in Television Program Making,” in Lecture Notes in Computer Science 1197, vol. 1242, Multimedia Applications, Services and Techniques-ECMAST '97, (1997) 51-65 (“Kate”);
- Exhibit 2004 – U.S. Patent No. 7,185,201 to Rhoads, et al. (“Rhoads”)
- Exhibit 2005 – Prosecution History for U.S. Patent No. 8,020,000
- Exhibit 2006 – U.S. Patent No. 8,020,000 Patent to Oostveen, et al.
- Exhibit 2007 – U.S. Patent Application Publication No. 2003/0021441 to Levy, et al. (“Levy”)
- Exhibit 2008 – File History for U.S. Patent Application No. 16/264,134 (“’134 Application File History” or “’053 Patent File History”)
- Exhibit 2009 – Prosecution History for U.S. Patent No. 9,066,114

A. Background and Qualifications

4. My qualifications to testify about the Challenged Patents, the prior art at issue, and the relevant technology are summarized below, and additional detail regarding my background, experience, and qualifications can be found in my Curriculum Vitae, which has been attached hereto as Appendix A.

5. I am currently a Professor of Electrical and Computer Engineering (“ECE”) and Statistics at the University of Illinois at Urbana-Champaign (UIUC).

6. I received an Ingénieur Civil Electricien from Faculté Polytechnique de Mons, Belgium in 1984, a Masters in Electrical Engineering from the

Washington University, St. Louis in 1986, and a PhD in Electrical Engineering from Washington University, St. Louis in 1990.

7. After working as a Research Scientist for Bell Communications Research in Morristown, New Jersey, I joined the University of Illinois as Assistant Professor (1996) and later became Associate Professor (1999) and Professor (2003) in the Department of Electrical and Computer Engineering, Research Professor in the Coordinated Science Laboratory, faculty member in the Beckman Institute's Image Formation and Processing Group, and affiliate professor in the Department of Statistics. I am also a member of the Information Trust Institute and the founding director of the Center for Information Forensics, a multidisciplinary research center involving twenty colleagues. I am also a member of the Advanced Digital Sciences Center, a multidisciplinary research center of UIUC in Singapore. My fields of professional interest and expertise are information theory, image and video processing, statistical signal processing and modeling, decision theory, information hiding and authentication, and the application of multiresolution signal analysis, optimization theory, and fast algorithms to these areas.

8. I have served as Associate Editor for the IEEE Transactions on Information Theory from 1996 till 1998, for the IEEE Transactions on Image Processing from 1999 till 2002, and then as Area Editor from 2002 till 2006. In

1999, I was co-chair of the IEEE Information Theory workshop on Detection, Estimation and Classification. I was a Guest Editor of the IEEE Transactions on Information Theory 2000 special issue on Information-Theoretic Imaging; Guest Editor of the IEEE Transactions on Signal Processing's 2003 special issue on Data Hiding; and member of the IEEE Image and Multidimensional Signal Processing (IMDSP) Society Technical Committee (1998-2003), the Information Forensics and Security (IFS) Technical Committee (2007-2010) and the Image, Video, and Multidimensional Signal Processing (IVMSP) Technical Committee (2010-present); and the Board of Governors of the IEEE Signal Processing Society (2005-2007). I was co-founder and Editor-in-Chief of the IEEE Transactions on Information Forensics and Security. I am a Fellow of IEEE (2003) and a recipient of a 1997 Career award from the National Science Foundation and recipient of the IEEE Signal Processing Society 1997 Best Paper award in the IMDSP area. I am also co-author of a paper that received the IEEE Signal Processing Society 2002 Young Author Best Paper award in the IMDSP area. I was selected as Beckman Associate of UIUC's Center for Advanced Study (2003) and Sony Faculty Scholar (2005-2007). I have also been plenary speaker for several international conferences and workshops and was a Distinguished Lecturer of the IEEE Signal Processing Society for 2012-2013.

9. My research and expertise have encompassed many areas relevant to the technology at issue in these proceedings, including signal, image and video processing, and automated content recognition, among others. Broadly, my research and expertise have encompassed the areas of signal processing, modeling, compression, imaging, information theory, statistics, stochastic processes, wavelets and multiresolution analysis, estimation, detection, pattern recognition, learning, visual analytics, content identification, watermarking, hashing, fingerprinting, steganography, information security, and optimization. My students and I research the theory and application of statistics and information theory to a variety of signal, image, and video processing problems. My work has included a number of primary research thrusts. First, it has included work in the area of content identification and robust hashing. We have worked to develop a fundamental theory of content identification from first principles, and to develop hashing codes that have outperformed the state of the art. Second, it has included information forensics and security, and, specifically, developing an information theory of watermarking and data hiding, in collaboration with others. This led to fruitful collaborations in which we constructed codes that can approach fundamental bounds and can resist arbitrary noise and a broad class of desynchronization attacks. We also analyzed security flaws for watermarking systems, and methods that can remedy those flaws. A third research thrust of mine, conducted primarily

at ADSC in Singapore, has included multimodal analytics, using techniques from statistical learning and computer vision. Our team won the 2012 HARL-ICPR competition on human activity recognition. A fourth and more recent research thrust has included learning codes for joint image/video compression and classification.

10. I have authored numerous publications in my areas research, including those listed in my Curriculum Vitae. Examples of courses I have taught include Introduction to Image and Video Processing (ECE418), Wavelets (ECE544), Statistical Imaging and Video Processing (ECE544), Statistical Learning (ECE544), Computational Inference and Learning (ECE566), and Making Sense of Big Data (ECE398BD/ECE314).

II. LEGAL FRAMEWORK

11. I am a technical expert and do not offer any legal opinions. However, counsel has informed me regarding certain legal principles regarding patentability and related matters under United States patent law, which I have applied in performing my analysis and arriving at my technical opinions in these matters.

A. Obviousness

12. I have been informed that a person cannot obtain a patent on an invention if the differences between the invention and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was

made to a person having ordinary skill in the art (“POSITA”). I have been informed that a conclusion of obviousness may be founded upon more than a single item of prior art. I have been further informed that obviousness is determined by evaluating the following factors: (1) the scope and content of the prior art, (2) the differences between the prior art and the claim at issue, (3) the level of ordinary skill in the pertinent art, and (4) secondary considerations of non-obviousness.

13. I have also been informed that it is improper for the obviousness inquiry to be done in hindsight. Instead, the obviousness inquiry should be done through the eyes of a POSITA at the time of the alleged invention. Specifically, I have been informed that it is improper, in an obviousness analysis, to apply hindsight by using the challenged patent’s specification or claims as a roadmap to assemble a collection of references or modifications to references. I am informed that the proper question is what would have been obvious to a person of ordinary skill in the art without the benefit of hindsight and specifically without access to the patent at issue.

14. In considering whether certain prior art renders a particular patent claim obvious, I have been informed that I should consider the scope and content of the prior art, including the fact that one of skill in the art would regularly look to the disclosures in patents, trade publications, journal articles, conference papers,

industry standards, product literature and documentation, texts describing competitive technologies, requests for comment published by standard setting organizations, and materials from industry conferences, as examples. I have been informed that for a prior art reference to be proper for use in an obviousness analysis, the reference must be “analogous art” to the claimed invention. I have been informed that a reference is analogous art to the claimed invention if: (1) the reference is from the same field of endeavor as the claimed invention (even if it addresses a different problem); or (2) the reference is reasonably pertinent to the problem faced by the inventor (even if it is not in the same field of endeavor as the claimed invention). In order for a reference to be "reasonably pertinent" to the problem, it must logically have commended itself to an inventor's attention in considering his problem. In determining whether a reference is reasonably pertinent, one should consider the problem faced by the inventor, as reflected either explicitly or implicitly, in the specification.

15. I have been informed that, in order to establish that a claimed invention was obvious based on a combination of prior art elements, a clear articulation of the reason(s) why a claimed invention would have been obvious must be provided. Specifically, I am informed that, under the U.S. Supreme Court’s decision in *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398 (2007), a combination of multiple items of prior art renders a patent claim obvious when

there was an apparent reason for one of ordinary skill in the art, at the time of the invention, to combine the prior art, which can include, but is not limited to, any of the following rationales: (A) combining prior art methods according to known methods to yield predictable results; (B) substituting one known element for another to obtain predictable results; (C) using a known technique to improve a similar device in the same way; (D) applying a known technique to a known device ready for improvement to yield predictable results; (E) trying a finite number of identified, predictable potential solutions, with a reasonable expectation of success; (F) identifying that known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art; or (G) identifying an explicit teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine the prior art references to arrive at the claimed invention. I am also informed that where there is a motivation to combine, claims may be rejected as prima facie obvious provided a POSITA would have had a reasonable expectation of success regarding the proposed combination.

16. I am informed that the existence of an explicit teaching, suggestion, or motivation to combine known elements of the prior art is a sufficient, but not a necessary, condition to a finding of obviousness. This so-called “teaching-

suggestion-motivation” test is not the exclusive test and is not to be applied rigidly in an obviousness analysis. In determining whether the subject matter of a patent claim is obvious, neither the particular motivation nor the avowed purpose of the patentee controls. Instead, the important consideration is the objective reach of the claim. In other words, if the claim extends to what is obvious, then the claim is invalid. I am further informed that the obviousness analysis often necessitates consideration of the interrelated teachings of multiple patents, the effects of demands known to the technological community or present in the marketplace, and the background knowledge possessed by a person having ordinary skill in the art. All of these issues may be considered to determine whether there was an apparent reason to combine the known elements in the fashion claimed by the patent.

17. I also am informed that in conducting an obviousness analysis, a precise teaching directed to the specific subject matter of the challenged claim need not be sought out because it is appropriate to take account of the inferences and creative steps that a POSITA would employ. The prior art considered can be directed to any need or problem known in the field of endeavor at the time of invention and can provide a reason for combining the elements of the prior art in the manner claimed. In other words, the prior art need not be directed towards solving the same specific problem as the problem addressed by the patent. Further, the individual prior art references themselves need not all be directed towards

solving the same problem. I am informed that, under the *KSR* obviousness standard, common sense is important and should be considered. Common sense teaches that familiar items may have obvious uses beyond their primary purposes.

18. I also am informed that the fact that a particular combination of prior art elements was “obvious to try” may indicate that the combination was obvious even if no one attempted the combination. If the combination was obvious to try (regardless of whether it was actually tried) or leads to anticipated success, then it is likely the result of ordinary skill and common sense rather than innovation. I am further informed that in many fields it may be that there is little discussion of obvious techniques or combinations, and it often may be the case that market demand, rather than scientific literature or knowledge, will drive the design of an invention. I am informed that an invention that is a combination of prior art must do more than yield predictable results to be non-obvious.

19. I am informed that for a patent claim to be obvious, the claim must be obvious to a POSITA at the time of the alleged invention. I am informed that the factors to consider in determining the level of ordinary skill in the art include (1) the educational level and experience of people working in the field at the time the invention was made, (2) the types of problems faced in the art and the solutions found to those problems, and (3) the sophistication of the technology in the field.

20. I am informed that it is improper to combine references where the references teach away from their combination. I am informed that a reference may be said to teach away when a POSITA, upon reading the reference, would be discouraged from following the path set out in the reference, or would be led in a direction divergent from the path that was taken by the patent applicant. In general, a reference will teach away if it suggests that the line of development flowing from the reference's disclosure is unlikely to be productive of the result sought by the patentee. I am informed that a reference teaches away, for example, if (1) the combination would produce a seemingly inoperative device, or (2) the references leave the impression that the product would not have the property sought by the patentee. I also am informed, however, that a reference does not teach away if it merely expresses a general preference for an alternative invention but does not criticize, discredit, or otherwise discourage investigation into the invention claimed.

21. I am informed that even if a prima facie case of obviousness is established, the final determination of obviousness must also consider “secondary considerations” if presented. In most instances, the patentee raises these secondary considerations of non-obviousness. In that context, the patentee argues an invention would not have been obvious in view of these considerations, which include: (a) commercial success of a product due to the merits of the claimed

invention; (b) a long-felt, but unsatisfied need for the invention; (c) failure of others to find the solution provided by the claimed invention; (d) deliberate copying of the invention by others; (e) unexpected results achieved by the invention; (f) praise of the invention by others skilled in the art; (g) lack of independent simultaneous invention within a comparatively short space of time; (h) teaching away from the invention in the prior art.

22. I am further informed that secondary-considerations evidence is only relevant if the offering party establishes a connection, or nexus, between the evidence and the claimed invention. The nexus cannot be based on prior art features. The establishment of a nexus is a question of fact.

B. Claim Construction

23. I have been informed by counsel that the first step in an unpatentability analysis involves construing the claims, as necessary, to determine their scope. Second, the construed claim language is then compared to the disclosures of the prior art. I am informed that claims are generally given their ordinary and custom meaning as understood by one of ordinary skill in the art at the time of the invention, in light of the patent specification.

24. For purposes of this declaration, I was not asked to offer any opinions on the issue of claim construction, and I was asked to simply assume that the express claim constructions offered in the Petitions were correct and to

apply those constructions in my analysis. *See* '114 Petition, 11-15; '962 Petition, 11-12; '831 Petition, 16-20. For all other claim terms not expressly construed by Petitioner or Patent Owner, I was instructed to apply the meanings of the claim terms of the challenged patent that were generally consistent with the terms' ordinary and customary meaning, as a person of ordinary skill in the art would have understood them at the time of the invention after reading the challenged patent (which I understand is the so-called *Phillips* claim construction standard). I have been instructed to assume for purposes of this proceeding that the time of each of the claimed inventions was July 11, 2003.

III. OPINIONS

A. Level of a Person Having Ordinary Skill in the Art

25. Unless expressly stated otherwise, my opinions provided in this Declaration are made based on the assumed date of invention aligning with the priority date of July 11, 2003 of the Challenged Patents, which counsel has instructed me to assume is the correct date. To the extent that any verb tense used in this Declaration or any deposition or testimony provided in this matter is in the present tense, e.g., “would understand” or “would be obvious,” such verb tense should be understood to be my opinion as of the Challenged Patents' priority date (again, unless expressly stated otherwise). I may, on occasion, merely use the present verb tense for ease of reading.

26. I understand that Petitioner's expert, Ahmed H. Tewfik, Sc.D., has proposed the following definition for the level of a person having ordinary skill in the art of the Challenged Patents as of July 11, 2003: "a bachelor's degree in computer science, electrical engineering or a similar field, and two years of work experience in multimedia signal processing, including watermarking, fingerprinting and their applications. A greater degree of education or work experience could substitute for a lesser degree of the other." See *'114 Tewfik Declaration*, ¶54, *'962 Tewfik Declaration*, ¶49, *'831 Tewfik Declaration*, ¶54. I was instructed by counsel to assume this definition was correct and to apply this definition for purposes of my analysis in this Declaration.

27. I met at least these minimum qualifications to be a person having ordinary skill in the art at least as of July 11, 2003, as set forth above in my description of my background and qualifications. Further, although my qualifications may exceed those of the hypothetical person having ordinary skill in the art defined above, my analysis and opinions regarding the Challenged Patents have been rendered from the perspective of a person having ordinary skill in the art at the time of the invention.

B. Summary of the Challenged Patents

1. Problem in the Prior Art to be Solved

28. As the background of the invention of the challenged patents discusses, various prior art approaches were known for triggering actions at detected time points in media content to enhance the content by “connecting” a piece of audio or video content to external actions, such as displaying additional information about the content at those points. ’114 Patent (Ex. 1001), 1:22-62.¹ For example, in the “trigg&link” system (i.e., trigger and link), television content was enhanced by having an icon appear on the screen when a particular actor appeared, and selection of the icon would cause further information about the actor to be displayed. *Id.* at 1:36-51. But, as the Challenged Patents describe, these time instants were required “to have been indicated in the video stream” itself by the program provider. *Id.* at 1:48-52; *see also* Ex. 1013 (reference describing that, in the trigg&link system, “[t]he program provider inserts triggers in the broadcast stream”). As another example, the Challenged Patents describe “Local insertion,”

¹ For ease of reference, all of my citations to the specification of the Challenged Patents are to the ’114 Patent, which shares a common specification with the ’831 Patent and the ’962 Patent. Thus, the citations to the specification of the ’114 Patent are intended to also apply to the ’831 Patent and the ’962 Patent, as well.

in which a national program producer would allow specific parts of a national broadcast to be replaced by a regional program, such as allowing local advertisements or local weather forecasts to replace national ones. Ex. 1001, 1:53-58. But this approach similarly required the national program producer to “indicate which segments are suitable for such local insertion” so that such “indicated segments” could then be replaced with local content by a local redistributor. *Id.* at 1:58-62.

29. As the Challenged Patents explain, in both of these prior art examples, it was necessary for the program maker or broadcaster to have marked or indicated specific time instants so as to allow actions to be triggered therefrom. *Id.* at 1:63-67. The Challenged Patents explain that this was a problem because “this requires the broadcaster’s cooperation to insert these triggers thereby making an enhancement service provider dependent on the broadcaster.” *Id.* at 2:1-4. The Challenged Patents describe that there were various known ways for such markers or triggers to be inserted into the content stream to then allow actions to be triggered therefrom, but the patent again emphasizes that these known ways “need the cooperation of all actors in the broadcast chain.” *Id.* at 2:14-17. Another known way described in the Challenged Patents to embed triggers in content included insertion of watermarks, which the patent describes as suffering from the

additional problem of the content having to be modified to insert the watermarks.
Id. at 2:30-34.

2. *Solution of the Patented Invention*

30. “[T]he invention” of the Challenged Patents, however, “solves [these] problems,” by “**relating**” the desired trigger actions to the multimedia content by an enhancement service provider first selecting desired trigger time points within the multimedia content and then fingerprinting the multimedia content at those time points, instead of having to rely on program makers or broadcasters to have previously marked, indicated, or inserted trigger time points into the content that an enhancement service provider could then use to trigger actions. *Id.* at 2:36-58 (emphasis added). A trigger time point is selected in the multimedia signal where a trigger action is desired and then the signal is fingerprinted at that selected time point so as to create a way to trigger an action at that time point. *Id.* Specifically, “for each given trigger time point,” the system “deriv[es] a fingerprint on the basis of a segment of the multimedia signal, where the segment of the multimedia signal is **unambiguously** related with the given trigger time point.” *Id.* By selecting a trigger time point and then deriving the fingerprint based on that selected time point in the multimedia signal, the invention avoids the need to rely on broadcasters, program makers, or other actors in the

broadcast chain to have previously marked, indicated, or inserted specified time points to serve as triggers.

C. Summary of the Prior Art References Relied on in the Petition

1. *Murphy*

a. **Murphy relies on program provider or broadcaster cooperation**

31. Murphy describes an “automatic recognition system” for use in either i) a transmitter for recognizing a program being broadcast to “generate useful ancillary data therefrom” (for example, to automatically add associated data to a signal being broadcast) or ii) in a receiver to automatically perform a function such as recording a program. *Murphy*, Abstract, 9:26-10:31, 13:8-25. Murphy requires the program maker or broadcaster to include distinctive and easily recognizable audible jingles (which Murphy calls “audio cues”) placed at the beginning or end of particular programs. *Murphy*, 3:10-20. These audio cues are placed in the multimedia signal by either the program maker or broadcaster. *Id.* (“It is standard practice for programme makers and broadcasters to place an audio cue such as a piece of music or jingle immediately before a particular programme....”).

32. Murphy’s automatic recognition system includes a “programme cue identifier 10” (depicted in Figure 1 of Murphy), which includes a database 12 that stores information on the audio cues the system is required to recognize. *Murphy*, 6:23-32. A “signature” of the audio cue is stored, rather than storing the raw audio

cue; and Murphy describes that the recognition system itself can generate the signature of the audio cue from the raw audio cue or the signature may be generated by an external system and then downloaded into the database of the recognition system. *Id.* at 6:32-7:5. The signature of the audio cue may be derived from a standalone version of the audio cue (which Murphy describes as a “clean, noise-free version of the audio cue”) or derived from a previously broadcast version of the audio cue (which Murphy describes as “an audio cue recorded from a previous broadcast signal”) (*id.* at 7:5-7), which demonstrates to a POSITA that the audio cue signature stored in the database is not generated from the same content being broadcast in the audio broadcast signal that the automatic recognition system analyzes when attempting to detect the presence of audio cues in the broadcast signal.

33. Murphy then describes that a live broadcast radio audio signal may then be continuously analyzed by comparing it to the stored audio cues corresponding to given programs so as to detect the presence of audio cues in the broadcast audio signal. *Id.* at 10:1-12. In a transmitter embodiment (depicted in Figure 3 of Murphy), “[w]hen a match it detected,” the system generates appropriate event flags and inserts a data signal containing information about the program into the audio signal. *Id.* at 10:13-31. For example, if the audio cue corresponds to the audio cue that is played at the start of a traffic announcement,

then the system sets a “traffic announcement flag” and inserts a “Traffic Programme Identification Code” into a data signal that is added to the audio signal for the program. *Id.* In a receiver embodiment (depicted in Fig. 4), a receiver 80 adapted to include the recognition system receives a broadcast audio signal that is analyzed to check for a match with any of the stored audio cues. *Id.* at 13:8-21. “When a match with a particular audio cue is detected,” an event flag is set, which may be a program start flag that causes the recording system of the receive to start recording the broadcast audio signal. *Id.* at 13:21-26.

34. Murphy relies on *precisely the same type of problematic prior art approach the Challenged Patents distinguished from and overcame with their invention*: trigger time points indicated or inserted by the program producer or broadcaster, which I discussed above. ‘114 Patent, 1:63-2:4 (describing the prior art “trigg&link” and “local insertion” approaches where it was “necessary to mark or associate specific time instants in the video stream at which additional information should be available” and describing mechanisms that “require[d] the broadcaster’s cooperation to insert these triggers thereby making an enhancement service provider dependent on the broadcaster.”). Murphy relies on the “standard practice of programme makers and broadcasters **to place an audio cue** such as a piece of music or jingle immediately before a particular program,” and Murphy relies exclusively on the detection of these previously inserted “audio cues” as

triggers to cause actions, such as causing a recorder to start recording. *Murphy*, 3:10-27, 5:9-11 (“The audio recognition system is designed to detect a match between an input audio signal and one or more known audio cues stored for recognition purposes.”), 13:21-26. “Audio cues” are precisely that—“cues” that serve to indicate or mark points within the content at which actions can be triggered, and *Murphy* is entirely reliant on these “cues” to have been included in the multimedia signal, either by the program maker or broadcaster, which necessarily requires their cooperation. This is precisely the type of prior art problem that the invention of the Challenged Patents overcomes. In *Murphy*, the triggers have been placed in the multimedia signal by the program maker or broadcaster, which necessarily requires their cooperation to both use those audio cues and to have included them at particular points within the program.

35. I note that *Murphy* states its system “does not require the broadcaster to adapt the signal to include additional data which can be detected and used to control the recording system 82 nor does it impose any particular information format onto broadcasters.” *Id.* at 14:20-25. This statement merely refers to the fact that *Murphy* does not rely on program makers or broadcasters to insert inaudible codes such as watermarks at time points in the signal, which is an approach *Murphy* describes as problematic (*Murphy*, 2:30-3:8). I note in passing that *Murphy*’s broadcaster-inserted audio cues could be thought of as audible

watermarks. The Challenged Patents similarly seek to avoid the need for insertion of watermarks (*'114 Patent*, 2:30-35), but that is only one form of potential broadcaster cooperation that the patents avoid. *'114 Patent*, 2:34-35 (noting that a disadvantage of watermarking is that it necessarily changes the video/audio), 2:43-45 (noting that it is a “*further* object” of the invention “to enable trigger actions without modifying the multimedia signal”). Murphy’s mere non-reliance on insertion of traditional inaudible watermarks, however, does not mean that Murphy removes reliance on cooperation from program makers and broadcasters to have placed the audible audio cues in the multimedia content. Murphy absolutely still requires the cooperation of the program maker or broadcaster to utilize the audio cues and also to have placed them at points within the audio broadcast, which entirely prevents an enhancement service provider from fingerprinting any time point of their choice in a stream and relating or associating a desired action to that time point—which, as discussed was the entire purpose and focus of the claimed invention in overcoming the prior art reliance on program provider and broadcaster cooperation to have already marked, indicated, or inserted time points within a multimedia signal from which actions can be triggered.

2. *Brunk*

36. *Brunk* generally relates to deriving and utilizing content signatures for identifying items of content. *Brunk*, Abstract, [0003], [0006]. Specifically, *Brunk*

describes deriving a signature from a content item, and the signature “may include a number derived from a signal (e.g., a content item) that serves as a statistically unique identifier of that signal,” which, in Brunk’s terms, “means that there is a high probability that the signature was derived from the digital signal in question.” *Brunk*, [0021]. Brunk provides a “hash” as one possible signature algorithm that can be used, and the hash is applied to a selected portion of a signal, such as the first 10 seconds, a selected video frame, etc. *Id.* Importantly, Brunk is not concerned with timing of actions or accurately identifying time points within a multimedia signal for purposes of timing actions, as Petitioner’s declarant, Dr. Tewfik, acknowledges. *See ‘114 Tewfik Decl.*, ¶ 144 (“Brunk does not contain an express disclosure of using cue signals to allow for detection of fingerprints to trigger actions at specific times based on the signatures. In other words, Brunk does not expressly disclose using fingerprints to identify a trigger time point of a sequence of trigger time points in the content for triggering an action.”) (citing *Brunk*, [0066]-[0068]). Instead, Brunk is directed to merely identifying items of content. As one example, Brunk would be applicable to, for example, identifying the name of a song by generating a content signature from a selected portion of the song.

D. Murphy does not teach “a reference time point identified by the reference fingerprint,” as required by claims 1, 8, and 10 of the ’114 Patent

37. I note that Claims 1, 8 and 10 recite “a reference time point identified by the reference fingerprint.” *See ’114 Patent*, Claims 1, 8, and 10. It is my opinion that Murphy does not teach this limitation. In particular, Murphy teaches that “[w]hen a match with a particular audio cue is detected an event flag is set [...] which, when detected by the recording system 82, causes the recording system 82 to start recording the signal.” *Murphy* 13:21-25. Murphy also teaches that an event flag to stop recording is generated “upon repeated occurrence of the audio cue.” I would understand this to indicate that actions are taken at *different* times for the *same* audio cue signature. I would further understand Murphy’s system to operate such that a (context-dependent) action is taken *every time* the audio cue signature is matched. At least because signatures are matched, causing actions to be taken, at different times in the same content (here, the same episode of the program), the signature cannot identify a reference time point because the time point in the content is ambiguous. Murphy, therefore, does not teach “a reference time point identified by the reference fingerprint,” as recited in claims 1, 8, and 10.

E. Murphy does not teach “reference fingerprints previously derived” from the “content being played back,” as required by claims 1, 8, and 10 of the ’114 Patent

38. I note that Claims 1, 8 and 10 recite the steps of “determin[e/ing] a plurality of trigger fingerprints *from content being played back*” and “access[ing] a database that includes a plurality of reference fingerprints *previously derived from the content.*” See ’114 Patent, Claims 1, 8, and 10. These limitations, therefore, recite that the reference fingerprints must have been previously derived from the same “content” that is subsequently “played back” used to determine trigger fingerprints when “played back.” Murphy, however, describes audio cues as being derived from “a clean, noise-free version of the signal” or “versions previously broadcast.” *Murphy*, 7:5-7. It is my opinion that, in neither of these scenarios, is the signature derived from the same content that is being played back, whether the audio cue signature is derived by the broadcaster for distribution or by the end user directly. *Murphy*, 14:6-15. In the “clean, noise-free” scenario, the audio cue signature is derived from a version of the audio cue that is not part of *any* content, much less the broadcast audio signal being transmitted received, played and/or recorded. In the “previously broadcast” scenario, the audio cue signature is derived from (at best) *different* content than the content in which the audio cue is recognized (i.e., the broadcast audio signal being transmitted received, played and/or recorded). I note that Murphy is concerned with live broadcasts and

includes no disclosure of repeatedly airing the same content. *See, e.g., Murphy*, 2:11-13 (rejecting a solution that requires newsreaders to manually take action to add program signal related information during broadcasts); 12: 2-5 (live stories and sports events causing schedule disruption).

F. Murphy Does Not Teach “an action to be triggered during playback of the segment,” as required by claims 11, 20, and 21 of the ’114 Patent

39. I understand that Claims 11, 20 and, 21 require “access[ing]” “a segment” “being associated with a corresponding action to be triggered *during playback of the segment*,” “generat[e/ing] a reference fingerprint of *the segment*,” “assign[ing] the reference fingerprint” “to the corresponding action that is to be triggered *during playback of the segment*,” and “causing the playback device to trigger the corresponding action *during playback of the segment*.” *See ’114 Patent*, Claims 11, 20, 21. It is my opinion that Murphy does not teach or suggest this limitation, at least because Murphy is limited to taking actions *after* an audio cue is recognized (i.e., after playing the segment containing the audio cue). I understand that Petitioner’s expert, Dr. Tewfik, contends that it would be obvious to modify Murphy’s system to incorporate a time-delay unit and configure the length of the time delay to “slightly less than the length of the segment” in order to render this claim limitation obvious. *Petition* 38; *Tewfik Dec.* ¶ 218. I disagree. The purpose of Murphy’s time-delay unit is “to avoid the possibility of losing the initial part of the

programme.” *Murphy*, 12:23-24. Modifying the time-delay unit to use a delay “slightly less than the length of the segment,” as Dr. Tewfik proposes, would *deliberately* lose the initial part of the program. For example, if the audio cue corresponded to the jingle for a particular program, then the proposed modification would involve deliberately configuring the system to record *only a portion* of the jingle. I would understand that this is less desirable than recording the entire jingle (as disclosed in *Murphy*) and less desirable than recording none of the jingle (as would occur if recording started directly after the audio cue). I would further understand *Murphy*’s stated goal “to avoid the possibility of losing the initial part of the programme” to *teach away* from the modification proposed by Dr. Tewfik, as the modification would deliberately “los[e] the initial part of the programme” by virtue of setting the time delay to start recording during the playback of the jingle.

G. Murphy does not teach “accessing segments of content” during fingerprint generation, as required by claims 11, 20, and 21 of the ’114 Patent

40. I understand that Claim 11 recites steps of “accessing segments of content, a segment among the segments being associated with a corresponding action to be triggered during playback of the segment,” “generating a reference fingerprint of the segment associated with the corresponding action,” and providing an identifier “causing the playback device to trigger the corresponding action during the playback of the segment.” *’114 Patent*, Claim 11. I further

understand that claims 20 and 21 require substantially similar limitations. '114 Patent, Claims 20, 21. It is my opinion that Murphy does not teach this limitation, at least because Murphy accesses the audio cue in isolation in order to generate the signature, rather than deriving it from the content. I understand that the Petition identifies Murphy's sampling windows (used when *detecting* fingerprints rather than generating them) with the claimed "segments" of content which the claims require be accessed during fingerprint generation. This identification does not stand up to scrutiny. First, as previously stated, the content is divided into windows when detecting fingerprints rather than generating them. *Murphy*, 9:5-7 ("The covariance matrix [i.e., the signature] of the audio cue is determined in a similar manner except that windowing of the audio cue is not required."). These windows will generally not align with the start of the audio cue during fingerprint detection, because Murphy teaches starting a sampling window only once a second (two-second windows with a one-second overlap; see *Murphy*, 7:30-33; 8:7-8). I would understand that these overlaps are necessary because it is impossible to know, *a priori*, where in the broadcast audio stream the audio cue will start. See also *Murphy*, 8:4-5. Murphy goes on to clarify that a window is begun only once per second due to limited processing power, stating that "moving the window by a single sample each time results in very high processing overheads and slow processing time." *Murphy*, 8:5-7. Because of this, and because the full sample rate

in Murphy is 16kHz, windows will generally not align with the start of the audio cue.

41. Next, I understand that Dr. Tewfik opines that “Murphy teaches that the same process of *accessing the content segment by segment occurs when fingerprints are generated* for storage in database 12.” *Tewfik Dec.*, ¶ 215. This is incorrect. In fact, Murphy specifically states that “windowing the audio cue is not required” when generating the covariance matrix (signature) of the audio cue. *Murphy*, 9:5-7. I would understand this to mean that the content is not accessed “segment by segment” as Dr. Tewfik contends. Furthermore, accessing the content segment by segment to generate a fingerprint would be inconsistent with the other teachings of Murphy as well, which indicate that the audio cue length is the same as a single window length (*Murphy*, 7:30-33) and a “clean, noise-free” version of the audio cue can be used for generating the audio cue signature (*Murphy*, 7:5-7). Together, I would understand these teachings to indicate that the sample of the audio cue used to generate the fingerprint is only the length as a single window. Because Dr. Tewfik conflates “windows” and “segments,” it would not make sense to access a single segment “segment-by-segment” as he contends.

H. The Claimed Invention of the Challenged Patents Differentiates Over Murphy in View of Brunk

42. It is further my opinion that the invention of the Challenged Patents differentiates over the proposed combination of references, and that Murphy in

view of Brunk is unsuited to solving the problem described by the Challenged Patents. For example, the Challenged Patents describe how the system enables “frame accurate” localization within multimedia content. *’114 Patent*, 5:2-4. I would understand this to mean that the system can trigger an action to within 1/24-1/60 of a second based on conventional video frame rates. By contrast, I would understand that Murphy’s system (1) cannot locate an arbitrary point within an audio signal (at least because Murphy is limited to locating “jingles” or other distinctive audio cues), and (2) may be delayed for as much as three seconds from the beginning of the audio cue before an action can be triggered (because Murphy must recognize an entire, two-second audio cue, and the recognition window may not start until up to one second into the audio cue). *Murphy*, 7:30-33; 8:7-8. Murphy’s system is thus between 72 and 180 times less accurate at triggering an action than the invention of the Challenged Patents.

43. Similarly, Murphy’s reliance on static, repeated jingles means that there are certain programs and scenarios where Murphy will not work for even Murphy’s intended purpose. For example, any program without a distinctive jingle cannot be recorded by Murphy’s receiver or used to automatically generate PTY information in the transmitter embodiment. If a program includes a “cold open” (i.e., content that preceded the theme song), that cold open will not be recorded by Murphy’s system. Finally, if the audio cue appears at any point in the program

other than the beginning and end, these additional audio cues will disrupt operation of Murphy's system. For example, I note that the Petitions give the example of 60 Minutes' "ticking clock" audio cue, which "appears at the beginning and upon return from commercial break." *Petition*, 18. I note that, per Murphy's description of how Murphy's system operates, this would cause recording to begin at the beginning of the program *and recording to stop upon return from commercial break*. See *Murphy*, 11:2-3 (end flag generated upon repeated occurrence of audio cue); 14:3 (end flag stops recording). As such, Petitioner's example illustrates how the design of Murphy's system is limited by its reliance on broadcaster-inserted audio cues.

I. **Murphy does not teach "deriving a second fingerprint based on a segment of the multimedia signal that corresponds to the second trigger time point" as required by Claims 1c, 8c, and 15c of the '962 Patent**

44. It is my opinion that Murphy does not teach deriving multiple fingerprints from the same multimedia signal. As a result, Murphy fails to teach the claim limitation "deriving a second fingerprint based on a segment of the multimedia signal that corresponds to the second trigger time point," where the multimedia signal is the same signal from which the first fingerprint was derived. Rather, Murphy's system describes building up the template database based on *individual audio cues*. *Murphy*, 14:6-15. For example, Murphy discloses

downloading, from the broadcaster, “the signature of the audio cue for the particular program.” *Id.* I would understand these signatures to be derived based on the isolated, “clean, noise-free version of the audio cue” (*Murphy*, 7:5-6) rather than an audio cue that is added to any particular episode of the program. Next, *Murphy* discloses that “the audio cue itself could be downloaded” for local signature generation. *Murphy*, 14:10-11. I would understand this downloaded audio cue to likewise be the “clean, noise-free version” so as to be able to derive the best quality signature. Finally, *Murphy* discloses that the user could “input to the recognition system a sample of the audio cue recorded off-air.” I would understand this to indicate that the user provides to the recognition system *only the sample of the audio cue*; in other words, only enough of the signal to derive a single fingerprint. This is my understanding at least because *Murphy* does not disclose any techniques for automatically extracting an audio cue from a larger sample in order to derive its signature. If *Murphy*’s system *could* locate a jingle in a longer audio sample in order to derive its signature, there would be no need for a signature because the same technique could locate the jingle in a live broadcast. Rather, *Murphy*’s system must be provided with the isolated audio cue in order to derive its signature for future recognition.

45. I understand that Dr Tewfik has opined that “the processor 14 of recognition system 10 generates signatures (deriving fingerprints) of the signal

provided by the user.” *Tewfik, Dec.* ¶ 155, citing *Murphy*, Figure 2. To the extent that Dr. Tewfik is opining that the system generates plural signatures or derives plural fingerprints from the same signal provided by the user, I disagree. In particular, Murphy specifically distinguishes the processes of generating signatures for *audio cues* from the process of generating signatures for a broadcast audio signal, saying “[t]he covariance matrix of the audio cue is determined in a similar manner except that windowing of the audio cue is not required.” I would understand this to be the case because windowing is only needed where the signal length is greater than the length of the audio cue. Thus, in generating the audio cue, windowing is not needed because the audio cue is provided *in isolation*. For at least this reason, Murphy’s system does not (and cannot) derive multiple fingerprints from the “sample of the audio cue” provided by the user. Finally, Murphy’s Figure 2 depicts the process of *matching* signatures, not deriving audio cue signatures to be matched later. *Murphy*, 5:3-4 (Figure 2 depicts processing an audio signal by the “recognition system”); 9:5-7 (windowing depicted in Figure 2 not needed in audio cue signature generation).

46. It is also my opinion that Murphy fails to teach this claim limitation because Murphy does not teach deriving fingerprints based on time points and, in particular, does not teach deriving a fingerprint based on a *received* time point. In particular, as I described above, Murphy is reliant on an isolated version of the

audio cue to derive the signature (because, as discussed, Murphy does not perform windowing when generating signatures for audio cues). Murphy thus generates the signature based on the *entire* received sample and has not need of (or use for) a time point indicating a segment of the received signal from which to derive a signature. In other words, Murphy's system is reliant on the user to manually isolate a sample for signature generation rather than automatically generating the signature based on the content as a whole and a received time point. As such, Murphy's recognition system does not have any trigger time point information as to where, in the broadcast audio signal, the audio cue appeared. Furthermore, as discussed above, in Murphy's system, no time point is associated with or used to generate or derive an audio cue signature. *See* ¶ 37, *supra*.

J. Murphy does not teach “receiving a second trigger time point,” as required by Claims 1, 8, and 15 of the '962 Patent

47. Claims 1, 8, and 15 recite the step of “receiving a second trigger time point of the plurality of time points.” *See* '962 Patent, claims 1, 8, 15. Murphy does not teach this step because Murphy's audio cue signatures are not associated with any time point. In particular, and as discussed above, Murphy teaches that the *same* audio cue signature can cause action at *different* times, such as starting recording and stopping recording. As such, it is unclear what time point *would be* received for an audio cue signature that is used to both start and stop recording. As such, it does not make sense for Murphy's system to “receive a second time point” (or any

time point), since Murphy's system does not make use of time points in the signal in its operation. *See* ¶ 37, *supra*.

48. Furthermore, Murphy does not include any disclosure of "receiving" a time point. The only information Murphy described being provided to the system for use in generating an audio cue signature is the audio cue itself. *Murphy*, 14:12. I note that Dr. Tewfik claims that the user tells the system "to start recording (second action) at the beginning of the desired programme (second trigger time point)." *Tewfik Dec.*, ¶ 152. This is unsupported in the disclosure of Murphy. In particular, as described above, Murphy only states that the user provides the audio cue. Audio cues can trigger different actions (start recording and stop recording) at different times (at the beginning and end of each episode of the program). Thus, when a user provides an audio cue to the system, it is unclear how the user indicates (even implicitly) that the audio cue is to be used "to start recording (second action) at the beginning of the desired programme (second trigger time point)," at least because the same audio cue is also used to *stop* recording at the *end* of the desired program. Rather, the audio cue identifies the program, which for which the user instructs recording by providing the associated code. *Murphy*, 13:30-33.

K. Murphy does not teach “receiving ... a second action,” as required by Claims 1, 8, and 15 of the ’962 Patent

49. Claims 1, 8, and 15 also recite the step of “receiving ... a second action.” *See* ’962 Patent, claims 1, 8, 15. It is my opinion that Murphy does not teach or suggest this limitation. First, I understand that Petitioner admits that Murphy does not teach or suggest storing the action in template database 12. At least for this reason, it would not make sense for Murphy to receive an action because it would have nothing to do with the action once received. I understand that Dr Tewfik contends that Murphy’s “appropriate audio cue information,” stored in the template database 12, includes a “corresponding trigger time point and action.” *Tewfik Dec.*, ¶ 151. However, Murphy makes it clear that this information is the PTY (program type) information, rather than any action: “Information concerning the programme, such as the programme type, with which the audio cues are associated is also stored.” *Murphy*, 7:7-11, 6:30-32, 10:7-9. As such, I would understand Murphy to teach that the only information received from the user is (1) the pre-processed audio cue (i.e., the fingerprint) (*Murphy*, 14:6-9); or (2) the audio cue to be pre-processed (*Murphy*, 14:10-11, 14:11-15). I would also understand Murphy to teach that the only information stored in the template database is audio cue signature data and program type information.

50. Furthermore, Murphy does not teach “receiving an action” with an audio cue signature because, as described above, the same audio cue signature can

cause multiple different actions. Murphy's system cannot distinguish between the different times these audio cues are matched, because the same signature is matched more than once. This is why Murphy's system is reliant on context (i.e., whether it is already recording) to determine whether to start recording or stop recording when it detects an audio cue. As such, Murphy is not reliant on action information provided by the user and, because it reuses audio cues for different purposes, could not use action information provided by the user if it were received.

L. Murphy does not teach “the action being associated with a time point indicating when, in the multimedia content, the action is to be performed” as required by claims 1, 11, 19, and 24 of the '831 Patent.

51. Claims 1, 11, 19, and 24 of the '831 Patent each require an action to be associated with a reference fingerprint and they require the action to be performed in response to determining a match; but, importantly, these claims also each require “the action being associated with a time point indicating when, in the multimedia content, the action is to be performed.” Murphy fails to teach this requirement of “the action being associated with a time point indicating when, in the multimedia content, the action is to be performed.” In particular, Murphy takes actions based on detecting event flags generated in response to audio cues being recognized. No time points are associated with the actions because Murphy makes no use of time points. *See* ¶ 37, *supra*.

52. I understand that Petitioner contends that Murphy teaches actions being triggered before or after the location of the audio cue signature in the audio signal by virtue of the time-delay unit disclosed as a part of the transmitter embodiment. *Petition*, 36 (“Murphy also teaches that actions can be triggered by a trigger fingerprint at time points before or after the time point of the trigger fingerprint in the signal.”). This would, however, be contrary to the express teachings of Murphy, which teaches that actions should be taken in “real time.” *Murphy*, 4:16-24. Incorporating a time-delay unit to allow Murphy’s system to take actions after the fingerprint is recognized would add additional delay to the processing delays inherent in Murphy’s system. This would thus be *adding* delay and rendering signal generation *less* real time and less preferable under the teachings of Murphy.

53. Murphy also fails to teach that the stored audio cue signatures are associated with any particular action to take responsive to detecting the audio cue. As described above, Murphy’s actions are *context-dependent*: “Once a start flag has been generated, an end flag could be generated on the repeated occurrence of the audio cue or of the occurrence of a second, play-out, audio cue.” *Murphy*, 11:1-4. I would understand these “play-out cues” to operate like any other cue, causing an end flag to be generated if a start flag has already been generated and causing a start flag to be generated otherwise. Thus, Murphy teaches that event flag

generation is dependent on the context of what previously occurred in the recognition system. Particular actions in Murphy are not taken because they are associated with a signature detected, but rather based on whether or not a start flag has already been generated.

54. I understand that the Petition contends in the alternative that Murphy in combination with Brunk teaches audio cue signatures associated with an action to be taken responsive to detecting that audio cue. *Petition*, 35. In particular, the Petition states that “in the combined system when a match is detected, the corresponding action is received from database.” However, modifying Murphy in this way would result in an inoperable system. In particular, Murphy’s stated goal is to be able to stop recording “on the repeated occurrence of an audio cue.” *Murphy*, 11:2-3. However, Murphy is reliant on audio cues (i.e. jingles) already present in a program. Thus, the combined system would cause Murphy’s system to take *the same* action (for example, start recording) every time the signature was matched. Murphy’s system, as modified, would thus start recording at the beginning of the program and attempt start recording *again* at the repeated occurrence of the jingle at the end of the program, contrary to the stated goal of Murphy.

M. Murphy does not teach “a plurality of reference fingerprints having been previously derived” from the multimedia content being played back, as required by claim 1 of the ’831 Patent

55. Claim 1c recites the step of “accessing a plurality of reference fingerprints, **each reference fingerprint** among the plurality of reference fingerprints **having been previously derived from a respective segment of the multimedia content.**” “*The multimedia content,*” thus, refers to the same multimedia content recited in Claim 1a that is being “play[ed] back” that is subsequently used to determine trigger fingerprints during playback by the playback device. This limitation, therefore, requires the reference fingerprints be previously derived from the same multimedia content that is being “play[ed] back.” The Petitioner relies only upon on Murphy to teach this requirement, but Murphy does not teach or suggest it. In particular, as discussed above, I would understand none of the disclosures of Murphy to teach or suggest deriving the audio cue signatures from the content that is actually played over the loudspeaker. *See ¶ 38, supra.*

[REMAINDER OF PAGE INTENTIONALLY LEFT BLANK]

I declare that all statements made herein of my knowledge are true, and that all statements made on information and belief are believed to be true, and that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Date: March 9, 2020

By: 
Dr. Pierre Moulin

APPENDIX A

PIERRE MOULIN

University of Illinois
Coordinated Science Laboratory
1308 W. Main St.
Urbana, IL 61801

tel (217) 244-8366
fax (217) 244-8371
Email: moulin@ifp.uiuc.edu
WWW: www.ifp.uiuc.edu/~moulin

Education D.Sc. in Electrical Engineering, *Washington University, St Louis* 08/1987–05/1990
M.Sc. in Electrical Engineering, *Washington University, St Louis* 08/1985–05/1986
Ingénieur Civil Electricien, *Faculté Polytechnique de Mons, Belgium* 09/1979–07/1984

Professional Experience

Professor of ECE and Statistics, *University of Illinois, Urbana, IL* 08/2003–present
Associate Professor, *University of Illinois, Urbana, IL* 08/1999–08/2003
Assistant Professor, *University of Illinois, Urbana, IL* 01/1996–08/1999
Research Scientist, *Bell Communications Research, Morristown, NJ* 10/1990–12/1995
Research Assistant, *Washington University, St Louis, MO* 10/1987–05/1990
Research Engineer, *Faculté Polytechnique de Mons, Belgium* 10/1984–06/1985
Summer Intern, *Standard Telecommunications Laboratories Ltd., U.K.* 07/1983–08/1983

Professional Activities

Editorial Board member for *Proceedings of IEEE*, 2007-2012
Co-Founder and Editor in Chief for *IEEE Transactions on Information Forensics and Security*, 2005-2008
Area Editor for *IEEE Transactions on Image Processing*, 2002-2006.
Guest Editor for *IEEE Transactions on Signal Processing* supplements on secure media, 2004-2005.
Guest Associate Editor for *IEEE Transactions on Signal Processing's* special issue on Data Hiding, Apr. 2003.
Guest Associate Editor for *IEEE Transactions on Information Theory's* special issue on Information-theoretic imaging, Aug. 2000
Associate Editor for *IEEE Transactions on Image Processing*, 1999–2002
Associate Editor for *IEEE Transactions on Information Theory*, 1996–1998
Member, IEEE Image and Multidim. Signal Processing Technical Committee, 1998-2003
Co-chair, Allerton Conference 2007, 2008
Member, Technical Program Committee of miscellaneous major IEEE conferences
Invited organizer, NSF Workshop on Signal Authentication, Orlando, FL, 2002.
Organizer, special sessions on watermarking at Allerton (10/01, 02, 03) and on information-theoretic imaging at Asilomar (11/02)
Invited lectures at MIT, Stanford, Berkeley, Carnegie-Mellon, Harvard, Illinois, Purdue, Michigan, UC San Diego, UC Santa Barbara, Maryland, National U. of Singapore, Nanyang Tech. U. (Singapore), Chinese U. Hong Kong, Hong Kong U. of Science and Technology, Hong Kong Polytechnic, National Taiwan U., Academia Sinica, National Chiao Tung U. (Taiwan), KAIST (Korea), INRIA Sophia-Antipolis (France), Delft (Netherlands), Louvain (Belgium), Tech. U. Munich (Germany), Bell Labs, IBM Research, Microsoft Research, KLA-Tencor, Qualcomm.
Keynote/Plenary talks at Int. Workshop on Digital Watermarking (IWDW), Seoul, Korea, 2002; Int. Symp. on Image and Signal Processing and Analysis (ISPA), Rome, 2003; Wav-ila Challenge, Barcelona, Spain, 2005; IEEE Int. Conf. on Acoustics, Speech and Signal

Processing (ICASSP), Toulouse, France, 2006; Workshop on Multimedia Content Representation, Classification and Security (MRCS), Istanbul, Turkey, 2006; 3rd Int. Conf. on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP), Kaohsiung, Taiwan, 2007; IEEE Int. Conf. on Image Processing (ICIP), Brussels, Belgium, 2011; EUROCON, Zagreb, Croatia, 2013; ICSIPA, Melaka, Malaysia, 2013; VCIP, Singapore, 2015; ISPACS, Taipei, 2019.

Tutorial lecturer at ICIP'01, ICASSP'02, ICIP'04, and ISIT'06.

Co-Chair, IEEE Info Theory Workshop in Detection, Classif. & Imaging (Santa Fe, 2/99); Allerton Conf. (Monticello, IL, 9/06, 9/07); IEEE Workshop on Information Forensics and Security (Tenerife, Spain, 11/12); Tech. Program co-chair, IEEE Symp. on Information Theory (Hong Kong, 6/15).

Member, *IEEE Int. Conf. on Image Processing* Organizing Committee (Chicago 1998; Technical Program Co-chair, Brussels 2011)

Reviewer and panelist for National Science Foundation, reviewer for IEEE Trans. on SP, IP, IT, COM, NN and JOSA, CVGIP, JVCIP, Addison-Wesley, IEEE Press, SIAM, and Annals of Statistics.

Consultant and expert witness

Founding Director of *Center for Information Forensics*, UIUC.

Honors and Awards

Ronald W. Pratt Faculty Outstanding Teaching Award, 2018
IEEE Information Theory Society Board of Governors, 2016—2018
IEEE Signal Processing Society Distinguished Lecturer, 2012–2013
UIUC Sony Faculty Scholar, 2005—2007
IEEE Signal Processing Society Board of Governors, 2005—2007
IEEE Fellow, 2003
Beckman Associate in UIUC Center for Advanced Study, 2003
Co-author, IEEE Signal Processing Society 2002 Young Author Best Paper Award
IEEE Signal Processing Society 1997 Best Paper Award
NSF Career award, 1998-2001
UIUC Incomplete List of Teachers Rated as Excellent, 1996, 1999, 2000, 2005, 2007, 2009, 2011, 2012, 2016, 2019
Rotary International Scholar, 1985-1986
A. Dosin Award Recipient, 1984
Co-winner 1989 US Amateur Team Chess Championship

University Service

ECE Promotions and Tenure Committee, 2009-Present
Chair, ECE Signal Processing and Circuits Area, 2007-2017
ECE Faculty Search Committee, 2006—2012 and 2015—Present
ECE Graduate Committee, 1997—2008 and 2011—2014
ECE Representative to Research Board, 2007-2008
Elected Member, UIUC Faculty Senate, 2004-2005
Member, Information Trust Institute educational committee, 2004 - 2006
Elected Member, ECE Faculty Advisory Committee, 2002-2003
Member, Computer Resources and Education Committee, 1999—2002
Chair, ECE Graduate Seminar Committee, 1998—1999
Member, ECE Graduate Seminar Committee, 1996—2003
Ad hoc subcommittee to evaluate various courses
Organizer, weekly Digital Signal Processing seminar series, 1997.

Ph.D. Graduates

	Name	Year	Placement
1.	Juan (Julia) Liu	2001	Research Scientist, PARC, Palo Alto, CA
2.	Prakash Ishwar	2002	Prof., Boston U.
3.	Kivanc Mihcak	2002	Assoc Prof., Bogazici U., Istanbul, Turkey
4.	Shawn C. Herman	2002	Numerica Inc., Fort Collins, CO.
5.	Tie Liu	2006	Assoc. Prof., Texas AM
6.	Ying (Grace) Wang	2006	Flarion Qualcomm, NJ
7.	Negar Kiyavash (co-advisor)	2006	Assoc. Prof., UIUC
8.	Maha El Choubassi	2008	Intel Research, Santa Clara, CA
9.	Shankar Sadasivam	2011	Qualcomm, San Diego, CA
10.	Yen-Wei Huang	2013	Google
11.	Honghai Yu	2015	A*STAR, Singapore
12.	Patrick Johnstone	2017	Postdoc, Rutgers U.

Postdocs

	Name	Year	Placement
1.	Aaron Lanterman	1998-2001	Prof, Georgia Tech
2.	Sviatoslav Voloshnykovskiy	1999	Assoc. Prof., U. of Geneva
3.	Jong C. Ye	1999-2001	Prof, KAIST, Korea
4.	Natalia Schmid	2001	Assoc. Prof., U. of West Virginia
5.	Bing Bing Ni	2010-2015	Asst. Prof., Shanghai Jiaotong U., China
6.	Jiwen Lu	2012-2015	Asst. Prof., Tsinghua U., Beijing, China
7.	Gang Wang	2012-2015	Assoc. Prof., Nanyang Tech. U., Singapore
8.	Zhang Le	2016-	ADSC Singapore
9.	Jagan Varadarajan	2013 - 2016	ADSC Singapore

Funding

NSF, DARPA, ARO, ONR, NATO, A*STAR, Microsoft, HP Labs, Xerox, KLA-Tencor

Sabbatical Leaves

National Taiwan University, 2019

Chinese University of Hong Kong, 2010 and 2014

MIT, 2005

Microsoft Research, 2001

Course Development

1. ECE418, Introduction to Image and Video Processing, 1996-1998
2. ECE544 Wavelets, 1998
3. ECE544 Statistical Image and Video Processing, 2006-2008
4. ECE544 Statistical Learning, 2012, 2015, 2016
5. ECE598PM (now ECE566) Computational Inference and Learning, 2014-2017
6. ECE398BD (now ECE314) Making Sense of Big Data, 2014.

Other Teaching

Introduction to Image and Video Processing (ECE418), Digital Signal Processing II (ECE551), Information Theory (ECE563), Random Processes (ECE534), Detection and Estimation Theory (ECE561)

Journal Papers

- P. R. Johnstone and P. Moulin, "Faster Subgradient Methods for Functions with Hölderian Growth," *Mathematical Programming*, Vol. 180, pp. 418—450, 2020.
- R. Trabelsi, J. Varadarajan, L. Zhang, I. Jabri, Y. Pei, F. Smach, A. Bouallegue, and P. Moulin, "Understanding the Dynamics of Social Interactions: A Multi-Modal Multi-View Approach" *ACM Transactions on Multimedia Computing, Communications, and Applications*, Feb. 2019.
- R. W. Yeung, C. Chen, Q. Chen, and P. Moulin, "On Characterizations of Markov Random Fields and Subfields," *IEEE Trans. Information Theory*, Aug. 2018.
- P. Moulin, "The Log-Volume of Optimal Codes on Memoryless Channels, Asymptotically Within a Few Nats," *IEEE Trans. Information Theory*, Vol. 63, pp. 2278—2313, April 2017.
- P. R. Johnstone and P. Moulin, "Local and global convergence of a general inertial proximal splitting scheme for minimizing composite functions," *Computational Optimization and Applications*, Feb. 2017.
- J. Lu, G. Wang and P. Moulin, "Localized Multi-Feature Metric Learning for Image Set Recognition," *IEEE Trans. Circ. Syst. Video Technology*, Vol. 26, No. 3, pp. 529—540, March 2016.
- B. Ni, V. R. Paramathayalan, T. Li, and P. Moulin, "Multiple Granularity Modeling: A Coarse-to-Fine Framework for Fine-grained Action Analysis," *International Journal on Computer Vision (IJCV)*, March 2016.
- H. Yu and P. Moulin, "SNR Maximization Hashing," *IEEE Trans. Information Forensics and Security*, Vol.10, No. 9, pp. 1927—1938, Sep. 2015.
- B. Ni, G. Wang, and P. Moulin, "Order Preserving Sparse Coding" *IEEE Trans. on Pattern Recognition and Machine Intelligence*, Aug. 2015, Vol. 37, No. 8, pp. 1615—1628.
- H. Yu and P. Moulin, "Regularized Adaboost Learning for Identification of Time-Varying Content," *IEEE Trans. Information Forensics and Security*, Vol. 9, No. 10, pp. 1606—1616, Oct. 2014.
- Y.-W. Huang and P. Moulin, "On the Fingerprinting Capacity Games for Arbitrary Alphabets and Their Asymptotics," *IEEE Trans. Information Forensics and Security*, Vol. 9, No. 9, pp. 1477—1490, Sep. 2014.
- B. Ni, P. Moulin and S. Yan, "Pose Adaptive Motion Feature Pooling for Human Action Analysis," *International Journal on Computer Vision (IJCV)*, July 2014.
- J. Lu, G. Wang and P. Moulin, "Human Identity and Gender Recognition from Gait Sequences with Arbitrary Walking Directions," *IEEE Trans. on Information Forensics and Security*, Vol. 9, No. 1, pp. 41—51, Jan. 2014.
- B. Ni, Y. Pei, S. Yan, and P. Moulin, "Multi-Level Depth and Image Fusion for Human Activity Detection," *IEEE Transactions on System, Man and Cybernetics, Part B (TSMC-B)*, Vol. 43, No. 5, pp.1383—1394, 2012.
- Y.-W. Huang and P. Moulin, "On the Saddle-point Solution and the Large-Coalition Behavior of Fingerprinting Games," *IEEE Trans. on Information Forensics and Security*, Vol. 7, No. 1, pp. 160—175, 2012.
- S. Sadasivam, P. Moulin and T. P. Coleman, "A Message Passing Approach to Combating

Desynchronization Attacks,” *IEEE Trans. on Information Forensics and Security*, Vol. 6, pp. 894–905, 2011.

J.-F. Jourdas and P. Moulin, “High-Rate Random-Like Fingerprinting Codes with Linear Decoding Complexity,” *IEEE Transactions on Information Forensics and Security*, Vol. 4, No. 4, Dec. 2009.

N. Kiyavash, P. Moulin and T. Kalker, “Regular Simplex Fingerprints and Their Optimality Properties,” *IEEE Transactions on Information Forensics and Security*, Vol. 4, No. 3, pp. 318–329, Sep. 2009.

N. Kiyavash and P. Moulin, “Performance of Orthogonal Fingerprints Under Worst-Case Noise Attacks,” *IEEE Transactions on Information Forensics and Security*, Vol. 4, No. 3, pp. 293–301, Sep. 2009.

M. El Choubassi and P. Moulin, “On Reliability and Security of Randomized Detectors Against Sensitivity Analysis Attacks,” *IEEE Trans. Information Forensics and Security*, Vol. 4, No. 3, pp. 273–283, Sep. 2009.

S. Sadasivam and P. Moulin, “On Estimation Accuracy of Desynchronization Attack Channel Parameters,” *IEEE Transactions on Information Forensics and Security*, Vol. 4, No. 3, pp. 284–292, Sep. 2009.

P. Moulin, “A Neyman-Pearson Approach to Universal Erasure and List Decoding,” *IEEE Trans. Information Theory*, Oct. 2009.

Y. Wang and P. Moulin, “Perfectly Secure Steganography: Capacity, Error Exponents, and Code Constructions,” *IEEE Transactions on Information Theory*, Special Issue on Security, Vol. 54, No. 6, pp. 2706–2722, June 2008.

M. El Choubassi and P. Moulin, “Noniterative Algorithms for Sensitivity Analysis Attacks,” *IEEE Trans. Information Forensics and Security*, Vol. 2, No.3, pp. 113–126, June 2007.

P. Moulin and Y. Wang, “Capacity and Random-Coding Exponents for Channel Coding with Side Information,” *IEEE Trans. on Information Theory*, Vol. 53, No. 4, pp. 1326–1347, Apr. 2007.

Y. Wang and P. Moulin, “Optimized Feature Extraction for Learning-Based Image Steganalysis,” *IEEE Trans. Information Forensics and Security*, Vol. 2, No. 1, pp. 31–45, March 2007.

P. Moulin, “Signal Transmission with Known-Interference Cancellation,” *IEEE Signal Processing Magazine*, Lecture Notes, Vol. 24, No. 1, pp. 134–136, Jan. 2007.

J. C. Ye, P. Moulin and Y. Bresler, “Asymptotic Global Confidence Regions for Parametric 3-D Shape Estimation,” *IEEE Trans. Image Processing*, Vol. 15, No. 10, pp. 2904–2919, Oct. 2006.

P. Moulin and A. K. Goteti, “Block QIM Watermarking Games,” *IEEE Trans. Information Forensics and Security*, Vol. 1, No. 3, pp. 293–310, Sep. 2006.

S. Jana and P. Moulin, “Optimality of KLT for Encoding Gaussian Vector-Scale Mixtures: Application to Reconstruction, Estimation and Classification,” *IEEE Trans. on Information Theory*, Vol. 52, No. 9, pp. 4049–4067, Sep. 2006.

T. Liu, P. Moulin and R. Koetter, “On Error Exponents of Modulo Lattice Additive Noise Channels,” *IEEE Trans. on Information Theory*, Vol. 52, No. 2, pp. 454–471, Feb. 2006.

- P. Moulin and R. Koetter, "Data-Hiding Codes," invited tutorial paper, *Proc. IEEE*, Vol. 93, No. 12, pp. 2083–2127, Dec. 2005.
- P. Ishwar and P. Moulin, "On the Existence and Characterization of the Maxent Distribution Under General Moment Inequality Constraints," *IEEE Trans. on Information Theory*, Vol. 51, No. 9, pp. 3322–3333, Sep. 2005.
- J. L. Cannons and P. Moulin, "Design and Statistical Analysis of a Hash-Aided Image Watermarking System," *IEEE Trans. on Image Processing*, Vol. 13, No. 10, pp. 1393–1408, Oct. 2004.
- P. Moulin and M. K. Mihçak, "The Parallel-Gaussian Watermarking Game," *IEEE Trans. on Information Theory*, Vol. 50, No. 2, pp. 272–289, Feb. 2004.
- P. Moulin, "Comments on "Why Watermarking is Nonsense," *IEEE Signal Processing Magazine*, Vol. 20, No. 6, pp. 57–59, Nov. 2003.
- P. Moulin and A. Ivanović, "The Zero-Rate Spread-Spectrum Watermarking Game," *IEEE Transactions on Signal Processing*, Vol. 51, No. 4, pp. 1098–1117, Apr. 2003.
- P. Moulin and J. A. O'Sullivan, "Information-Theoretic Analysis of Information Hiding," *IEEE Trans. on Information Theory*, Vol. 49, No. 3, pp. 563–593, March 2003.
- P. Ishwar and P. Moulin, "On the Equivalence Between Set-Theoretic and Maxent MAP Estimation," *IEEE Trans. on Signal Processing*, Vol. 51, No. 3, pp. 698–713, March 2003.
- J. C. Ye, Y. Bresler and P. Moulin, "Cramer-Rao Bounds for Parametric Shape Estimation in Inverse Problems," *IEEE Trans. on Image Processing*, Vol. 12, No. 1, pp. 71–84, Jan. 2003.
- J. C. Ye, Y. Bresler and P. Moulin, "A Self-Referencing Level-Set Method for Image Reconstruction from Sparse Fourier Samples," invited paper, *Int. J. of Computer Vision*, special issue on level-set methods, Dec. 2002.
- P. Moulin and M. K. Mihçak, "A Framework for Evaluating the Data-Hiding Capacity of Image Sources," *IEEE Trans. on Image Processing*, Vol. 11, No. 9, pp. 1029–1042, Sep. 2002.
- A. Jain, P. Moulin, M. I. Miller and K. Ramchandran, "Information-Theoretic Bounds on Target Recognition Performance Based on Degraded Image Data," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 24, No. 9, pp. 1153–1166, Sep. 2002.
- J. Liu and P. Moulin, "Information-Theoretic Analysis of Interscale and Intrascale Dependencies Between Image Wavelet Coefficients," *IEEE Trans. on Image Processing*, Vol. 10, No. 10, pp. 1647–1658, Nov. 2001.
- P. Moulin, invited discussion of "Regularization of Wavelet Approximations," by A. Antoniadis and J. Fan, *Journal of the American Statistical Association*, Vol. 96, No. 455, pp. 959–960, Sep. 2001.
- P. Moulin, "The Role of Information Theory in Watermarking and Its Application to Image Watermarking," invited paper, *Signal Processing*, Vol. 81, No. 6, pp. 1121–1139, June 2001.
- J. Liu and P. Moulin, "Complexity-Regularized Image Denoising," *IEEE Trans. on Image Processing*, Vol. 10, No. 6, pp. 841–851, June 2001.
- J. C. Ye, Y. Bresler and P. Moulin, "Cramer-Rao Bounds for Parametric Boundaries of

Targets in Inverse Scattering Problems,” *IEEE Trans. on Antennas and Propagation*, May 2001.

M. K. Mihçak, P. Moulin, M. Anitescu, and K. Ramchandran, “Rate–Distortion–Optimal Subband Coding Without Perfect Reconstruction Constraints,” *IEEE Trans. on Signal Processing*, Vol. 49, No. 3, pp. 542–557, Mar. 2001.

P. Ishwar and P. Moulin, “On Spatial Adaptation of Motion Field Smoothness in Video Coding,” *IEEE Trans. Circ. Syst. Video Tech.*, Vol. 10, No. 6, pp. 980–989, Sep. 2000.

P. Moulin and J. Liu, “Statistical Imaging and Complexity Regularization,” *IEEE Trans. on Information Theory*, Special issue on information-theoretic imaging, Vol. 46, No. 5, pp. 1762–1777, Aug. 2000.

J. C. Ye, Y. Bresler and P. Moulin, “Asymptotic Global Confidence Regions in Parametric Shape Estimation Problems,” *IEEE Trans. on Information Theory*, Special issue on information-theoretic imaging, Vol. 46, No. 5, pp. 1881–1895, Aug. 2000.

P. Moulin, M. Anitescu, and K. Ramchandran, “Theory of Rate–Distortion–Optimal, Constrained Filter Banks — Application to FIR and IIR Biorthogonal Designs,” *IEEE Trans. on Signal Processing*, Vol. 48, No. 4, pp. 1120–1132, April 2000.

M. K. Mihçak, I. Kozintsev, K. Ramchandran and P. Moulin, “Low-Complexity Image Denoising Based on Statistical Modeling of Wavelet Coefficients,” *IEEE Signal Processing Letters*, Vol. 6, No. 12, pp. 300–303, Dec. 1999.

R. Krishnamurthy, J. W. Woods and P. Moulin, “Frame Interpolation and Bidirectional Prediction of Video Using Compactly-Encoded Optical Flow Fields and Label Fields,” *IEEE Trans. Circ. Syst. Video Tech.*, Vol. 9, No. 5, pp. 713–726, Aug. 1999.

P. Moulin and J. Liu, “Analysis of Multiresolution Image Denoising Schemes Using Generalized–Gaussian and Complexity Priors,” *IEEE Trans. on Info. Theory*, Special Issue on Multiscale Analysis, Vol. 45, No. 3, pp. 909–919, Apr. 1999.

V. Pavlovic, P. Moulin and K. Ramchandran, “An Integrated Framework for Adaptive Subband Image Coding,” *IEEE Trans. on Signal Processing*, Vol. 47, No. 4, pp. 1024–1038, Apr. 1999.

P. Moulin and M. K. Mihçak, “Theory and Design of Signal-Adapted FIR Paraunitary Filter Banks,” *IEEE Trans. on Signal Processing*, Special Issue on Wavelets and Filter Banks, Vol. 46, No. 4, pp. 920–929, Apr. 1998.

P. Moulin, R. Krishnamurthy, “Multiscale Modeling and Estimation of Motion Fields for Video Coding,” *IEEE Trans. on Image Processing*, Vol. 6, No. 12, pp. 1606–1620, Dec. 1997.

P. Moulin, M. Anitescu, K.O. Kortanek and F. Potra, “The Role of Linear Semi–Infinite Programming in Signal–Adapted QMF Bank Design,” *IEEE Trans. on Signal Processing*, Vol. 45, No. 9, pp. 2160–2174, Sep. 1997.

P. Moulin, “A Multiscale Relaxation Technique for SNR Maximization in Nonorthogonal Subband Coding,” *IEEE Trans. on Image Processing*, Vol. 4, No. 9, pp. 1269–1281, Sep. 1995.

P. Moulin, A.T. Ogielski, G. Lilienfeld and J.W. Woods: “Video Signal Processing and Coding on Data–Parallel Computers,” *Digital Signal Processing: A Review Journal*, Vol. 5, No. 2, pp. 118–129, Apr. 1995.

P. Moulin, invited discussion of “Wavelet Shrinkage: Asymptotia?” by D. Donoho and I. Johnstone, *Journal of the Royal Statistical Society B*, Vol. 57, No. 1, 1995.

P. Moulin, “Wavelet Thresholding Techniques for Power Spectrum Estimation,” *IEEE Trans. on Signal Processing*, Vol. 42, No. 11, pp. 3126–3136, Nov. 1994.

P. Moulin, “A Wavelet Regularization Method for Diffuse Radar–Target Imaging and Speckle–Noise Reduction,” *Journal of Mathematical Imaging and Vision*, Special Issue on Wavelets, Vol. 3, No. 1, pp. 123–134, 1993.

J. A. O’Sullivan, P. Moulin, and D. L. Snyder, “An Application of Splines to Maximum Likelihood Radar Imaging,” *International Journal of Imaging Systems and Technology*, Vol. 4, pp. 256–264, 1993.

P. Moulin, J. A. O’Sullivan, and D. L. Snyder, “A Method of Sieves for Multiresolution Spectrum Estimation and Radar Imaging,” *IEEE Trans. on Information Theory*, Special Issue on Wavelets and Multiresolution Analysis, pp. 801–813, Mar. 1992.

Book

P. Moulin and V. Veeravalli, *Statistical Inference for Engineers and Data Scientists*, Cambridge University Press, 2019.

Book Chapters

B. Ni, G. Wang, and P. Moulin, “RGBD-HuDaAct: A color-depth video database for human daily activity recognition,” *Consumer Depth Cameras for Computer Vision: Advances in Computer Vision and Pattern Recognition*, pp 193-208, Springer-Verlag, 2013.

P. Moulin, “Information-Hiding Games,” Springer-Verlag Lecture Notes in Computer Sciences, Vol. 2613, 2003.

P. Moulin, “Multiscale Image Decompositions and Wavelets,” *Handbook of Image and Video Processing*. A. C. Bovik, Ed., Academic Press, 2000.

K.O. Kortanek and P. Moulin, “Using Semi–Infinite Programming in Orthogonal Wavelet Filter Design,” *Semi–Infinite Programming and Its Applications*, Kluwer Academic series on Nonconvex Optimization and its Applications, 1998.

P. Moulin, J. A. O’Sullivan, and D. L. Snyder, “A Sieve–Constrained Maximum–Likelihood Method for Target Imaging,” in *Radar and Sonar II*, Springer–Verlag, The IMA Volumes in Mathematics and Its Applications, Vol. 39, Eds. F. Grunbaum et al., pp. 95–122, 1992.

P. Moulin, “Adaptive Multiresolution Image Restoration and Compression,” in *IEEE Case Studies in Medical Instrument Design*, Eds. H. Troy Nagle and W.J. Tomkins, pp. 247–254, IEEE, New York, 1992.

Conference Papers

F. Zhang and P. Moulin, “A New Variational Method for Deep Supervised Image Hashing,” to appear in Proc. ICASSP, Barcelona, Spain, May 2020.

P. Moulin, “Game-theoretic design of universal adversarial perturbation detectors,” presented at *Information Theory and Applications*, San Diego, CA, Feb. 2020.

T. Jayashankar, P. Moulin, T. Blu, C. Gilliam, LAP-Based Video Frame Interpolation, *Proc. ICIP 2019*, Taipei, Taiwan, Sep. 2019.

P. Moulin and A. Goel, “Random Ensemble of Locally Optimum Detectors for Detection of Adversarial Inputs,” *IEEE GlobalSIP Conf., Signal Processing for Adversarial Machine Learning*, Anaheim, CA, Nov. 2018.

- P. Moulin, "Gaussian Mixture Models for Detection of Adversarial Inputs," *Information Theory and Applications*, San Diego, CA, Feb. 2018.
- P. Moulin and A. Goel, "Locally Optimal Detection of Adversarial Inputs to Image Classifiers," *Proc. ICME*, Hong Kong, July 2017.
- R. W. Yeung, C. Chen, Q. Chen, and P. Moulin, Information-Theoretic Characterizations of Markov Random Field and Subfield, *Proc. ISIT*, Aachen, Germany, June 2017.
- P. Moulin, "Source Coding with Distortion Profile Constraints", *Proc. ISIT*, Aachen, Germany, June 2017.
- P. Moulin, "Lower Bounds on Rate of Fixed-Length Source Codes under Average- and ϵ -Fidelity Constraints," *Proc. ISIT*, Aachen, Germany, June 2017.
- Zhang Le, J. Varadarajan, P. N. Suganthan, N. Ahuja, and P. Moulin, "Robust Visual Tracking Using Incremental Oblique Random Forest," *CVPR*, June 2017.
- J. Varadarajan, R. Subramanian, N. Ahuja, P. Moulin, J.-M. Odobez, "Active Online Anomaly Detection Using Dirichlet Process Mixture Model and Gaussian Process Classification," *Proc. Winter Conference on Applications of Computer Vision*, Santa Rosa, CA, March 2017.
- P. R. Johnstone and P. Moulin, "Convergence Rates of Inertial Splitting Schemes for Nonconvex Composite Optimization," *Proc. ICASSP*, New Orleans, LA, March 2017.
- P. Moulin, "Refined asymptotic analysis of source coding under excess-distortion constraint," *ITA*, San Diego, CA, Feb. 2017.
- P. Moulin, "Asymptotically Achievable Error Probabilities for Multiple Hypothesis Testing," *Proc. IEEE Int. Symp. Information Theory*, Barcelona, Spain, July 2016.
- P. Moulin, "Strong Large Deviations for Composite Multiple Hypothesis Testing," *ITA*, San Diego, CA, Feb. 2016.
- H. Yu and P. Moulin, "Multi-Feature Hashing Based on SNR Maximization," *Proc. ICIP*, Quebec City, Canada, Sep. 2015.
- T. Blu, P. Moulin, and C. Gilliam, "Approximation Order of the LAP Optical Flow Algorithm," *Proc. ICIP*, Quebec City, Canada, Sep. 2015.
- P. Moulin, "Coding Redundancy of Finite-Blocklength Universal Channel Coding," *Proc. IEEE Int. Symp. Information Theory*, Hong Kong, June 2015.
- P. Moulin and P. R. Johnstone, "Strong Large Deviations for GLRT and Variants in Composite Hypothesis Testing," *Proc. IEEE Int. Symp. Information Theory*, Hong Kong, June 2015.
- J. Lee, M. Raginsky, and P. Moulin, "On MMSE Estimation from Quantized Observations in the Nonasymptotic Regime," *Proc. IEEE Int. Symp. Information Theory*, Hong Kong, June 2015.
- B. Ni and P. Moulin, "Motion Part Regularization: Improving Action Recognition via Trajectory Selection," *Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- J. Lu, G. Wang, W. Deng, P. Moulin, and J. Zhou, "Multi-Manifold Deep Learning for Image Set Classification," *Proc. CVPR*, June 2015.

- V. E. Liong, J. Lu, G. Wang, P. Moulin, and J. Zhou, “Deep Hashing for Compact Binary Codes Learning,” *Proc. CVPR*, June 2015.
- H. Yu and P. Moulin, “SNR Maximization Hashing,” *Proc. ICASSP*, Brisbane, Australia, Apr. 2015.
- P. R. Johnstone and P. Moulin, “Convergence of an Inertial Proximal Method for L_1 Regularized Least Squares,” *Proc. ICASSP*, Brisbane, Australia, Apr. 2015.
- S. D. Chen and P. Moulin, “A Classification Centric Quantizer for Efficient Encoding of Predictive Feature Errors,” *Proc. Asilomar Conf.*, Monterrey, CA, Nov. 2014.
- J. Lu, G. Wang, W. Deng, and P. Moulin, “Simultaneous feature learning and dictionary learning for image set base face recognition,” *ECCV*, Lecture Notes in Computer Science Vol. 8689, pp. 265—280, Zurich, Switzerland, Sep. 2014.
- Y. Zhou, B. Ni, S. Yan, P. Moulin, and Q. Tian, “Pipelining Localized Semantic Features for Fine-Grained Action Recognition,” *ECCV*, Lecture Notes in Computer Science Vol. 8692, pp. 481—496, Zurich, Switzerland, Sep. 2014.
- V. Tan and P. Moulin, “Second- and Higher-Order Asymptotics For Erasure and List Decoding,” *Proc. ISIT*, Honolulu, HI, pp. 1887—1891, July 2014.
- Y.-W. Huang and P. Moulin, “Strong Large Deviations for Composite Hypothesis Testing,” *Proc. ISIT*, Honolulu, HI, pp. 556—560, July 2014.
- B. Ni, T. Li, and P. Moulin, “Beta Process Multiple Kernel Learning,” *CVPR*, Columbus, OH, pp. 963—970, June 2014.
- B. Ni, V. R. Paramathayalan and P. Moulin, “Multiple Granularity Analysis for Fine-Grained Action Detection,” *CVPR*, Columbus, OH, pp. 756—763, June 2014.
- S. D. Chen and P. Moulin, “A Two-Part Predictive Coder for Multitask Signal Compression,” *Proc. ICASSP*, Florence, Italy, pp. 2035—2039, May 2014.
- R. Naini and P. Moulin, “Fingerprint Information Maximization for Content Identification,” *Proc. ICASSP*, Florence, Italy, pp. 3809—3813, May 2014.
- P. Moulin and P. R. Johnstone, “Kullback-Leibler Divergence and the Central Limit Theorem,” *Proc. ITA*, San Diego, CA, Feb. 2014.
- J. Lu, G. Wang, and P. Moulin, “Image set classification using holistic multiple order statistics features and localized multi-kernel metric learning,” *IEEE International Conference on Computer Vision (ICCV)*, pp. 329—336, 2013.
- B. Ni and P. Moulin, “Manipulation Pattern Discovery: A Nonparametric Bayesian Approach,” *ICCV*, 2013.
- P. Moulin, “Asymptotic Neyman-Pearson Games for Converse to the Channel Coding Theorem,” *Proc. ISIT*, Istanbul, July 2013.
- P. R. Johnstone, A. Emad, P. Moulin, and O. Milenkovic, “RFIT: A New Algorithm for Matrix Completion,” *Proc. SPARS*, Lausanne, Switzerland, July 2013.
- D. Sungatullina, J. Lu, G. Wang, and P. Moulin, “Multiview discriminative learning for age-invariant face recognition,” *IEEE Int. Conf. on Automatic Face and Gesture Recognition*, pp. 1—6, Shanghai, China, Apr. 2013.
- H. Yu and P. Moulin, “Regularized Adaboost for Content identification,” *Proc. ICASSP*,

Vancouver, Canada, May 2013.

H. Yu, P. Moulin and S. Roy, "RGB-D Video Content Identification," *Proc. ICASSP*, Vancouver, Canada, May 2013.

D. Sungatullina, J. Lu, G. Wang, and P. Moulin, "Multiview Discriminative Learning for Age-Invariant Face Recognition," *Proc. 10th IEEE International Conference on Automatic Face and Gesture Recognition*, Shanghai, China, Apr. 2013.

P. Moulin, "A New Metaconverse and Outer Region for Finite-Blocklength MACs," *Proc. Information Theory and Applications (ITA) Workshop*, San Diego, CA, Feb. 2013.

A. Singh, N. Ahuja, P. Moulin, "Online learning with kernels: Overcoming the growing sum problem," *Proc. IEEE Machine Learning for Signal Processing (MLSP)*, Santander, Spain, pp. 449-454, Sep. 2012.

B. Ni, M. Xu, J. Tang, S. Yan, and P. Moulin, "Omni-range spatial contexts for visual classification," *Proc. CVPR*, pp. 3514-3521, Providence, RI, June 2012.

B. Ni, P. Moulin, S. Yan, "Order-Preserving Sparse Coding for Sequence Classification," *Proc. ECCV* pp. 173-187, Providence, RI, June 2012.

B. Ni, N. C. Dat, P. Moulin, "RGBD-camera based get-up event detection for hospital fall prevention," *Proc. ICASSP*, pp. 1405-1408, Kyoto, Japan, March 2012.

R. Dubey, B. Ni, and P. Moulin, "A Depth Camera Based Fall Recognition System for the Elderly," *Proc. 9th Int. Conf. Image Analysis and Recognition (ICIAR)*, LNCS Vol. 7325, pp. 106-113, Aveiro, Portugal, June 2012.

P. Moulin, "The log-volume of optimal constant-composition codes for memoryless channels, within $O(1)$ bits," *Proc. ISIT*, pp. 826-830, Cambridge, MA, July 2012.

Y.-W. Huang and P. Moulin, "Finite blocklength coding for multiple access channels," *Proc. ISIT*, pp. 831-835, Cambridge, MA, July 2012.

Y.-W. Huang and P. Moulin, "On fingerprinting capacity games for arbitrary alphabets and their asymptotics," *Proc. ISIT*, pp. 2571-2575, Cambridge, MA, July 2012.

R. Naini and P. Moulin, "Real Adaboost for content identification," *Proc. IEEE Stat. Signal Processing Workshop*, Ann Arbor, MI, Aug. 2012.

R. Naini and P. Moulin, "Model-based decoding metrics for content identification," *Proc. ICASSP*, pp. 1829-1832, Kyoto, Japan, March 2012.

P. Moulin, "The log-volume of optimal codes for memoryless channels, Up to a Few Nats," *Proc. Information Theory and Applications Workshop*, San Diego, CA, Feb. 2012.

B. Ni, G. Wang, and P. Moulin, "RGBD-HuDaAct: A color-depth video database for human daily activity recognition," *Proc. 1st IEEE workshop on Consumer Depth Cameras for Computer Vision*, held in conjunction with *Int. Conf. on Computer Vision (ICCV) Workshop*, pp. 1147-1153, Barcelona, Nov. 2011.

S. Sadasivam and P. Moulin, "A Universal Divergence-Rate Estimator for Steganalysis in Timing Channels," *Proc. 2nd Workshop on Information Security and Forensics*, Seattle, WA, Dec. 2010.

Y.-W. Huang and P. Moulin, "On the Saddle-point Solution and the Large-Coalition Behavior of Fingerprinting Games," *Proc. 2nd Workshop on Information Security and Forensics*, Seattle, WA, Dec. 2010.

- R. Naini and P. Moulin, "Interference Management through Mobile Relays in Ad Hoc Networks," *Proc. ASILOMAR Conf.*, Pacific Grove, CA, Nov. 2010.
- G. Gigaud and P. Moulin, "A Geometrically Resilient SURF-Based Image Fingerprinting Scheme," *Proc. IEEE Int. Conf. Image Processing (ICIP)*, Hong Kong, Sep. 2010.
- G. Gigaud and P. Moulin, "Traitor Tracing Aided by Compressed SURF Features," *Proc. Conf. Info. Sci. and Sys.*, Princeton, NJ, March 2010.
- P. Moulin, "Statistical Modeling and Analysis of Content Identification and Retrieval," *Proc. Information Theory and Applications Workshop*, San Diego, CA, Feb. 2010.
- Y.-W. Huang and P. Moulin, "Capacity-Achieving Fingerprint Decoding," *Proc. 1st Workshop on Information Security and Forensics*, London, UK, Dec. 2009.
- S. Sadasivam and P. Moulin, "Combating Desynchronization Attacks on Blind Watermarking Systems: A Message Passing Approach," *Proc. ICIP*, Cairo, Egypt, Nov. 2009.
- I. Ezzeddine and P. Moulin, *Proc. IEEE Information Theory Workshop*, Taormina, Italy, Oct. 2009.
- S. D. Chen, V. Monga, and P. Moulin, "Meta-classifiers For Multimodal Document Classification," *Proc. IEEE Multimedia Signal Processing Workshop*, Rio de Janeiro, Brazil, Oct. 2009.
- Y.-W. Huang and P. Moulin, Saddlepoint Solution of the Fingerprinting Capacity Game under the Marking Assumption, *Proc. Int. Symp. Information Theory*, Seoul, Korea, July 2009.
- P. Moulin, "A Strong Converse for the Gel'fand-Pinsker Channel," *Proc. Int. Symp. Information Theory*, Seoul, Korea, July 2009.
- P. Moulin, "Optimal Gaussian Fingerprint Decoders" *Proc. ICASSP*, Taipei, Taiwan, Apr. 2009.
- P. Moulin and Y. Wang, "Information-Theoretic Analysis of Spherical Fingerprinting," *Proc. ITA Workshop*, San Diego, CA, Feb. 2009.
- M. El Choubassi and P. Moulin, "Joint Estimation-Detection Games for Sensitivity Analysis Attacks," *Proc. SPIE Conf.*, San Jose, CA, Jan. 2009.
- P. Moulin, "A Neyman-Pearson Approach to Universal Erasure and List Decoding," *Proc. IEEE Symp. Information Theory*, pp. 61—65, July 2008.
- P. Moulin, "Universal Fingerprinting: Capacity and Random-Coding Exponents," *Proc. IEEE Symp. Information Theory*, pp. 220—224, July 2008.
- J.-F. Jourdas and P. Moulin, "Towards Optimal Design of Multimedia Fingerprinting Codes," *Proc. ICASSP*, Las Vegas, NV, pp. 1649—1652, April 2008.
- J.-F. Jourdas and P. Moulin, "A Low-Rate Fingerprinting Code And Its Application To Blind Image Fingerprinting," *Proc. SPIE*, San Jose, CA, Jan. 2008.
- J.-F. Jourdas and P. Moulin, "A High-Rate Fingerprinting Code," *Proc. SPIE*, San Jose, CA, Jan. 2008.
- S. Sadasivam, P. Moulin, and R. Koetter, "Graphical Models for Desynchronization-Resilient Watermark Decoding," *Proc. IEEE Int. Conf. on Image Proc.*, San Antonio, TX, pp. V.477—480, Sep. 2007.

- M. El Choubassi and P. Moulin, "Sensitivity Analysis Attacks Against Randomized Detectors," *Proc. IEEE Int. Conf. on Image Proc.*, San Antonio, TX, pp. II.129—132, Sep. 2007.
- N. Kiyavash and P. Moulin, "Digital Fingerprinting: on Colluders Worst-Case Noise," *IEEE Workshop on Statistical Signal Processing*, Madison, WI, July 2007.
- P. Moulin and N. Kiyavash, "Expurgated Gaussian Fingerprinting Codes," *Proc. IEEE Int. Symp. on Information Theory*, Nice, France, June 2007.
- P. Moulin and N. Kiyavash, "Performance of Random Fingerprinting Codes Under Arbitrary Nonlinear Attacks," *Proc. ICASSP*, Honolulu, HI, pp. II.157—160, April 2007.
- N. Kiyavash and P. Moulin, "Sphere Packing Lower Bound on Fingerprinting Error Probability," *Proc. SPIE*, San Jose, CA, Jan. 2007.
- Y. Wang and P. Moulin, "Capacity and Optimal Collusion Attack Channels for Gaussian Fingerprinting Games," *Proc. SPIE*, San Jose, CA, Jan. 2007.
- S. Sadasivam and P. Moulin, "Cramer-Rao Bound on Watermark Desynchronization Parameter Accuracy," *Proc. SPIE*, San Jose, CA, Jan. 2007.
- L. Perez-Freire, P. Moulin, and F. Perez-Gonzalez, "Security of spread-spectrum-based data hiding," in *Security, Steganography, and Watermarking of Multimedia Contents IX*, Edward J. Delp III and P. W. Wong, Eds., vol. 6505. San Jose, California, USA: SPIE, January 2007.
- P. Moulin, "On the Optimal Structure of Watermark Decoders Under Geometric Attacks," *Proc. IEEE Int. Conf. on Image Proc.*, Atlanta, GA, Oct. 2006.
- P. Moulin, "Universal Decoding of Watermarks Under Geometric Attacks" *Proc. IEEE Int. Symp. on Information Theory*, July 2006.
- Y. Wang and P. Moulin, "Capacity and Random-Coding Error Exponent for Public Fingerprinting Game," *Proc. IEEE Int. Symp. on Information Theory*, July 2006.
- S. Ray, P. Moulin and M. Médard, "On Jamming in the Wideband Regime," *Proc. IEEE Int. Symp. on Information Theory*, July 2006.
- S. Ray, P. Moulin and M. Médard, "On Optimal Signaling and Jamming Strategies in Wideband Fading Channels," *Proc. IEEE Conf. on Signal Processing Applications for Wireless Communications (SPAWC)*, June-July 2006.
- N. Kiyavash and P. Moulin, "On Optimal Collusion Strategies for Fingerprinting", *Proc. ICASSP*, Toulouse, France, May 2006.
- N. Kiyavash and P. Moulin, "A Framework for Optimizing Nonlinear Collusion Attacks on Fingerprinting Systems", *Proc. CISS*, Princeton, NJ, March 2006.
- P. Moulin and R. Koetter, "A Framework for Design of Good Watermark Identification Codes," *Proc. SPIE*, San Jose, CA, Jan. 2006.
- M. El Choubassi and P. Moulin, "On the fundamental tradeoff between watermark detection performance and robustness against sensitivity analysis attacks," *Proc. SPIE*, San Jose, CA, Jan. 2006.
- Y. Wang and P. Moulin, "Statistical modeling and steganalysis of DFT-based image steganography," *Proc. SPIE*, San Jose, CA, Jan. 2006.

- N. Kiyavash and P. Moulin, "Regular Simplex Fingerprints and Their Optimality Properties," *Proc. Int. Workshop on Digital Watermarking*, Siena, Italy, Sep. 2005.
- P. Moulin and Y. Wang, "Improved QIM Strategies for Gaussian Watermarking," *Proc. Int. Workshop on Digital Watermarking*, Siena, Italy, Sep. 2005.
- P. Moulin and A. K. Goteti, "Minmax Strategies for QIM Watermarking Subject to Attacks with Memory," *Proc. IEEE ICIP*, Sep. 2005.
- P. Moulin and Y. Wang, "On Achievable Error Exponents for Watermarking," *Proc. SPIE Conf.*, San Jose, CA, Jan. 2005
- M. El Choubassi and P. Moulin, "A New Sensitivity Analysis Attack," *Proc. SPIE Conf.*, San Jose, CA, Jan. 2005
- A. K. Goteti and P. Moulin, "QIM Watermarking Games," *Proc. ICIP*, Singapore, Oct. 2004.
- P. Moulin and A. Briassouli, "A Stochastic QIM Algorithm for Robust, Undetectable Image Watermarking," invited paper, *Proc. ICIP*, Singapore, Oct. 2004.
- T. Liu, P. Moulin and R. Koetter, "On Error Exponents of Nested Lattice Codes for the AWGN Channel," *Proc. Information Theory Workshop*, San Antonio, TX, Oct. 2004.
- P. Moulin and Y. Wang, "Error Exponents for Channel Coding with Side Information," *Proc. Information Theory Workshop*, San Antonio, TX, Oct. 2004.
- P. Moulin, A. K. Goteti and R. Koetter, "Optimal Sparse QIM Codes for Zero-Rate Blind Watermarking," *Proc. ICASSP*, Montreal, Canada, May 2004.
- A. K. Goteti and P. Moulin, "Two Private, Perceptual Data-Hiding Games," *Proc. ICASSP*, Montreal, Canada, May 2004.
- P. Moulin and Y. Wang, "New Results on Steganographic Capacity," *Proc. Conf. on Info. Syst. and Sci.*, Princeton, NJ, March 2004.
- Y. Wang and P. Moulin, "Steganalysis of Block-Structured Stegotext," *Proc. SPIE Vol. 5306*, San Jose, CA, Jan. 2004.
- P. Moulin, "Embedded-Signal Design for Channel Parameter Estimation Part I: Linear Embedding," *Proc. IEEE Workshop on Statistical Signal Processing*, St Louis, MO, Sep. 2003.
- P. Moulin, "Embedded-Signal Design for Channel Parameter Estimation Part II: Quantization Embedding," *Proc. IEEE Workshop on Statistical Signal Processing*, St Louis, MO, Sep. 2003.
- Y. Wang and P. Moulin, "Steganalysis of Block-DCT Steganography," *Proc. IEEE Workshop on Statistical Signal Processing*, St Louis, MO, Sep. 2003.
- J. Liu and P. Moulin, "On the Risk of Translation-Averaged Wavelet Estimators," *Proc. SPIE*, San Diego, Aug. 2003.
- S. Jana and P. Moulin, "Optimal Transform Coders for Classification/Reconstruction with Joint Decoding at the Receiver," *Proc. IEEE Int. Symp. on Information Theory*, Yokohama, Japan, July 2003.
- T. Liu and P. Moulin, "Error Exponents for Zero-Rate Watermarking Game with Squared-Error Constraints," *Proc. IEEE Int. Symp. on Information Theory*, Yokohama, Japan,

July 2003.

T. Liu and P. Moulin, "Error Exponents for One-Bit Watermarking," *IEEE Proc. ICASSP*, Hong Kong, April 2003.

A. Briassouli and P. Moulin, "Detection-Theoretic Analysis of Warping Attacks in Spread-Spectrum Watermarking," *IEEE Proc. ICASSP*, Hong Kong, April 2003.

S. Jana and P. Moulin, "Optimal Transform Coding of Gaussian Mixtures for Joint Classification/Reconstruction," *Proc. Data Compression Conference*, Snowbird, UT, March 2003.

N. Schmid, Y. Bresler and P. Moulin, "Complexity Regularized Shape Estimation from Noisy Fourier Data," *Proc. IEEE Int. Conf. on Image Processing*, Rochester, NY, 2002.

J. C. Ye, Y. Bresler and P. Moulin, "Cramer-Rao Bounds for Parametric Shape Estimation in Inverse Problems," *Proc. IEEE Int. Conf. on Image Processing*, Rochester, NY, 2002.

A. Ivanović and P. Moulin, "Game-Theoretic Performance Analysis of Image Watermark Detectors," *Proc. IEEE Int. Conf. on Image Processing*, Rochester, NY, 2002.

M. K. Mihçak and P. Moulin, "Information-Embedding Codes Matched to Local Gaussian Image Models," *Proc. IEEE Int. Conf. on Image Processing*, Rochester, NY, 2002.

P. Moulin, A. Briassouli and H. Malvar, "Detection-Theoretic Analysis of Desynchronization Attacks in Watermarking," *Proc. DSP'02*, Santorini, Greece, July 2002.

P. Moulin and M. T. Orchard, "Undergraduate Education in Image and Video Processing," *Proc. Int. Conf. on Ac., Sp., and Signal Processing*, Orlando, FL, May 2002.

S. C. Herman and P. Moulin, "A Particle Filtering Approach to Joint Radar Tracking and Automatic Target Recognition", *Proc. IEEE Aerospace Conference*, Big Sky, Montana, March 2002.

P. Moulin and A. Briassouli, "The Gaussian Fingerprinting Game," *Proc. CISS'02*, Princeton, NJ, March 2002.

P. Moulin and J. A. O'Sullivan, "Optimal Key Design in Information-Embedding Systems," *Proc. CISS'02*, Princeton, NJ, March 2002.

P. Moulin and A. Ivanović, "The Fisher Information Game for Optimal Design of Synchronization Patterns in Blind Watermarking," *Proc. IEEE Int. Conf. on Image Processing*, Thessaloniki, Greece, Oct. 2001.

P. Moulin and A. Ivanović, "Game-Theoretic Analysis of Watermark Detection," *Proc. IEEE Int. Conf. on Image Processing*, Thessaloniki, Greece, Oct. 2001.

J. C. Ye, Y. Bresler and P. Moulin, "A Self-Referencing Level-Set Method for Image Reconstruction from Sparse Fourier Samples," *Proc. IEEE Int. Conf. on Image Processing*, Thessaloniki, Greece, Oct. 2001.

J. Liu and P. Moulin, "Statistical Image Restoration Based on Adaptive Wavelet Models," *Proc. IEEE Int. Conf. on Image Processing*, Thessaloniki, Greece, Oct. 2001.

J. Liu and P. Moulin, "Approximation-Theoretic Analysis of Translation-Invariant Wavelet Denoising," *Proc. IEEE Int. Conf. on Image Processing*, Thessaloniki, Greece, Oct. 2001.

J. A. O' Sullivan and P. Moulin, "Some Properties of Optimal Information Hiding and Information Attacks," *Proc. 39th Allerton Conf.*, pp. 92-101, Monticello, IL, Oct. 3-5,

2001.

J. C. Ye, Y. Bresler and P. Moulin, "A Self-Referencing Level-Set Method for Image Reconstruction from Sparse Fourier Samples," *Proc. IEEE Int. Workshop on Variational and Level Set Methods*, Vancouver, B.C., July 2001.

P. Moulin and A. Ivanović, "The Watermark Selection Game," *Proc. Conf. on Info. Syst. and Sci.*, Baltimore, MD, March 2001.

P. Moulin, "The Parallel-Gaussian Watermarking Game," *Proc. Conf. on Info. Syst. and Sci.*, Baltimore, MD, March 2001.

J. Liu and P. Moulin, "Approximation-Theoretic Analysis of Translation-Invariant Wavelet Denoising," invited paper, *Proc. Conf. on Info. Syst. and Sci.*, Baltimore, MD, March 2001.

P. Moulin, M. K. Mihçak and G.-I. Lin, "An Information-Theoretic Model for Image Watermarking," *IEEE Int. Conf. on Image Processing*, Vancouver, BC, Sep. 2000.

R. Venkatesan, M. Jakubowski, S.-M. William Koon and P. Moulin, "Robust Image Hashing," *IEEE Int. Conf. on Image Processing*, Vancouver, BC, Sep. 2000.

P. Ishwar and P. Moulin, "Shift-Invariant Image Denoising — A Multiple-Domain Maxent MAP Framework," *IEEE Int. Conf. on Image Processing*, Vancouver, BC, Sep. 2000.

J. Liu and P. Moulin, "Analysis of Interscale and Intrascale Dependencies Between Image Wavelet Coefficients," *IEEE Int. Conf. on Image Processing*, Vancouver, BC, Sep. 2000.

S. Jana and P. Moulin, "Optimal Design of Transform Coders and Quantizers for Image Classification," *IEEE Int. Conf. on Image Processing*, Vancouver, BC, Sep. 2000.

M. K. Mihçak, A. F. Docimo, P. Moulin and K. Ramchandran, "Design and Analysis of a Forward-Adaptive Wavelet Image Coder," *IEEE Int. Conf. on Image Processing*, Vancouver, BC, Sep. 2000.

M. Kesal, M. K. Mihçak, R. Kötter and P. Moulin, "Iteratively Decodable Codes for Watermarking Applications," *Proc. 2nd Symposium on Turbo Codes and Related Topics*, pp. 113-116, Brest, France, Sep. 2000.

P. Ishwar and P. Moulin, "Image Modeling and Restoration – Information fusion, Set-theoretic methods, and the Maximum Entropy Principle," *Proc. MAXENT'00*, Paris, France, July 2000.

P. Moulin and J. A. O'Sullivan, "Information-Theoretic Analysis of Information Hiding," *Proc. IEEE Int. Symp. on Information Theory*, Sorrento, Italy, June 2000.

P. Moulin and J. Liu, "Statistical Imaging and Complexity Regularization," *Proc. IEEE Int. Symp. on Information Theory*, Sorrento, Italy, June 2000.

P. Moulin and J. A. O'Sullivan, "Information-Theoretic Analysis of Watermarking," *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'00)*, Istanbul, Turkey, June 2000.

P. Ishwar and P. Moulin, "Fundamental Equivalences between Set-Theoretic and Maximum-Entropy methods in Multiple-Domain Image Restoration," *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'00)*, Istanbul, Turkey, June 2000.

J. C. Ye, Y. Bresler and P. Moulin, "Global Confidence Regions in Parametric Shape Estimation," *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'00)*,

Istanbul, Turkey, June 2000.

A. Jain, P. Moulin, M. I. Miller and K. Ramchandran, "Information-Theoretic Bounds on Target Recognition Performance," *SPIE's 14th International Symposium on AeroSense 2000*, Orlando, FL, April 2000.

J. C. Ye, Y. Bresler and P. Moulin, "Wavelet Descriptors for Parametric Shape Estimation in Nonlinear Inverse Scattering Problems," *SPIE's International Symposium on AeroSense 2000*, Orlando, FL, April 2000.

P. Ishwar and P. Moulin, "Multiple-Domain Image Modeling and Restoration," *Proc. IEEE Int. Conf. on Image Proc.*, Vol. I, pp. 362–366, Kobe, Japan, 1999.

J. Liu and P. Moulin, "Image Denoising Based on Scale-Space Mixture Modeling of Wavelet Coefficients," *Proc. IEEE Int. Conf. on Image Proc.*, Vol. I, pp. 386–390, Kobe, Japan, 1999.

S. Jana and P. Moulin, "Optimal Design of Transform Coders for Image Classification," *Proc. Conf. on Info. Syst. and Sci.*, Baltimore, MD, March 1999.

A. Jain, P. Moulin, K. Ramchandran and M. I. Miller, "Performance Bounds for ATR based on Compressed Data," *Proc. Army Compression Workshop*, Huntsville, AL, Dec. 1998.

P. Moulin, M. Anitescu, and K. Ramchandran, "Some Fundamental Properties of MMSE Filter Banks," *Proc. IEEE Sig. Proc. Symp. on Time-Frequency and Time-Scale Analysis*, Pittsburgh, PA, Oct. 1998.

P. Moulin and J. Liu, "Analysis of Multiresolution Image Denoising Schemes Using Generalized - Gaussian Priors," *Proc. IEEE Sig. Proc. Symp. on Time-Frequency and Time-Scale Analysis*, Pittsburgh, PA, Oct. 1998.

J. Liu and P. Moulin, "A New Complexity Prior for Multi-Resolution Image Denoising," *Proc. IEEE Sig. Proc. Symp. on Time-Frequency and Time-Scale Analysis*, Pittsburgh, PA, Oct. 1998.

J. Liu and P. Moulin, "Complexity-Regularized Image Restoration," *Proc. IEEE Int. Conf. on Image Proc. (ICIP'98)*, Chicago, IL, Oct. 1998.

M. K. Mihçak, K. Ramchandran and P. Moulin, "Adaptive Wavelet Packet Image Coding Using an Estimation-Quantization Framework," *Proc. IEEE Int. Conf. on Image Proc. (ICIP'98)*, Vol. I, pp. 92–96, Chicago, IL, Oct. 1998.

P. Moulin and M. Miller, "Detection Bounds for Automatic Target Recognition," *Proc. Int. Symp. on Information Theory (ISIT'98)*. Boston, MA, Aug. 1998.

J. A. O'Sullivan, P. Moulin and M. Ettinger, "Information Theoretic Analysis of Steganography," *Proc. Int. Symp. on Information Theory (ISIT'98)*, Boston, MA, Aug. 1998.

X. Yang, K. Ramchandran and P. Moulin, "Optimization of High-Energy-Compaction, Nearly-Orthonormal, Linear-Phase Filter Banks," *Proc. ISCAS'98*, Monterrey, CA, June 1998.

P. Ishwar, K. Ratakonda, P. Moulin and N. Ahuja, "Image Denoising in Multiple Compaction Domains," invited paper, *Proc. ICASSP'98*, Vol. 3, pp. 1889–1892, Seattle, May 1998.

P. Moulin, M. Anitescu and K. Ramchandran, "Asymptotic Performance of Subband Coders Using Constrained, Signal-Adapted FIR Filter Banks," *Proc. 32nd Conf. on*

Info. Syst. and Sci. (CISS), Princeton, NJ, March 1998.

J. Liu and P. Moulin, "Complexity-Regularized Image Denoising," *Proc. IEEE Int. Conf. on Image Proc. (ICIP'97)*, pp. II. 370—373, Santa Barbara, CA, Oct. 1997.

P. Ishwar and P. Moulin, "Switched Control-Grid Interpolation for Motion-Compensated Video Coding," *Proc. IEEE Int. Conf. on Image Proc. (ICIP'97)*, pp. III. 650—653, Santa Barbara, CA, Oct. 1997.

R. Krishnamurthy, J. W. Woods and P. Moulin, "Compactly-Encoded Optical Flow Fields for Bidirectional Prediction and Video Coding," *Proc. VCIP'97*, Feb. 1997.

P. Moulin, K. Ramchandran and V. Pavlovic, "Transform Image Coding based on Joint Adaptation of Filter Banks and Tree Structures," *Proc. IEEE Int. Conf. on Image Proc. (ICIP'96)*, Lausanne, Switzerland, Sep. 1996.

R. Krishnamurthy, P. Moulin and J. W. Woods, "Multiscale Motion Estimation for Scalable Video Coding," *Proc. ICIP'96*, Lausanne, Switzerland, pp. I. 965—968, Sep. 1996.

P. Moulin, "Signal Estimation Using Adapted Tree-Structured Bases and the MDL Principle," *Proc. IEEE Sig. Proc. Symp. on Time-Frequency and Time-Scale Analysis*, pp. 141—143, Paris, June 1996.

P. Moulin, "Design of Signal-Adapted FIR Paraunitary Filter Banks," *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'96)*, Atlanta, GA, May 1996.

P. Moulin and R. Krishnamurthy, "Multiscale Motion Fields for Video Coding," *Proc. 9th IEEE Workshop on Image and Multidim. Signal Processing*, Belize, March 1996.

R. Krishnamurthy, P. Moulin and J. Woods, "Optical Flow Techniques Applied to Video Coding," *Proc. ICIP'95*, Washington, D.C., Oct. 1995.

P. Moulin, "Model Selection Criteria and the Orthogonal Series Method for Function Estimation," *Proc. IEEE Int. Symp. on Information Theory (ISIT'95)*, Whistler, B.C., Sep. 1995.

P. Moulin, "A New Look at the Signal-Adaptive QMF Bank Design," *Proc. ICASSP'95*, Detroit, MI, pp. 1312—1315, Apr. 1995.

P. Moulin, "A Relaxation Technique for Minimizing the L^2 Reconstruction Error in Nonorthogonal Subband Coding," *Proc. ICIP'94*, Austin, TX, Nov. 1994.

P. Moulin, "The Use of Wavelets for Spectral Density Estimation with Local Bandwidth Adaptation," *Proc. ISIT'94*, p. 40, Trondheim, Norway, June 1994. Also invited paper, *Interface'94*, Research Triangle Park, June 1994.

P. Moulin, "Optimization Techniques for Efficient Image Compression in Nonorthogonal Multiresolution Bases," *Proc. 1994 NJIT Symp. on Wavelets*, Apr. 1994.

P. Moulin, "Wavelets as a Regularization Technique for Spectral Density Estimation," *Proc. SIAM Annual Meeting*, p. 69, Philadelphia, July 1993.

P. Moulin, I. Rhee and P. Caputo, "Video Signal Processing on SIMD Parallel Computer Architectures," *Proc. 8th IEEE Workshop on Image and Multidim. Signal Processing*, pp. 246—247, Cannes, France, Sep. 1993.

P. Moulin and A. Loui, "Hierarchical-Finite-Element Estimation and Parameterization of Optical Flow Field," *Proc. 8th IEEE Workshop on Image and Multidim. Signal Processing*, pp. 130—131, Cannes, France, Sep. 1993.

- P. Moulin and A. Loui: "Application of a Multiresolution Optical Flow Based Method for Motion Estimation to Video Coding," *Proc. ISCAS'93*, pp. I.1–4, Chicago, IL, May 1993.
- P. Moulin, "Optimal L^2 Approximation of Images in Nonorthogonal Multiresolution Bases," *Proc. ICASSP'93*, pp. V.189–192, Minneapolis, MN, Apr. 1993.
- P. Moulin, "Model-Based Multiresolution Restoration of Speckle Images: Application to Radar Imaging," invited paper, *Proc. IEEE Int. Symp. on Information Theory*, p. 67, San Antonio, TX, Jan. 1993.
- P. Moulin, "Adaptive Multiresolution Image Restoration and Compression," invited paper, *IEEE EMBS Workshop on Case Studies in Medical Instrument Design*, Orlando, FA, Oct. 91.
- P. Moulin, "Wavelets as a Regularization Technique for Spectral Density Estimation," *Proc. IEEE Signal Processing Conf. on Time-Frequency and Time-Scale Analysis*, pp. 73–76, Victoria, B.C., Oct. 1992.
- P. Moulin, "An Adaptive Finite-Element Method for Image Representation," *Proc. Int. Conf. on Pattern Recognition*, pp. III.70–74, The Hague, The Netherlands, Aug. 1992.
- P. Moulin, "An Adaptive Spline Method for Signal Restoration," *Proc. ICASSP'92*, San Francisco, CA, Mar. 92, pp. IV.169–172.
- P. Moulin, J. A. O'Sullivan, and D. L. Snyder, "A Method of Sieves for Estimating the Spectral Density," *Proc. IEEE Int. Symp. on Information Theory*, p. 67, Budapest, Hungary, June 91.
- J. A. O'Sullivan, P. Moulin, D. L. Snyder, and S. P. Jacobs: "Computational Considerations for Maximum-Likelihood Radar Imaging," *Proc. 24th CISS Conf.*, Princeton, NJ, March 1990.
- P. Moulin, D. L. Snyder, and J. A. O'Sullivan: "A Method of Sieves for Regularizing Maximum-Likelihood Spectral Estimates," *Proc. IEEE Int. Symp. on Information Theory*, p. 112, San Diego, CA, Jan. 1990.
- P. Moulin, D. L. Snyder, and J. A. O'Sullivan: "A Sieve-Constrained Maximum-Likelihood Estimator for the Spectrum of a Gaussian Process," *Proc. 28th Allerton Conf.*, Urbana-Champaign, IL, Sept. 1989.
- J. A. O'Sullivan, D. L. Snyder, and P. Moulin, "The Role of Spectrum Estimation in Forming High-Resolution Radar Images," *Proc. ICASSP 1989*, Glasgow, U.K., May 1989.
- P. Moulin, J. A. O'Sullivan, and D. L. Snyder, "Maximum-Likelihood Spectrum Estimation of Periodic Processes from Noisy Data," *Proc. CISS Conf.*, Johns Hopkins University, Baltimore, MD, March 1989.
- J. A. O'Sullivan, P. Moulin, and D. L. Snyder, "Cramer-Rao Bounds for Constrained Spectrum Estimation with Application to a Problem in Radar Imaging," *Proc. 27th Allerton Conf.*, Urbana-Champaign, IL, Sept. 1988.
- R. Boite and P. Moulin, "Comparaison du rapport signal à bruit pour une représentation numérique en virgule fixe et en virgule flottante," *Proc. IASTED Symp.*, Paris, France, June 1985.

Conference Papers Submitted for Publication

A. Goel and P. Moulin, "Locally optimal detection of stochastic universal adversarial perturbations," submitted to ICIP, 2020.