Seeing the Invisibles:
A Backstage Tour of Information Forensics

Min Wu

Media and Security Team (MAST)
ECE Department / UMIACS
University of Maryland, College Park

http://www.ece.umd.edu/~minwu/research.html

Include joint research with Wei-Hong Chuang, Adi Hajj-Ahmad, Ravi Garg, Hongmei Gou, Shan He, K.J. Ray Liu, Christine McKay, Hui Su, Ashwin Swaminathan, Wade Trappe, Avinash Varna, Jane Wang, Chau-Wai Wong, and Hong Zhao.

THE A. JAMES CLARK SCHOOL *of* ENGINEERING | UNIVERSITY OF MARYLAND

A Decade+ of Research on Info. Forensics

- Our life and work are forever changed …
  - By advances in electronics, computing, communications, sensing, and signal processing
- So much multimedia info and content right at our hand …
  - Keep us entertained
  - Used as important evidences & records
- Gather traces of evidence
  - Origin, history, integrity, etc.
  - A broad range of applications

Role Play as Sherlock Holmes in our Digital Era

1. Leak: A proprietary image sent to 10 people got leaked out

   ➔ Who leaked the info?

Role Play as Sherlock Holmes in our Digital Era

1. Leak: A proprietary image sent to 10 people got leaked out

   ➔ Who leaked the info?

2. Source: Picture of a heavily guarded xPhone7 prototype showed up on web

   ➔ Is it a real photo or a graphic rendition? Who in the company took it using his/her camera?

3. When/Where: an audio clip with incriminating words showed up

   ➔ were its content and recording time true as claimed?

Role Play as Sherlock Holmes in our Digital Era

1. Leak: A proprietary image sent to 10 people got leaked out
   ➔ Who leaked the info?
2. Source: Picture of a heavily guarded xPhone7 prototype showed up on web
   ➔ Is it a real photo or a graphic rendition? Who in the company took it using his/her camera?
3. When/Where: an audio clip with incriminating words showed up
   ➔ were its content and recording time true as claimed?

Role Play as Sherlock Holmes in our Digital Era

1. Leak: A proprietary image sent to 10 people got leaked out
   ➔ Who leaked the info?
2. Source: Picture of a heavily guarded xPhone7 prototype showed up on web
   ➔ Is it a real photo or a graphic rendition? Who in the company took it using his/her camera?
3. When/Where: an audio clip with incriminating words showed up
   ➔ were its content and recording time true as claimed?

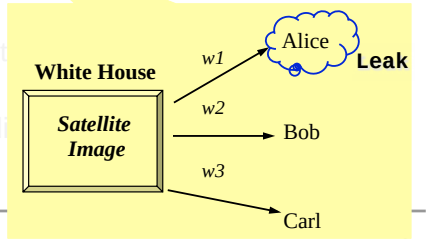

Example: Replacing the Last Bit of Pixels

Replace LSB with Pentagon's MSB

Robust Watermarking via Spread Spectrum Embedding

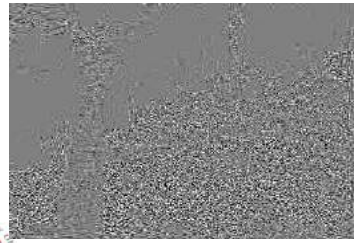

10011010 …

© Copyright …

- *Embedding domain tailored to media characteristics & application requirement*

Fingerprinting Topographic Map

- Traditional protection: intentionally alter geospatial content
- Embed much less intrusive digital fingerprint for a modern protection

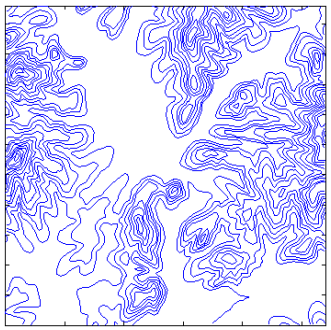

1100x1100 Original Map

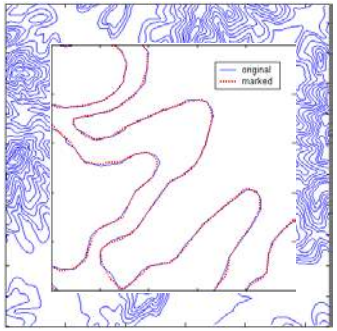

Fingerprinted Map

• 9 long curves are marked; 1331 control points used to carry the fingerprint

Collusion-Resistant Fingerprinting: Examples

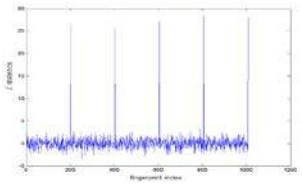

5-User Averaging Attack

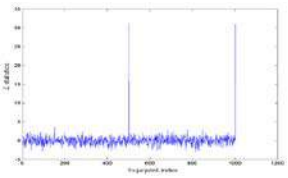

2-User Interleaving Attack

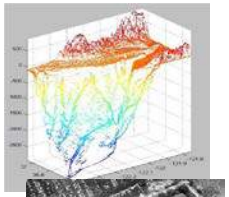

=> Can survive combination attacks of collusion + print + scan and detect fingerprint from hard copies

Joint Coding-Embedding via Anti-Collusion Codes

User-1 **(-1,-1, -1, -1, 1, 1, 1, 1, …, 1)**

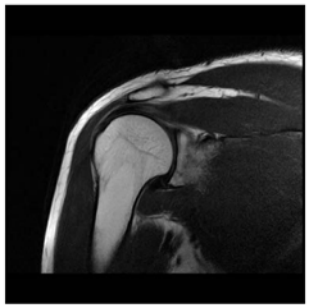

(-1, 1, 1, 1, 1, 1, …, -1, 1, 1, 1) *User-4*

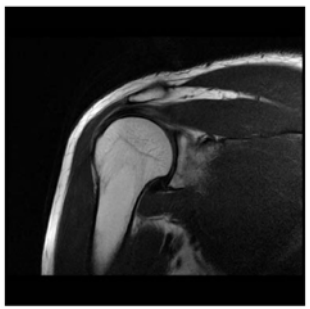

Embed fingerprint via HVS-based spread spectrum embedding

Collude by Averaging

Extracted fingerprint code **(-1, 0, 0, 0, 1, …, 0, 0, 0, 1, 1, 1)**

Uniquely Identify User 1 & 4

16-bit ACC for detecting up to 3 colluders out of 20

Applications: From Government to Hollywood

Insert special signals to identify recipients

- Deter leak of proprietary documents
- Consider imperceptibility, robustness, traceability

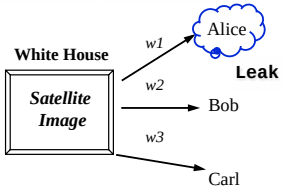

Rights management by copyright industry
($500+Billion ~ 5% U.S. GDP)

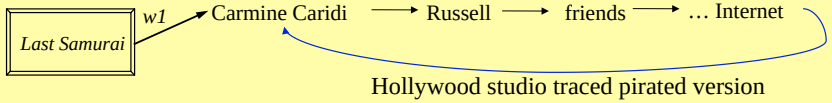

- Successfully catch media pirates => Wide adoption now
  - Alleged Movie Pirate Arrested by FBI – Oscar screening copies (Jan. 2004)
  - Track down illicit sharing of digital TV access – Daqing, China (Nov. 2007): found and convicted perpetrator

Explore More …

- "**Information Forensics: An Overview of the First Decade**," by M. Stamm, M. Wu, K.J.R. Liu, invited article for inaugural issue, *IEEE Access*, May 2013.
  - Joint Coding and Embedding; fundamental tradeoffs
  - Large scale video fingerprinting; special media type such as maps
  - Group based fingerprinting; behavior forensics
- Recent directions: probabilistic fingerprint/tracing code
- Industry R&D: by movie industry (e.g. Technicolor R&D, MovieLab) and printing/marketing (e.g. Digimarc Inc.)

Role Play as Sherlock Holmes in our Digital Era

1. Leak: A proprietary image sent to 10 people got leaked out

   ➔ Who leaked the info?
2. Source: Picture of a heavily guarded xPhone7 prototype showed up on web

   ➔ Is it a real photo or a graphic rendition? Who in the company took it using his/her camera?

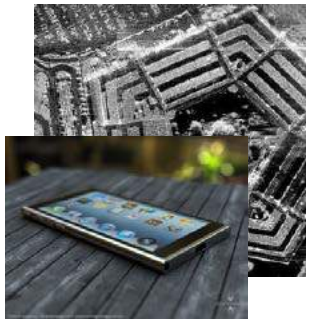

Metadata in header can be forged ☹

Exploit Intrinsic Fingerprints via Component Forensics

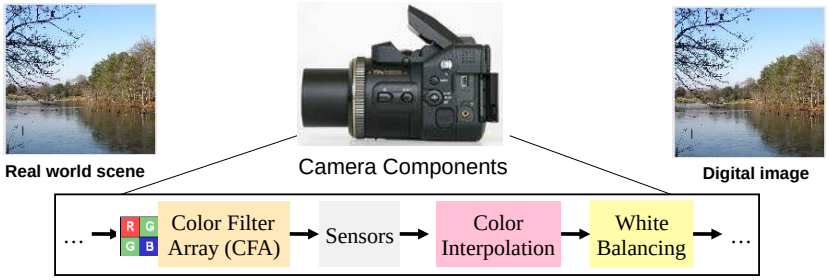

- Break down the info. processing chain into individual components
- Identify algorithms and parameters employed in major components of a digital device or processing system
- Concept extensible to general info proc. chain beyond multimedia

Intrinsic Traces Representing Group Properties

- "Digital / software" components of device or processing system

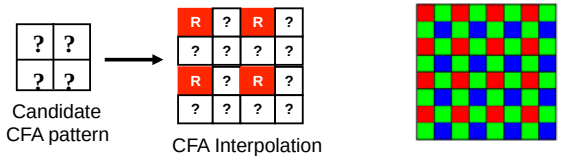

- Ensemble properties of analog components:

  e.g. statistical noise profile of sensors

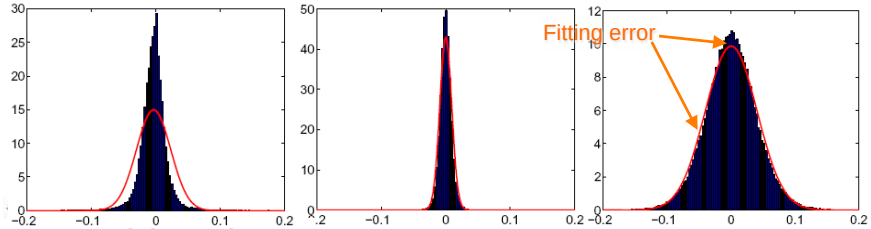

Photography vs Computer Graphics?

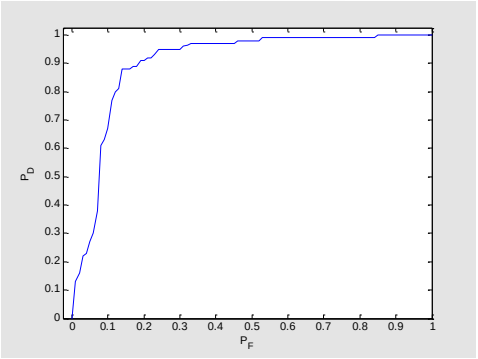

P_F – prob. of false alarm
(% of photographic images misclassified)

P_D – prob. of correct decision
(% of computer generated images classified correctly)

*Photorealistic image source – Columbia univ. database

Note: PRNU pattern amplified by a factor of 42,500

Intrinsic Traces Link to Individual Devices

- Individual variation from "analog" part of sensors
  - "Unreproducible" properties due to manufacturing variability
- Challenges to overcome
  - Picture content variability, post-processing, anti-forensics, etc.

-123	-106	-34	16
63	-11	-4	39
60	59	16	-15
-83	64	7	-20

43	-7	12	-36
-9	-10	65	51
71	-9	40	101
53	-13	54	74

-14	-27	-5	-42
65	37	56	63
95	68	41	-11
61	40	2	-51

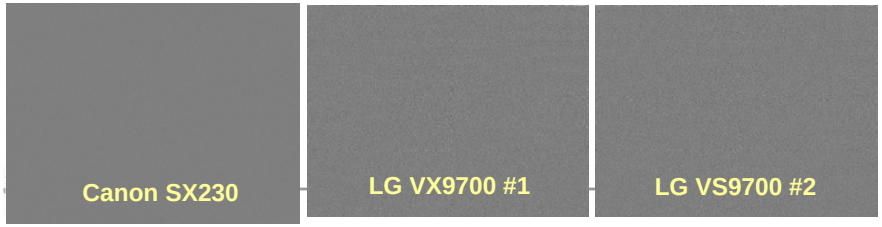

Matching Sensor Noise via Correlation Metrics

- Measure similarity by normalized cross-correlation (NCC)
  - NCC gives a sharp peak at the right alignment of right match
  - May include weights from measurement reliabilities
- Robust peak measurement
  - Normalize with "ambient" correlation values (Fridrich et al.)

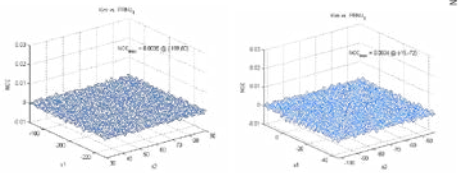

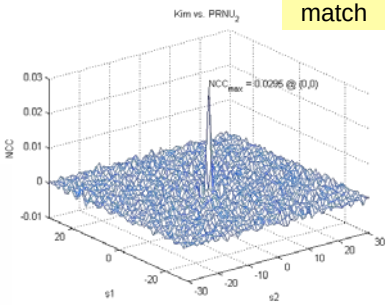

match

mis-match

Tampering Detection

Explore intrinsic fingerprints left by various processing modules

- *To infer the algorithms and parameters employed in various components of the digital device and processing systems*
- *New traces or vanished old traces suggests potential post-camera operations*

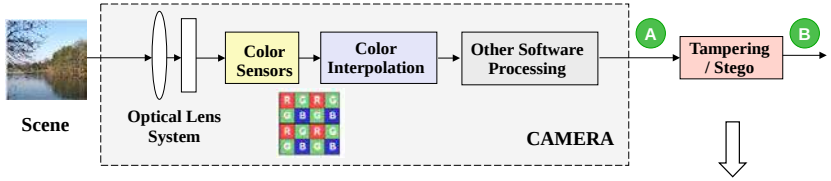

From direct camera output

After post-camera filtering

Black: Sony P72; White: Canon Powershot S410
Grey: Classified as other cameras with low confidence

Explore More …

- Early years' research:

  Special Issue on Digital Forensics
  by IEEE Signal Processing Magazine
  (March 2009)
  Edited by E. Delp, N. Memon and M. Wu

  and a concurrent special issue in
  IEEE Security & Privacy Magazine

- "**Information Forensics: An Overview of the First Decade**," by M. Stamm, M. Wu, K.J.R. Liu, *IEEE Access* (invited), May 2013.

- New DARPA "MediFor" program

Role Play as Sherlock Holmes in our Digital Era

1. Leak: A proprietary image sent to 10 people got leaked out

   ➔ Who leaked the info?

2. Source: Picture of a heavily guarded xPhone6 prototype showed up on web

   ➔ Is it a real photo or a graphic rendition? Who in the company took it using his/her camera?

3. When/Where: an audio clip with incriminating words showed up

   ➔ were its content and recording time true as claimed?

Role Play as Sherlock Holmes in our Digital Era

1. Leak: A proprietary image sent to 10 people got leaked out

   ➔ Who leaked the info?

2. Source: Picture of a heavily guarded xPhone7 prototype showed up on web

   ➔ Is it a real photo or a graphic rendition? Who in the company took it using his/her camera?

Ubiquitous Forensic Fingerprints from Power Grid

- Electric Network Frequency (ENF): 50 or 60 Hz nominal
  - Change slightly due to demand-supply
  - Main trends consistent in same grid

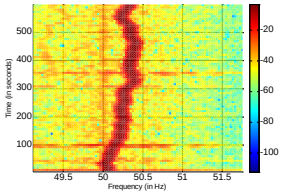

Power ENF signal

Ubiquitous Forensic Fingerprints from Power Grid

- Electric Network Frequency (ENF): 50 or 60 Hz nominal
  - Change slightly due to demand-supply
  - Main trends consistent in same grid
- ENF can bee "seen" or "heard" in sensor recordings
  - Power grid influences electronic sensing (E/M interference, vibration etc)
  - Help determine recording time/location, detect tampering, etc.

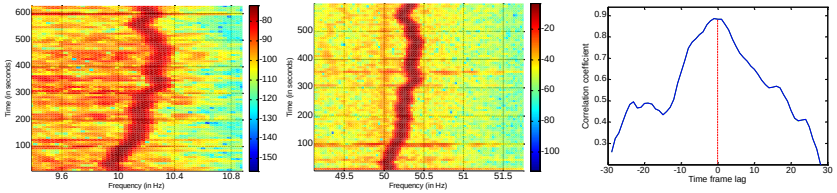

Video ENF signal | Power ENF signal | Normalized correlation

ENF matching result demonstrating similar variations in the ENF signal extracted from video and from power signal recorded in India

Tampering Detection

- Adding a clip into original video leads to discontinuity in ENF
  - Clip insertion can also be detected by comparing the video ENF signal with the power ENF at corresponding time

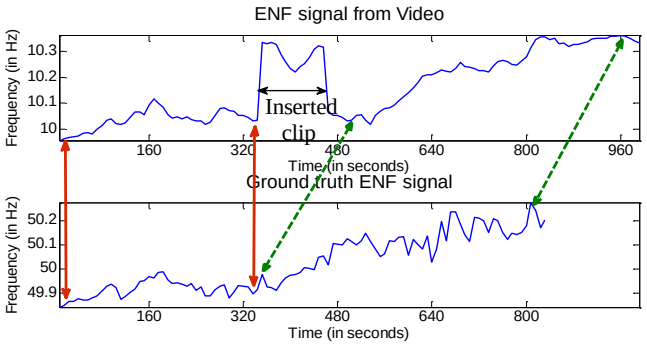

"Forensic Binding" of Audio and Visual Tracks

- High correlation of ENFs in audio & video captured at same time
  => can extend to synch multiple media streams

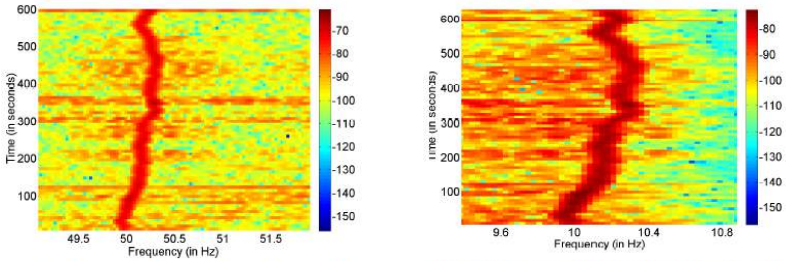

(a) ENF signal from the audio track (b) ENF signal from the video track

- Anti-Forensic Study: possible to remove narrow-band ENF; but much harder to tamper/transplant a valid ENF w/o being caught

From Time Stamps to Location

- Match with ENF references over times + grids
  - Verify or exhaustively search for the matching location on grid level

What if no concurrent references available?
- Explore overall characteristics of ENF in a grid
- Also reduce computation of exhaustive search

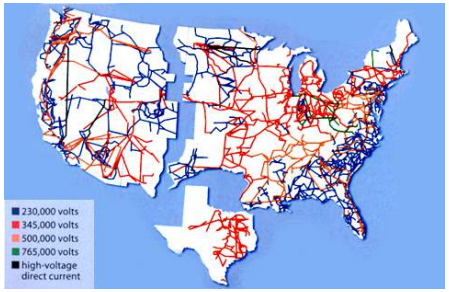

Explore Machine Learning to Infer Location

- Inter-Grid location-of-recording estimation from sensing signals containing ENF traces
  - Identified useful features for average 94% accuracy on audio

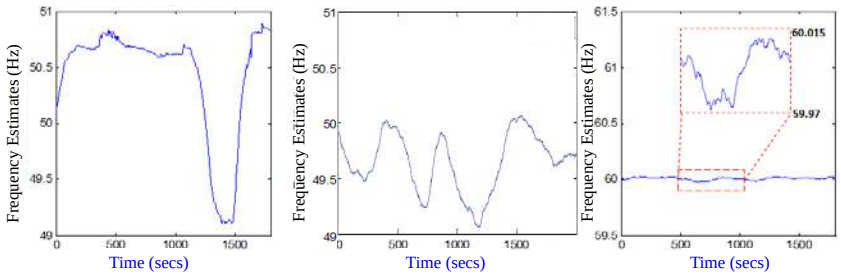

LEBANON | INDIA | Eastern US

From Time Stamps to Location

- Match with ENF references over times + grids
  - Verify or exhaustively search for the matching location on grid level

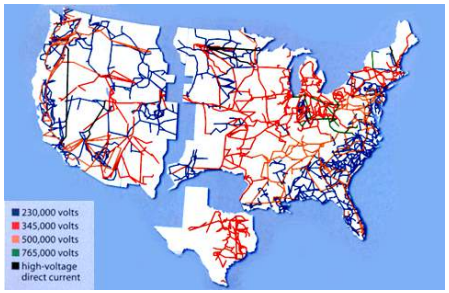

What if no concurrent references available?
- Explore overall characteristics of ENF in a grid
- Also reduce computation of exhaustive search

Can we determine where within a grid?
- E.g. DC or NYC in US East?
- Look at subtle traces in ENF
- Relate ENF correlation with distance

Can ENF Pinpoint to Locations Within a Grid?

- Main trend of ENF is known to be the same in a grid
- "Microscopic" traces
  - Aggregated effect of local events and propagations from elsewhere
- Our multi-location studies in U.S. east and west grids
  - Relate pairwise ENF correlations between query and anchor points with geographic and wireline distances

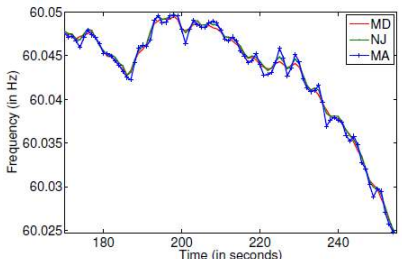

(a) ENF signals from different locations of US Eastern grid

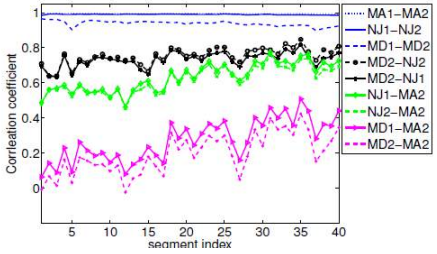

(b) Correlation between ENF signals after high-pass filtering

ENF in Historical Recordings

- Two ENFs may appear in digitized tape recordings
  1) original ENF; and
  2) ENF at time of digitization

  => Provide digital preservation guidelines to better utilize invisible traces
- Distortions and artifacts
  - Drifting; low SNR; etc.
- Ongoing: create a historical ENF database
  - Timestamp recordings of historic importance

Digitized Kennedy White House recording (1962 Cuban missile crisis)

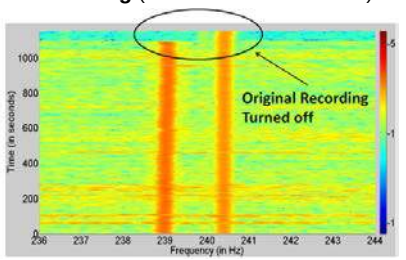

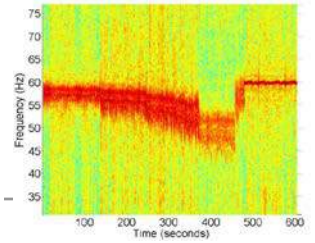

NASA audio from Apollo 11 (1969 1st moon-landing)

Speed Restoration: ENF as Intrinsic Freq. Reference

NASA Apollo 11 Mission recording

Before Restoration | After Restoration

Spectrogram

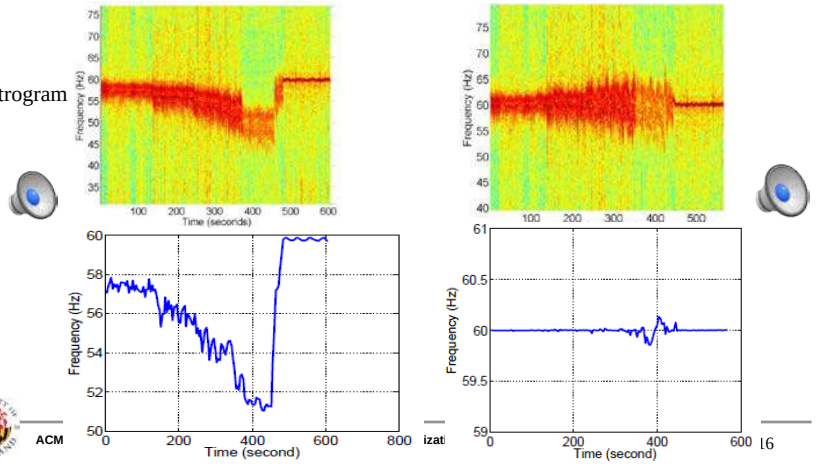

Immersive Media: Synch Streams via ENF in Audio

Demo-1:
Videos in Gym

(different viewing angles)

Demo-2:
Videos at different locations of Lab Building

Video after synchronization

2 synched stopwatches (as ground truth)

ENF Research & References

Estimate ENF Signal
(instantaneous freq.):

Robust, high resolution;

Exploit harmonics

SPL'13. APSIPA '12,
ACM MM'11, TIFS'13

Visual Modality:

Handle aliasing
– exploit rolling shutters;

Handle motion

TIFS'13, ICIP'14,
ACM MM'14 Immersive Media

Modeling & Analysis:

Statistically modeling of ENF signals;

Anti-Forensics.

WIFS'12, CCS'12, twoTIFS'13

Novel Applications:

Location & integrity;

Stream alignment;

Digital humanity (historic audio)

ICASSP'12-13, WIFS'13 / TIFS'15, iConf'14, ACM MM'14 Immersive

"Information Forensics: An Overview of the First Decade,"

by M. Stamm, M. Wu, K.J.R. Liu,
IEEE Access, invited article
for inaugural issue, May 2013.

MediFor: Newly Launched DARPA Program on media forensics aiming at restoring "Seeing is Believing"

See also Special Issue on Digital Forensics
by IEEE Signal Processing Magazine (March 2009)
Edited by E. Delp, N. Memon and M. Wu;
and a concurrent special issue in IEEE Security & Privacy Magazine

THANKS to many who paved ways …

Include joint work with collaborators:

K.J. Ray Liu, Wade Trappe, Jane Wang, Hong Zhao; Kari Klaus, Douglas Oard.

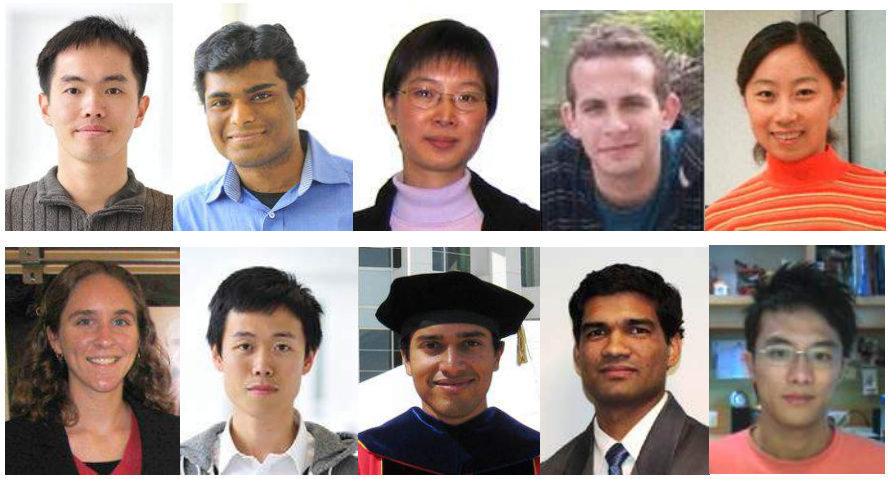

Include joint work with graduate & undergrad students

Wei-Hong Chuang, Ravi Garg, Hongmei Gou, Adi Hajj-Ahmad, Shan He, Christine McKay, Hui Su, Ashwin Swaminathan, Avinash Varna, Chau-Wai Wong.

Min Wu: Info Forensics & Media Security

- Provide assurance for proper use of content
  - Answer who has done what, when and how.
- Cross-disciplinary and balancing theory & practice
  - Analytic modeling for fundamental understanding
  - Design effective and efficient algorithms with synergy from signal proc., comm, machine learning, crypto …

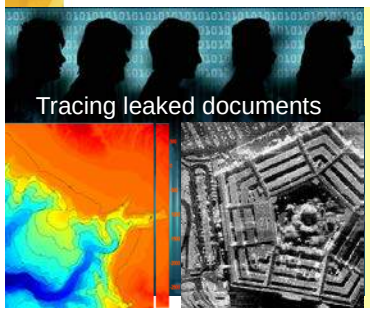

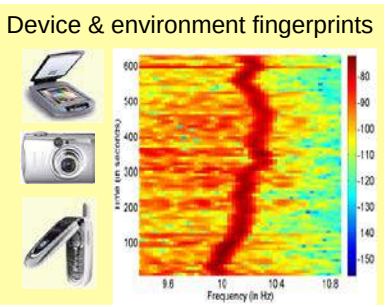

Learn more from this URL