Gwenaël Doërr

Technicolor R&D France
Security & Content Protection Labs
gwenael.doerr@technicolor.com

Checking Up the Health of Multimedia Security

technicolor
Security Laboratories

Agenda

Content Protection Ecosystem

Review of the Anti-piracy Arsenal

Signal Processing in the Presence of an Adversary

Research Outlook

Questions and Answers

Information Forensics and Security

technicolor
Security Laboratories

The Challenging Transition to Digital

Key specificities of digital content

- Clones rather than copies i.e. no more generational degradation
- Assets can be tangible or intangible
- Ease of dissemination i.e. the world is at your doorstep

Apparition of a bestiary of pirates (Courtesy: Irdeto)

On the cost of piracy…

CNBC's Crime Inc #10: Hollywood Robbery (August 2012)

technicolor Security Laboratories

Threat Analysis

Filming

Insider leak (dailies theft)

Post-processing (color, CGI, edits, etc.)

Insider leak

Movie

Insider leak

Distributor
TV, cable, satellite, IP

Insider leak

Hospitals, prisons
Broadcast, VOD

Hotels

Home, pubs, clubs
Broadcast, VOD

Stream redirection
HDMI strip
Screencasting
Camcording

Filmprint theft

Theater release

Camcording

Insider leak

Airplane, cruise

N/A (Screencasting)

Insider leak at replication plant

Hardcopy release
(DVD and Blu-ray)
(rental and retail)

DVD/Blu-ray rip

Promotional copies
(Oscars, journalists, etc)

DVD/Blu-ray rip

technicolor
Security Laboratories

In-Theater Camcording over the Years

Number of pirate samples over time
Source: Raw data from MPAA piracy report (January 2012)

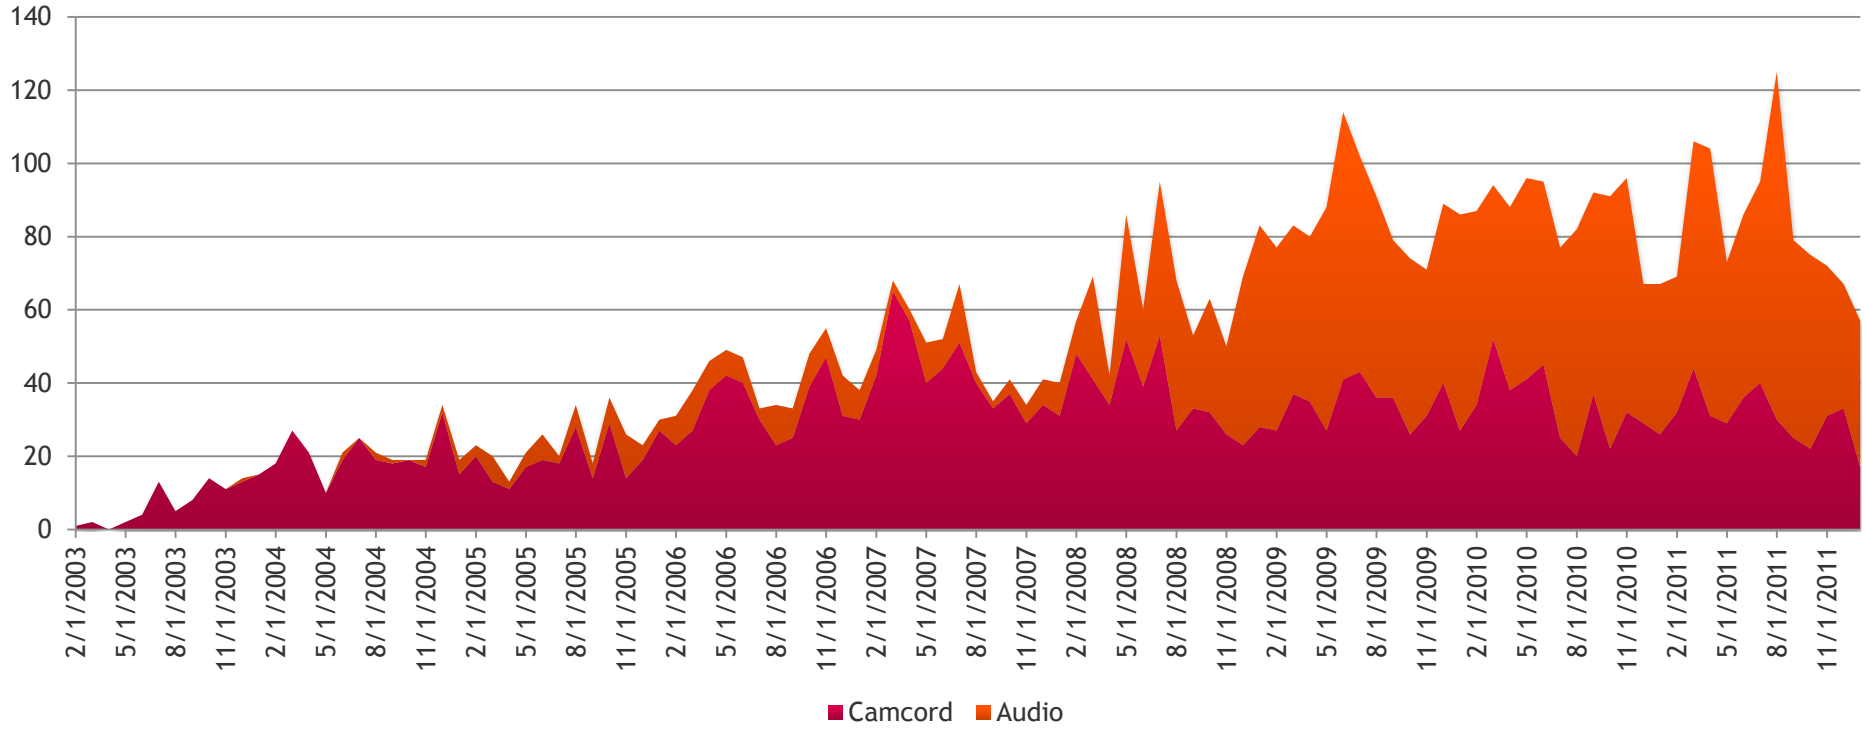

technicolor
Security Laboratories

Time-to-Black-Market

Number of days elapsed between US theatrical release and piracy detection
Source: Raw data from MPAA piracy report (January 2012)

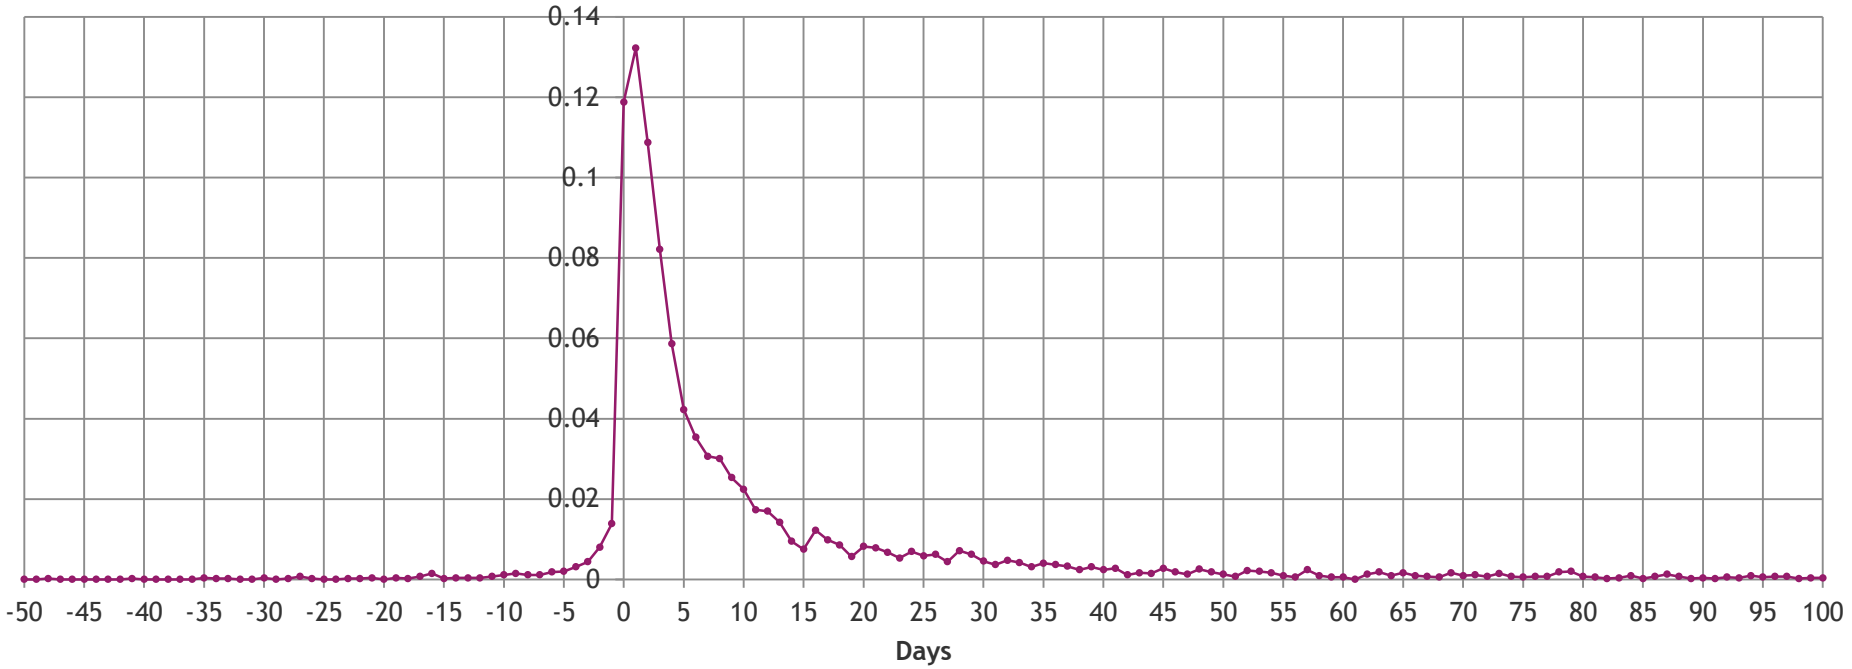

technicolor
Security Laboratories

Anti-Piracy Arsenal

Regulate

- WIPO 1996 (DMCA, EUCD, Hadopi, etc.)
- SOPA/PIPA

Inform / Educate

- FA©T anti-piracy information campaigns
- Hard-to-counterfeit security features
  - Intaglio, color-shifting inks, holograms, CDIs

Prevent

- Content encryption aka. CAS and DRM
- Anti rip

Interfere / Jam

- Anti-recording e.g. Macrovision
- Anti-camcording

Monitor / Scout

- Data loss prevention systems
- Content fingerprinting

Trace

- Digital watermarking
- Passive forensics

The Forensics Landscape

Outsourced
Internet Crimes Group, Deluxe (WEBWatch), Irdeto (Bay TSP), private investigators, blueSpike

Operated by MPAA
Manual process mostly relying on textual metadata collated in spreadsheets/databases e.g. the pirate release group name

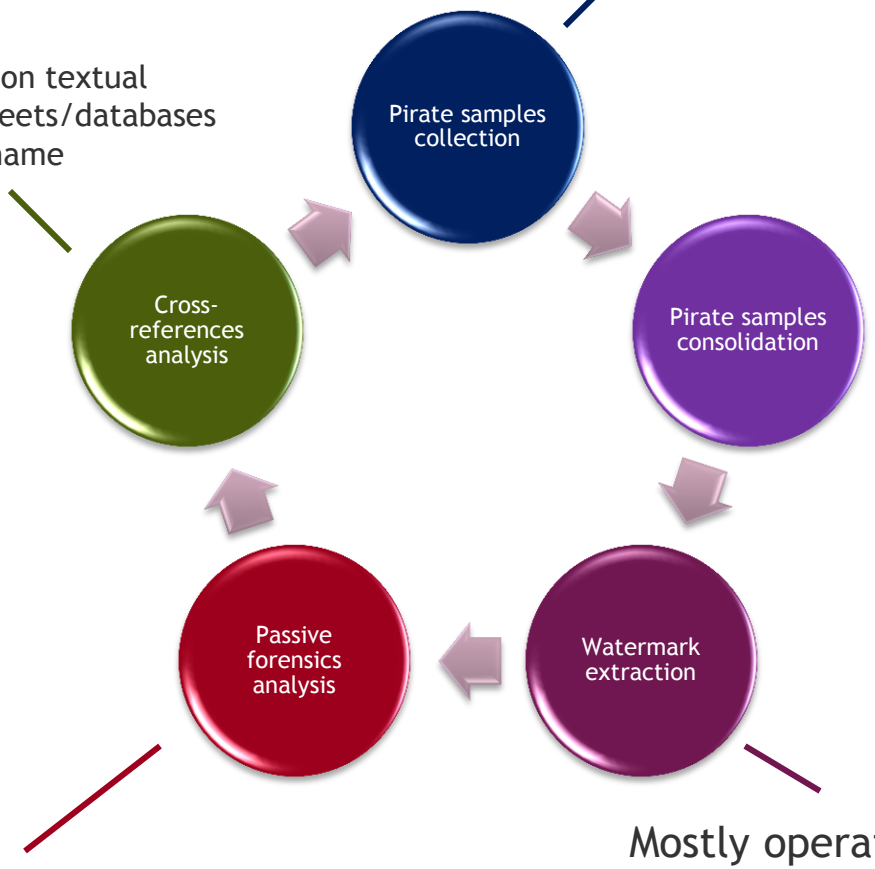

Outsourced
Manual Process?

Mostly operated by Deluxe/Technicolor
with detectors bought/licensed from WM vendors e.g. Civolution, Verimatrix, Irdeto, Verance, (Technicolor) or through web portals

Mostly operated by Deluxe/Technicolor
Manual process typically involving tedious audio-visual inspection

technicolor
Security Laboratories

Multimedia Encryption

Bulk encryption of the content essence

- Symmetric/asymmetric encryption
- Key distribution schemes for broadcast
- Usage rights transported jointly/separately

Selective encryption

- Preserve the battery of low-powered devices
- Provide preview to trigger purchase

Content Fingerprinting

Robust DNA-like compact representation

- Two contents should ‘hash’ to the same fingerprint as long as they are perceptually similar

Baseline framework

- Robust representation: filter banks, transforms, features extraction
- Quantization: ad-hoc, K-means, etc
- Binarization

Properties: granularity, robustness, discriminability, scalability

Applications

- Content identification: automated rights clearance, data loss prevention, broadcast monitoring
- Content realignment

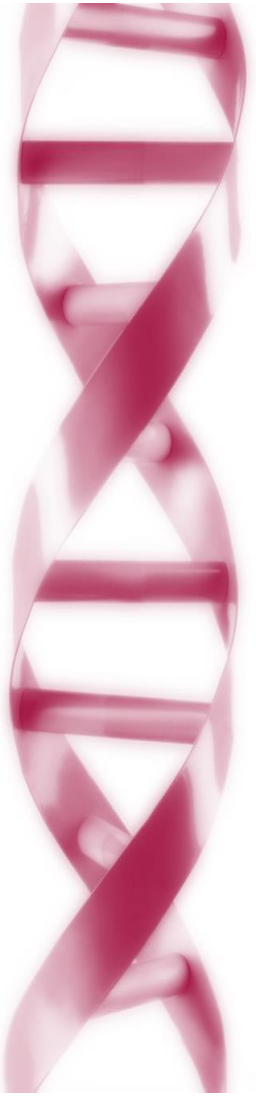

Digital Watermarking

Digital watermarking is a technique which imperceptibly alter digital content to hide a secret message in a robust manner. It is in some sense similar to invisible ink and paper watermarks.

Baseline framework

- Content adaptation: transform, perceptual model
- Communications layer: watermark modulation, resynchronization

Properties: capacity, fidelity, robustness, complexity, security

Applications

- Content protection: traitor tracing, copy control, broadcast monitoring
- Content enrichment a.k.a. second screen

Passive Forensics

Isolation of tell-tale statistical discrepancies

- Sensor forensics
- Processing pipeline forensics

Applications

- Content authentication
  - Reality check after Photoshopping
- Piracy path characterization
  - Compensation of piracy artifacts
  - Adjustment of the tracing piracy
  - Metadata for cross-referencing

Adversary-aware Signal Processing

Potential for money and/or strict laws ⇒ opponents and attacks

- Reverse multimedia scrambling techniques
- Wash out digital watermarks
- Reconstruct content from fingerprint
- Clean-up forensic statistical digital traces
- Etc

Objective of the adversary: learn or infer hidden parameters of the system to modify its expected behavior

- Leverage on a priori knowledge about content/secret statistics
- Sensitivity analysis to learn decision boundaries ⇒ switch decisions

Strong links to game theory

- Trade-off robustness ↔ security

Oracle Attack

Step over the secret boundary of a binary decision

- Watermarked or not, authentic or not, key point or not

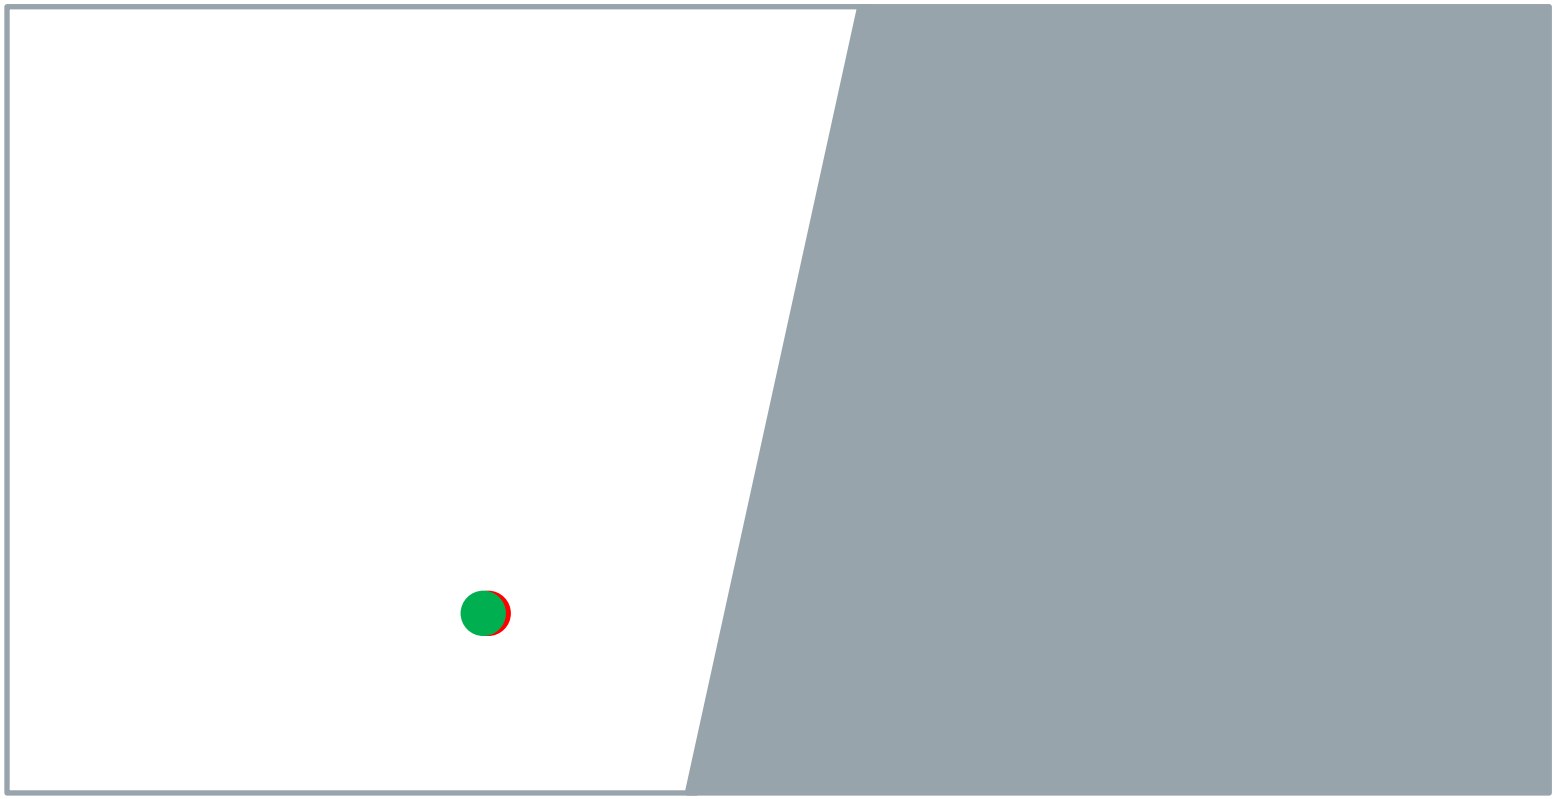

Secret Estimation from Multiple Observations

Setup: access to several contents watermarked with the same key

⇒ Look for peculiar persistent statistical properties

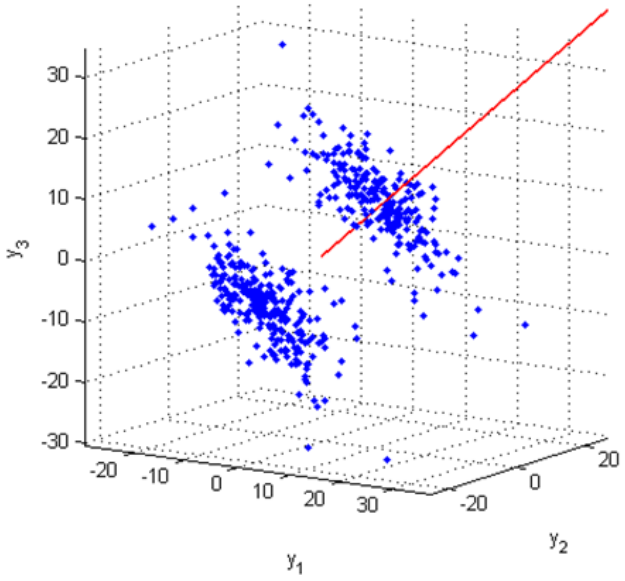

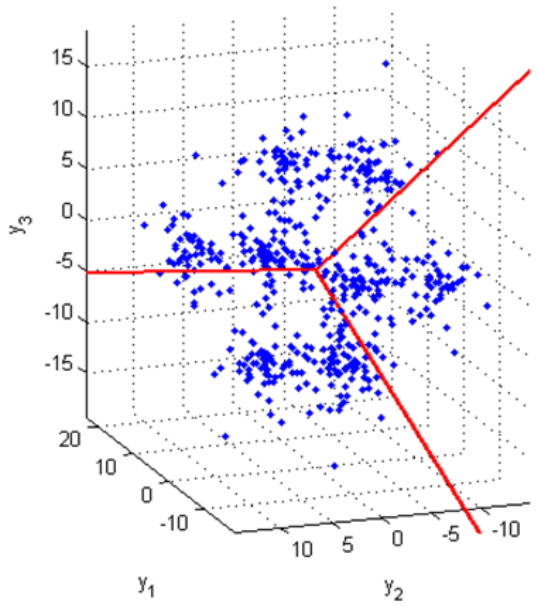

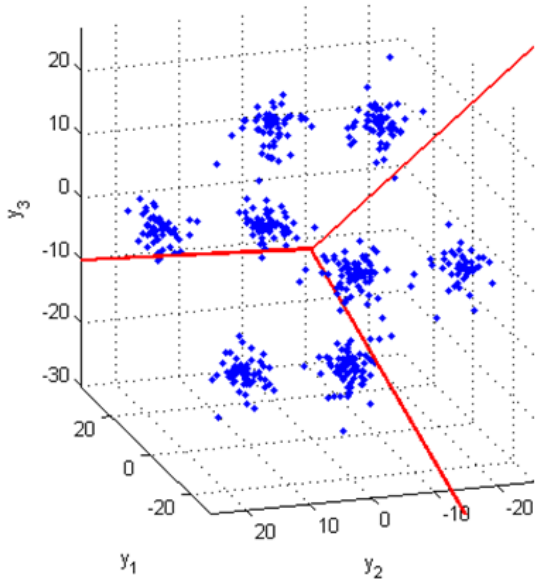

Exploit this knowledge to attack the system

technicolor
Security Laboratories

Reconstruction from SIFT

Original image

Reconstruction from SIFT description

+ inpainting

P. Weinzaepfel, Hervé Jégou, and Patrick Pérez, "*Reconstructing an Image from its Local Descriptors*", CVPR 2011

Defense Mechanims

Obfuscation techniques

- Security by obscurity
- Key-dependent parameterization of the system
  - Random permutations, projections, quantization

Cryptographic techniques

- Homomorphic encryption
- Zero-knowledge protocols
- Etc.

Obfuscation Techniques

1. Random tiling of the image
2. Compute some statistics for each tile
   e.g. mean, variance, etc
3. Randomized rounding

R. Venkatesan, S.-M. Koon, M. H. Jakubowski, and P. Moulin, "*Robust Image Hashing*", ICIP 2000

1. Generate low-pass pseudo-random patterns
2. Project the content onto those patterns
3. Take the sign of the correlation value
4. Generate the binary digest with a heuristic design

J. Fridrich and M. Goljan, "*Robust Hash Functions for Digital Watermarking*", ICIT 2000

CBIR: Randomizing the Quantizer

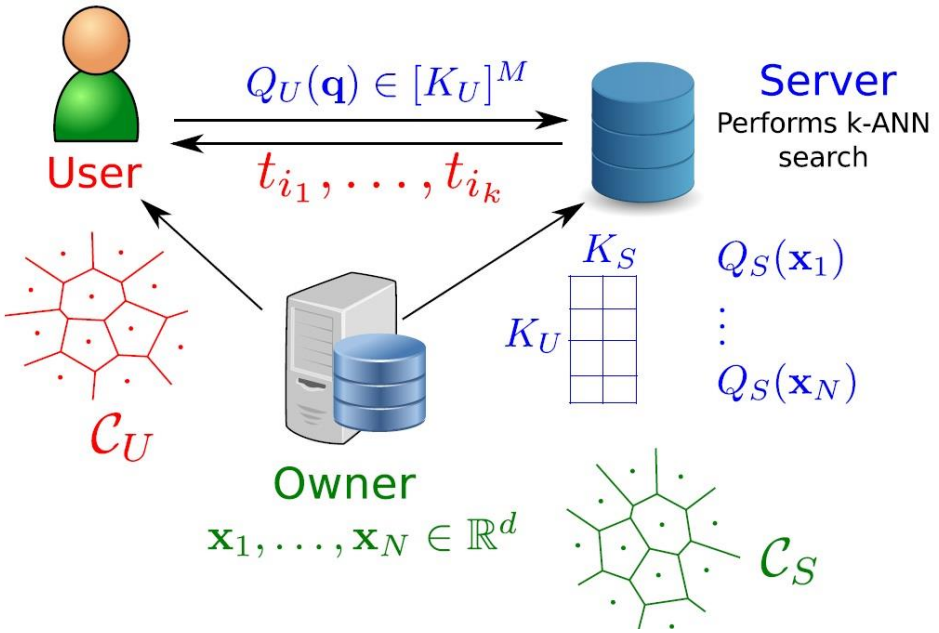

Baseline idea: randomize the quantizer & use different quantizer for Server and User

Randomized quantizers

- Random training subset
- Random initialization vector
- Stop before convergence

Curious but honest Server

- ~~Reconstruct $\mathbf{x}_i$ from h($\mathbf{x}_i$)~~
- Cluster the database vectors
- ~~Reconstruct $\mathbf{q}$ from h($\mathbf{q}$)~~
- *Detect similar queries*

B. Mathon, T. Furon, L. Amsaleg, and J. Bringer, "*Secure and Efficient Approximate Nearest Neighbors Search*", ACM IHMMSec 2013

technicolor Security Laboratories

The Issue of Security Assessment

How much security is provided by heuristic obfuscation techniques?

- Different keys ⇒ different obfuscated objects?

Several metrics based on information theory

☺ Mutual information, differential entropy, etc

☹ No security proof

What does it mean to be “more secure”?

Conventional Cryptography

Baseline principle: layered architecture to separate the signal processing layer from the cryptographic layer

Cryptographic hash functions
(typically used for authentication)

- High sensitivity: $a \approx b \Rightarrow h(a) \neq h(b)$
- Non invertibility
- Small collision probability

Visual hash: content fingerprint + hash function

- Inherits robustness from the fingerprint and security from the hash
- Does not really work in practice
  - Content fingerprinting is not strictly robust (even with ECC decoder hack)

Homomorphic Encryption

$$E_K(A+B) = E_K(A) \times E_K(B)$$

Linear operations directly in the encrypted domain

- Signal processing in the encrypted domain
- Privacy enhancement technologies

☺ Provides all the security features that you could dream of

☺ Recent leap forward with Gentry's fully homomorphic scheme

😐 Many operations not supported e.g. thresholding, trigonometry, …

☹ Overhead: big and slow!

R.L. Lagendijk, Z. Erkin, and M. Barni, "*Encrypted Signal Processing for Privacy Protection*", IEEE SPM, 2013

Traitor Tracing Codes

Rationale: embed an identifier unique to each recipient to pinpoint the source of a leak

Threat: several users colluding to produce a pirate copy

Risk: framing innocent users

Marking assumption: colluders can only modify bits that differ in their copies

Traitor tracing codes

- Cryptographers: provably secure, decoding, long
- Statisticians: error possible, exhaustive search

Research Outlook

Bad news: most low-hanging fruits have already been picked up

Multimedia encryption

- Format-preserving encryption for collaborative creation
- Impact on the content creation workflow

Digital watermarking

- Dealing with correlated samples
- Dealing with content-dependent transforms
- Perceptual models for stereo, HDR, UWG, HOA, …
- Real multi-dimensional watermark modulation
- Explaining the discrepancy between theory and practice
- Registration mechanisms

Content fingerprinting

- Registration-geared fingerprints
- Near-duplicates management

Passive forensics

- Piracy path modeling
- Piracy path identification
- Piracy path characterization

Concluding Remarks

Common pitfalls

- False sense of security by invoking crypto argument
- Inclination to fall in a cats and mouse loop
- Find a solution to a non-existing problem
- Overlooking the impact of security on performances
- Search for perfect security

Challenging marketing strategy

- Return on investment vs. non loss
- History of overselling multimedia security

Small research community at the intersection of multiple disciplines

Questions

technicolor

Security Laboratories