

Resilient Rights Protection for Sensor Streams *

Radu Sion, Mikhail Atallah, Sunil Prabhakar
Center for Education and Research in Information Assurance
Indiana Center for Database Systems
Computer Sciences, Purdue University
{sion, mja, sunil}@cs.purdue.edu

Abstract

Today's world of increasingly dynamic computing environments naturally results in more and more data being available as fast streams. Applications such as stock market analysis, environmental sensing, web clicks and intrusion detection are just a few of the examples where valuable data is streamed. Often, streaming information is offered on the basis of a non-exclusive, single-use customer license. One major concern, especially given the digital nature of the valuable stream, is the ability to easily record and potentially "re-play" parts of it in the future. If there is value associated with such future re-plays, it could constitute enough incentive for a malicious customer (Mallory) to duplicate segments of such recorded data, subsequently re-selling them for profit. Being able to protect against such infringements becomes a necessity.

In this paper we introduce the issue of rights protection for streaming data through watermarking. This is a novel problem with many associated challenges including: operating in a finite window, single-pass, (possibly) high-speed streaming model, surviving natural domain specific transforms and attacks (e.g. extreme sparse sampling and summarizations), while at the same time keeping data

alterations within allowable bounds. We propose a solution and analyze its resilience to various types of attacks as well as some of the important expected domain-specific transforms, such as sampling and summarization. We implement a proof of concept software (wms.*) and perform experiments on real sensor data from the NASA Infrared Telescope Facility at the University of Hawaii, to assess encoding resilience levels in practice. Our solution proves to be well suited for this new domain. For example, we can recover an over 97% confidence watermark from a highly down-sampled (e.g. less than 8%) stream or survive stream summarization (e.g. 20%) and random alteration attacks with very high confidence levels, often above 99%.

1 Introduction

Protecting rights over outsourced digital content becomes essential when considering areas where the data is sensitive and valuable. One example is the outsourcing of data for data mining. In this scenario data is produced/collected by a data collector and then sold to parties specialized in mining it. Different rights protection avenues are available, each with its advantages and drawbacks. Enforcement by legal means is usually ineffective, unless augmented by a digital counterpart such as Information Hiding. Digital Watermarking deploys Information Hiding as a method of Rights Protection to conceal an indelible "rights witness" (watermark) within the digital Work to be protected. The soundness of such a method relies on the assumption that altering the Work in the process of hiding the mark does not destroy the value of the Work, and that it is difficult for a malicious adversary ("Mallory") to remove or alter the mark beyond detection without destroying the value of the Work. The ability to resist attacks from such an adversary (mostly aiming at removing the embedded watermark) is one of the major concerns in the design of a sound solution.

A considerable amount of effort has been invested in the problem of watermarking multimedia data (im-

Portions of this work were supported by Grants EIA-9903545, IIS-0325345, IIS-0219560, IIS-0312357, IIS-9985019 and IIS-0242421 from the National Science Foundation, Contract N00014-02-1-0364 from the Office of Naval Research, by sponsors of the Purdue Center for Education and Research in Information Assurance and Security, and by Purdue Discovery Park's e-enterprise Center.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

ages, video and audio). More recently, the focus of watermarking for digital rights protection is shifting toward other data domains such as natural language text [2], software, algorithms [7] [15] and relational data [11] [18] [19]. Since these data domains often have very well defined restrictive semantics (as compared to those of images, video, or music) and may be designed for machine ingestion, the identification of the available “bandwidth” for watermarking is as important a challenge as the algorithms for inserting the watermarks themselves.

In this paper we introduce and study the problem of watermarking sensor streams data, which to the best of our knowledge, has not been addressed. Streaming data sources represent an important class of emerging applications [3] [4]. These applications produce a virtually endless stream of data that is too large to be stored directly. Examples include output from environmental sensors such as temperature, pressure, brightness readings, stock prices etc. Recent efforts in the broader area of streaming data deal with the database challenges of its management [5] [9] [10] [13].

Existing work on itemized data types [11] [18] [19] relies upon the availability of the entire dataset during the watermarking process. While this is generally a reasonable assumption, it does not hold true for the case of streaming data [3]; since the streamed data is typically available as soon as it is generated, it is desirable that the watermarking process be applied immediately on subsets of the data. Additionally, the attack and transformation models in existing research does not apply here. For example a process of summarization would defeat any of the above schemes. Yet another difference from previous research is the lack of a “primary key” reference data set, an essential, required, part in both [11] and [19]. Due to these differences, earlier work on watermarking relational data sets is not applicable to streams.

But why is watermarking streaming data important? Couldn’t we simply watermark the data once it is stored? This surely would work and enable rights protection for the stored result. But it would not deter a malicious customer (Mallory), with direct stream access, to duplicate segments of the stream and re-sell them or simply re-stream the data for profit. The main rights protection scenario here (see Figure 1) is to prevent exactly such leaks from a licensed customer.




Figure 1: Sensor Streams Watermarking Scenario.

Our contributions include (i) the proposal and definition of the problem of watermarking sensor streams,

(ii) the discovery and analysis of new watermark embedding channels for such data, (iii) the design of novel associated encoding algorithms, (iv) a proof of concept implementation of the algorithms and (v) their experimental evaluation. The algorithms introduced here prove to be resilient to important domain-specific classes of attacks, including stream re-sampling, summarization (replacing a stream portion by its average value) and random changes. For example, sampling the data stream down to less than 8% still yields a court-time confidence of watermark embedding of over 97%. Summarization (e.g. 20%) and random data alterations are also survived very well, often with a false-positive detection probability of under 1%.

The paper is structured as follows. Section 2 outlines the major challenges in this new domain. It proposes an appropriate data and transform model, discusses associated attacks and overviews related work. In Section 3 an initial solution is provided. Further resilience-enhancing improvements and attack handling capabilities are gradually introduced in Section 4. Section 5 analyzes the ability to convince in court to survive attacks and natural domain transformations. Section 6 presents **wms.***, a proof-of-concept java implementation of our solution; our experimental setup and results are introduced. Section 7 concludes.

2 Challenges

2.1 The Adversary

As outlined above, the nature of most “fast” time-series data applications imposes a set of strict requirements on any on-the-fly data processing method, such as watermarking. For one, it has to be able to keep up with the incoming data rate and, the fact that only a finite window of memory (e.g. of size ϖ , see below) is available for processing makes certain history-dependent computations difficult or simply impossible. At the same time, metrics of quality can only be handled within this space; any preservation constraints can be formulated only in terms of the current available data window. Including any history information will come at the expense of being unable to store as much new incoming data.

Moreover, the effectiveness of any rights protection method is directly related to its ability to deal with normal domain specific transformations as well as malicious attacks. There are several transforms relevant in a streaming scenario, including the following: (A1) summarization, (A2) sampling, (A3) segmentation (we would like to be able to recover a watermark from a finite segment of data drawn from the stream), (A4) linear changes ¹ (there might be value in actual *data trends*, that Mallory ² could still exploit, by scaling the initial values), (A5) addition of stream values and (A6) random alterations.

¹Taken care of by the initial normalization step.

²The traditional name of *the* maliciously acting party.

While we discuss most of these and other attacks in the next sections, let us note here that with respect to (A5), Mallory is bound to add only a limited amount of data (in order to preserve the value in the original stream) and these new values are to be drawn from a similar data distribution, lest they become easy to identify in the detection process as not conforming to the known original distribution. Also (A6) can be naturally modeled by a combination of (A2) and (A5).

2.2 Model

For the purpose of simplicity let us define a simple data stream as an (almost) infinite timed sequence of $(x[t])$ values “produced” by a set of data sources of a particular type (e.g. temperature sensors, stock market data). $x[t]$ is a notation for the value yielded by our source(s) at time t . Unless specified otherwise, let us denote a stream as $(x[], \varsigma)$ where ς is the number of incoming data values per time unit (*data rate*)³.

Note: While a time-stamp t can be assigned naturally to each and every data value when produced by a data source, it often becomes irrelevant after such domain-specific transformations as sampling and summarization which destroy the exact association between the value $x[t]$ and the time it was initially generated, t . Thus, the notation $x[t]$ is merely used to distinguish separate values in the stream and is not intended for suggesting the preservation of the time-stamp-value in the resulting stream (which is ultimately just a sequence of values).

Any stream processing is necessarily both time and space bound. The time bounds derive from the fact that it has to keep up with incoming data. We are going to model the space bound by the concept of a window of size ϖ . At each given point in time, no more than ϖ of the stream $(x[t])$ values (or equivalent amounts of arbitrary data) can be stored locally, at the processing point. Unless specified otherwise, as more incoming data becomes available, the default behavior of the window model is to “push” older items out (i.e. to be transmitted further, out of the processing facility) and “shift” the entire window (e.g. to the right) to free up space for new entries.

For simplicity, without sacrificing generality, for the remainder of the paper we are going to assume the stream values being normalized in the interval $(-0.5, +0.5)$. This assumption does not need to hold in general but instead just simplifies the task of understanding the algorithms.

For the purpose of the current framework, we define the *uniform random sampling* of degree χ of a stream $(x[], \varsigma)$ as another stream $(x'[], \varsigma')$ with $\varsigma' = \frac{\varsigma}{\chi}$ such that for each sample data item $x'[t]$, there exists a contiguous subset of $(x[])$, $(x[t_1], x[t_2])$ such that $x'[t] \in (x[t_1], x[t_2])$, $\{x'[t-1], x'[t+1]\} \not\subseteq (x[t_1], x[t_2])$, and t is uniformly distributed in (t_1, t_2) . In other

³The proposed solution does not rely on any characteristic of the actual stream data rate. For space and simplicity purposes in this paper we are discussing streams with fixed data rates.

words, it is constructed by randomly choosing one value out of every χ values in the original. A subtle variation of *uniform random sampling* is the case when $x'[t]$ is not randomly chosen but rather always the first element in it’s corresponding χ sized subset (e.g. $t = t_1$). We call this *fixed random sampling* of degree χ .

We define the *summarization* of degree ν of a stream $(x[], \varsigma)$ as another stream $(x'[], \varsigma')$ with $\varsigma' = \frac{\varsigma}{\nu}$ such that for each two adjacent sample data items $x'_1[t], x'_2[t + \nu]$, there exist two contiguous, adjacent, non-overlapping ν -sized subsets of $(x[])$, $(x[t - \nu + 1], x[t - \nu + 2], \dots, x[t])$, $(x[t + 1], x[t + 2], \dots, x[t + \nu])$ such that $x'_1[t] = \frac{\sum_{i \in (1, \nu)} x[t - \nu + i]}{\nu}$ and $x'_2[t + \nu] = \frac{\sum_{i \in (1, \nu)} x[t + i]}{\nu}$. In other words, for a continuous chunk of ν elements from the original stream summarization outputs its average.




Figure 2: (a) A sample stream. If all the extremes are considered to be major, then the resulting label bits for K are shown (for $\varrho = 2$, Section 4.1) (b) δ -Radius characteristic subset of extreme η .

We define an *extreme* η in a stream simply as either a local minimum or local maximum value. We define the extreme’s *characteristic subset of radius δ* , noted $\Xi(\eta, \delta)$ (see Figure 2 (b)), as the subset of stream items forming complete “chunks”, immediately adjacent to η and conforming to the following criteria: item i , with value $v_i \in \Xi(\eta, \delta)$ iff $|\eta - v_i| < \delta$ and all the items “between” i and the extreme η , also belong to $\Xi(\eta, \delta)$.

A *major extreme* of degree χ and radius δ is defined as an extreme η such that at least one item in $\Xi(\eta, \delta)$ can be found in *any* uniform random sampling of degree χ of $(x[])$ (i.e. some items in $\Xi(\eta, \delta)$ “survive” sampling of χ degree). For example, in Figure 2 (a), intuitively, it seems likely that extremes such as F, I and J have a smaller chance of surviving sampling than C, E or G. This is so because of the temporal shape of the stream’s evolution. C, E, G seem to yield characteristic subsets much “fatter” than F, I, J. Intuitively, δ needs to be chosen such that the characteristic subsets are going to be of an average size greater than χ , to handle sampling of degree χ .

To model the “fluctuating” nature of a stream, let $\varepsilon(\chi, \delta)$ be the average number of stream data items encountered/read per *major extreme* (i.e. before encountering a major extreme) of degree χ and radius δ . $\frac{1}{\varepsilon(\chi, \delta)}$ defines the average “frequency of major extremes” in terms of the number of observed data items.

For any numeric value x let $b(x)$ be the number of bits required for its accurate representation and

$msb(x, b)$ its most significant b bits. If $b(x) < b$ we left-pad x with $(b - b(x))$ zeroes to form a b -bit result. Similarly, $lsb(x, b)$ is used to denote the least significant b bits of x . Let wm be a watermark to be embedded, $wm[i]$ the i -th bit of wm .

In our solution we leverage the one-way cryptographic hash, a special de-facto secure construct. If $crypto_hash()$ is a cryptographic secure one-way hash, of interest are two of its properties: (i) it is computationally infeasible, for a given value V' to find a V such that $crypto_hash(V) = V'$ (one-wayness), and (ii) changing even one bit of the hash input causes random changes to the output bits (i.e. roughly half of them change even if one bit of the input is flipped). Examples of potential candidates for $crypto_hash()$ are the MD5 (used in the proof of concept implementation) or SHA hash. For more details on cryptographic hashes consult [17]. Let $H(V, k) = crypto_hash(k; V; k)$ (where “;” denotes concatenation).

2.3 Related Work

Could existing work in non-media data sets watermarking such as relational data [11] [18] [19] be adapted to the new domain? Our work in [19] requires access to the entire data set in an almost random access model, which is certainly not possible here at embedding time. Also, these efforts seem to make extensive use of the existence of a primary key (or an additional attribute, in [18]), thus rendering a direct adaptation impossible. Moreover, the expected attacks and transformations are different. For example a process of summarization would defeat any of the above schemes. Nevertheless it might be worth noting that, if a primary key is assumed to exist, e.g. if there is a guarantee that the time-stamp information for each stream value is going to be preserved in the result, then both the bit alteration method proposed by Kiernan et al in [11] (for numeric types) and the solution in [18] (for discrete data) could be adapted to work on a single attribute, namely the stream value. The result would likely be resilient to (time-stamp preserving) sampling, but fail with respect to any other attack or transformation.

But what about multimedia watermarking? Given the “streaming” nature of our data, would it not be possible to simply adapt an existing audio (or media) watermarking algorithm [6] [8] [12] [16] [20] since audio data is also an example of a data stream? In other words, why is our problem different? While there seem to be similarities between watermarking audio and sensor data for example, at a closer inspection these similarities prove to be just appearances. A multitude of differences are to be found between the two frameworks mainly deriving from different data models, associated semantic scopes and the itemized nature of sensor stream data.

In theory, a sensor stream could be viewed as an audio signal for example and processed as such. However, for all practical purposes such an approach would

not suit reality and/or often yield undesired results. For example, while in sensor data streams, summarization and sampling are routinely expected natural operations, audio streams are not to be summarized, and sampling in the audio domain entails an entirely different semantic. Summarization for example would not be survived by any of the existing [8] efforts. Moreover, data quality to be preserved in audio streaming is usually related to the human auditory system and its limitations. Any watermark-related alteration can be induced as long as the stream still “sounds” good. In the case of sensor streams (e.g. temperature) on the other hand, many scenarios involve widely different quality metrics, that often need to also consider overall stream characteristics ⁴.

A more in-depth comparison is out of scope here. In summary, while experiences in the multi-media domain are valuable, due to the nature of this new application domain, a solution for watermarking sensor streams needs to naturally handle attacks and transformations such as the ones outlined in Section 2.1.

3 An Initial Solution

This Section outlines the main solution and then gradually improves it to a more robust and resilient version, by identifying and fixing potential flaws.

3.1 Overview

At an overview level, watermark embedding proceeds as follows: (a) first a set of “major” extremes (actual stream items) are identified in the data stream, extremes that feature the property that they (or a majority thereof) can be recovered after a suite of considered alterations (possibly attacks) such as (random) sampling and summarization. Next, (b) a certain criterion is used to select some of these extremes as recipients for parts of the watermark. Finally (c), the selected ones are used to define subsets of items considered for 1-bit watermark embedding of bits of the global watermark. The fact that these extremes can be recovered ensures a consistent overlap (or even complete identity) between the recovered subsets and the original ones (in the un-altered data). In the watermark detection process (d) *all* the extremes in the stream are identified and the selection criteria in step (b) above is used once again to identify potential watermark recipients. For each selected extreme, (e) its corresponding 1-bit watermark is extracted and ultimately the global watermark is gradually re-constructed, by possibly also using error correction (e.g. majority voting).

Thus, one of the main insights behind our solution is the use of extreme values in the stream’s evolution as watermark bit-carriers. The intuition here lies in the fact that much of the stream value lies in exactly its fluctuating behavior and the associ-

⁴e.g. the total alteration introduced per data item should not exceed a certain threshold.

ated extremes, likely to be largely preserved in value-preserving, domain-specific transforms.

3.2 Embedding

Using the notation in Section 2.2, let $\alpha, \beta \in \mathbb{Z}$ such that $\alpha + \beta \leq b(x[])$, where $b(x[])$ is the bit-size of the values in the considered stream ($x[]$). Let χ be a secret integer and $\delta \in (0, 1)$ chosen such that $\delta < 2^{(b(x[]) - \alpha)}$ (i.e. all elements within a characteristic subset $\Xi(\eta, \delta)$ have the same most significant α bits). $\alpha, \beta, \delta, \chi$ are secret. We use the term “advance the window” to denote reading in more new data items while discarding old ones from the current data window.

```

wm_embed( $\delta, \alpha, \beta, wm, k_1, \phi$ )
  while (true) do
     $\eta \leftarrow$  first major extreme in win[]
    compute  $\Xi(\eta, \delta)$ 
     $i \leftarrow H(msb(\eta, \alpha), k_1) \bmod \phi$ 
    if  $i \leq b(wm)$  then
       $bit \leftarrow H(msb(\eta, \alpha), k_1) \bmod \beta$ 
      foreach  $v \in \Xi(\eta, \delta)$  do
         $v[bit - 1] \leftarrow false$ 
         $v[bit] \leftarrow wm[i]$ 
         $v[bit + 1] \leftarrow false$ 
      advance win[] past  $\eta$ 

```

Figure 3: Initial Embedding Algorithm

In the **initial step** of our embedding algorithm we first identify the first major extreme of degree χ and radius δ in the current window. If no major extremes can be found for given δ and χ values, one could consider instead extremes with characteristic subsets smaller than χ that guarantee an acceptable chance (e.g. 70%) of survival in case of sampling (i.e. $\frac{subset_size}{\chi} > 70\%$?).

Once a major extreme (η) is identified in the current window, in the **second step**, a *selection criterion* is used to determine whether η is going to be used in the embedding process or not: if $H(msb(\eta, \alpha), k_1) \bmod \phi = i$ and $i \leq b(wm)$, then η is considered for embedding bit i of the watermark, $wm[i]$. $\phi \in (b(wm), b(wm) + k_2)$ ($k_2 > 0$) is a secret unsigned integer fixed at embedding time, ensuring that only a limited number (a ratio of $\frac{b(wm)}{\phi}$) of these major extremes are going to be selected for embedding. We used a similar “fitness” selection criteria in [18]. Its power derives strength from both the one-wayness and randomness properties of the deployed one-way cryptographic hash, forcing Mallory into a “guessing” position with respect to watermark encoding location. The reason behind the use of the most significant bits of η in the above formula, is resilience to minor alterations and errors due to sampling. As discussed above, the assumption is that for any value $x \in \Xi(\eta, \delta)$, $msb(x, \alpha) = msb(\eta, \alpha)$.

If η is the result of the previous selection step, in the **third step** we embed bit $wm[i]$ into $\Xi(\eta)$. This is done by first, selecting a certain bit position $bit = H(msb(\eta, \alpha), k_1) \bmod \beta$ for embedding. Next, for each value $v \in \Xi(\eta, \delta)$ and in η itself, that bit position is set to

$wm[i]$ and the adjacent bits are set to false (to prevent overflow in case of summarization). In other words $v[bit - 1] = false$, $v[bit] = wm[i]$ and $v[bit + 1] = false$. The reasoning behind modifying an entire subset of items ($\Xi(\eta, \delta)$) is to survive summarizations. This is the case if the bit encoding is such that the average of any combination of ($\nu < |\Xi(\eta)|$ or less) items in $\Xi(\eta, \delta)$, would preserve the embedded bit. It is easy to show that this is indeed the case. Finally, the window is advanced past η and the process re-starts.

3.3 Detection

In the detection process the watermark is gradually reconstructed as more and more of the stream data is processed. The reconstruction process relies on an array of majority voting “buckets” as follows. For each bit $wm[i]$ in the original watermark wm , let $wm[i]^T$ and $wm[i]^F$ be “buckets” (unsigned integers) which are incremented accordingly each time we recover a corresponding true/false bit $wm^{det}[i]$ from the stream. In other words, if the detection process yields at some point $wm^{det}[i] = false$, then the $wm[i]^F$ value is incremented. Similarly, for $wm^{det}[i] = true$, $wm[i]^T$ is incremented. In the end, the actual $wm[i]$ will be estimated by the difference between $wm[i]^T$ and $wm[i]^F$, i.e. if $wm[i]^T - wm[i]^F > v$ then the estimated value for this particular bit becomes $wm^{est}[i] = true$ and conversely if $wm[i]^F - wm[i]^T > v$ then $wm^{est}[i] = false$, where $v > 0$. If detection would be applied on random, un-watermarked data, the probability of detecting $wm^{det}[i] = false$ would equal the probability of $wm^{det}[i] = true$, thus yielding virtually identical (v is used to distinguish this exact case) values for $wm[i]^T$ and $wm[i]^F$. In this case, $wm^{est}[i]$ would be un-defined, thus the data considered un-watermarked. The watermark effectively lies in a statistical bias in the *true/false* distribution for each bit encoding.

```

wm_detect( $\delta, \alpha, \beta, wm, k_1, \phi$ )
  while (true) do
     $\eta \leftarrow$  first extreme in win[]
     $i \leftarrow H(msb(\eta, \alpha), k_1) \bmod \phi$ 
    if  $i \leq b(wm)$  then
       $bit \leftarrow H(msb(\eta, \alpha), k_1) \bmod \beta$ 
      if ( $\eta[bit] = true$ ) then
         $wm[i]^T \leftarrow wm[i]^T + 1$ 
      else
         $wm[i]^F \leftarrow wm[i]^F + 1$ 
      advance win[] past  $\eta$ 

wm_construct( $wm[]^T, wm[]^F, v$ )
  for ( $i \leftarrow 0; i < b(wm); i \leftarrow i + 1$ )
    if ( $wm[i]^T - wm[i]^F > v$ ) then
       $wm[i] \leftarrow true$ 
    else
      if ( $wm[i]^F - wm[i]^T > v$ ) then
         $wm[i] \leftarrow false$ 
      else
         $wm[i] \leftarrow undefined$ 
    return  $wm[]$ 

```

Figure 4: Initial Detection Algorithm

Detection starts by identifying the first extreme

η in the current window. The selection criteria deployed in the embedding phase is tested on η . If $H(msb(\eta, \alpha), k_1) \bmod \phi = i$ and $i \leq b(wm)$, then η was likely used in embedding bit i of the watermark, $wm[i]$. This bit is then extracted from bit-position $H(msb(\eta, \alpha), k_1) \bmod \beta$ and depending on its value, the corresponding bucket $wm[i]^T$ or $wm[i]^F$ is incremented. Finally, the window is advanced past η and the process re-starts.

4 Improvements

Given the initial solution introduced above, we devised a set of improvements aimed at boosting its resilience level including: the ability to handle correlation detection attacks, handling repeated labels, label reconstruction after attacks, introducing a certain hysteresis in the label reconstruction scheme, aimed at defeating targeted extreme values altering attacks, alternative encodings, handling ability of offline multi-pass detection, multi-layer marks aiming to better handle summarization. Due to space constraints we now discuss some of the more important ones.

4.1 Defeating Correlation Detection

One particular issue of concern in the above solution is the fact that because there exists a correlation between the watermarking alteration (the $wm[i]$ bit) and its actual location (determined by $H(msb(\eta, \alpha), k_1)$), Mallory can mount a special attack with the undesirable result of revealing the mark embedding locations. The attack proceeds by first realizing that, despite the one-wayness of the deployed hash function $H()$, in fact, η is the only variable that determines *both* the bit embedding location as well as its value. Mallory can now simply build a set of “hash buckets” for each separate value of $msb(\eta, \alpha)$ (if α is secret the job becomes harder but not impossible) and count, for each extreme η encountered, which of the lower β bits of η is set (resp. reset) more often. For each η for which a bias in a bit position is discovered, that particular bit position is considered mark-carrying and randomized.

Thus, the problem lies here in the correlation between the actual bit location and the bit value, correlation induced by the fact that a single variable (η) determines both of these. A fix could possibly rely on a separate source of information to determine the location of the embedded bit, independently of the bit value. Also, this source of information would need to be consistently recoverable at detection time. For example, if time-stamp information would be assumed available, i.e. if all the processing and the attacks on the data stream could be assumed to preserve the time-stamp to value association, then the actual time-stamp would present an ideal candidate, effectively labeling each and every stream extreme uniquely while at the same time not being correlated (directly) to their values. This unique label could then be used in computing the bit position for embedding. In the se-

lection of the bit embedding location, instead of using $bit = H(msb(\eta, \alpha), k_1) \bmod \beta$ which yields a result correlated to the actual embedded bit value ($wm[i]$, where $i = H(msb(\eta, \alpha), k_1) \bmod \phi$) we propose to use $bit = H(msb(label(\eta), \alpha), k_1) \bmod \beta$ where $label(\eta)$ is the (virtually) unique label of extreme η . A labeling scheme like this would make “bucket counting” attacks impossible. In our model however, timestamps are not assumed to be preserved. Can we envision a different labeling scheme (at least) for extremes, that would survive the attacks and transformations outlined in Section 2.1? We propose to build it from scratch.

One of the challenging aspects of such a labeling scheme becomes clear when one considers data segmentation. To support segmentation, it needs to function based solely on information available close (in terms of stream location) to the considered to-be-labeled extreme. Also, labels computed at detection time from potential segments of sampled and/or summarized data, need to (at least) converge to the original ones, as more and more watermarked data is available. Let λ be the (secret) bit length of the labels resulting in our labeling scheme. Let $\varrho > 1$ be a (secret) unsigned integer. We propose the following labeling scheme: given two extremes i and a subsequent $i + \varrho$, we define $label_bit(i, i + \varrho) = true$ iff $msb(abs(val(i)), \alpha) < msb(abs(val(i + \varrho)), \alpha)$ and *false* otherwise. We define the label for extreme $i + \lambda$, $label(val(i + \lambda))$ is defined as the bit string composed of the concatenation of “1” (binary true) followed by each and every $label_bit(j, j + \varrho)$ in ascending order of $j \in (i, i + \lambda)$. In other words, an extreme is labeled by a certain differential interpretation of some of the preceding extreme values, e.g. in Figure 2 (a), the label for extreme **K** becomes “110100” ($\varrho = 2$).

Before going any further, let us analyze what happens if an important extreme is “lost”, e.g. if one extreme i is altered so much that its α most significant bits flip the $msb(abs(val(i)), \alpha) < msb(abs(val(i + \varrho)), \alpha)$ inequality, corrupting its corresponding label bit. What happens is in fact not too damaging: labels that were constructed using this particular extreme will be corrupted, until the detection process encounters again a continuous sequence of extremes not altered beyond recognition. But Mallory cannot afford altering extremes to such extents, and the secrecy of ϱ makes a random alteration attack the only choice.

In summary, the main purpose of such a labeling scheme is to ensure that Mallory cannot mount the “bucket counting” type of statistical analysis attack as outlined above. Different labels for adjacent extremes together with the use of one-way hashing completely defeat such an attack. The labeling scheme provides an independent, un-correlated source of information for determining the bit position to be altered.

4.2 Reconstructing Labels

Labeling, while providing a defense for the correlation attack, introduces the requirement to be able to iden-

tify major extremes at detection times, possibly in a summarized and/or sampled stream. This becomes a challenge as the definition of “major” does not make sense anymore in the context of a sampled version of the original stream. We propose the following solution. In a first stage, the degree of the transformation performed is determined. In a second stage, the definition of majority of an extreme is updated to reflect the fact that the considered stream is already transformed. A major extreme of degree χ and radius δ in the original stream $(x[], \varsigma)$, becomes a major extreme of degree $\frac{\chi}{\gamma}$ and radius δ in the transformed stream $(x'[], \frac{\varsigma}{\gamma})$, where γ is the degree of the transformation (e.g. summarization, sampling) applied to $(x[], \varsigma)$. Once we know γ identifying major extremes in the transformed stream is simply a matter of considering this updated definition. In a dynamic stream, with consistent stream data rates, γ can be determined by simply dividing the original stream rate to the current (transformed) stream rate, $\gamma = \frac{\varsigma}{\varsigma'}$. The more challenging scenario is to determine the value of γ corresponding to a (possibly transformed) stream $(x'[], \varsigma')$ for which only a segment is available. A reasonable assumption that can be made is that the transform was applied uniformly to the entire stream. In this case, one solution would start by preserving some information about the initial stream, namely the average size of the characteristic subsets of extremes, for a given δ . Then, in the transformed segment, extremes are identified and their average characteristic subset size for the same δ is computed. It is to be expected (arguably) that in a transformed (sampled and/or summarized) stream these sizes would shrink according to the actual transform degree. Dividing the original average characteristic subset size by the sampled stream average would thus yield an estimate of the transform degree γ . In our proof of concept implementation this method is used successfully. Space considerations prevent further elaboration.

4.3 Defeating Bias Detection

But what prevents Mallory from identifying all the major extremes for which there exists a majority of (possibly all) items in the characteristic subset with a certain bit position set to the same identical value? These extremes would then be (rightfully so) considered watermark carrying and Mallory could mount a simple attack of randomizing those bit positions. We propose a new approach that survives summarization and results in alterations effectively appearing random to the eyes of an attacker. Let $\Xi(\eta, \delta) = \{x_1, x_2, \dots, x_a\}$. For each $i \leq j \in [1, a]$, let $m_{ij} = \frac{\sum_{u \in [i, j]} x_u}{|j-i+1|}$. Then we define the *characteristic subset bit encoding convention* as follows: (i) we say that a bit value of “true” is embedded in $\Xi(\eta, \delta)$ iff $\forall j, i$ we have $lsb(H(lsb(m_{ij}, \beta), label(\eta)), \zeta) = 2^\zeta - 1$; similarly, (ii) we say that “false” is embedded iff $\forall j, i$ we have $lsb(H(lsb(m_{ij}, \beta), label(\eta)), \zeta) = 0$, where $\zeta > 0$ is a secret fixed at em-

bedding time. The embedding method simply alters the least significant β bits in the values in $\Xi(\eta, \delta)$ until the criteria is satisfied for the desired to-be-embedded $wm[i]$ bit value. It is to be noted that these alterations should aim to minimize the Euclidean distance (or possibly any other desired distance metric) from the starting point defined by $\{x_1, x_2, \dots, x_a\}$. We call this a “multi-hash encoding”.

The use of m_{ij} ensures survival to summarization, while the cryptographic hash provides the appearance of randomness. But is it feasible to assume that one could find such a point in the a -dimensional space defined by the items in $\Xi(\eta, \delta)$? How many computations are required to at least find one? There are $\frac{a(a+1)}{2}$ possible m_{ij} averages (including all $m_{ii} = x_i$ values). For each we consider the last ζ bits of its hash, effectively getting an output space of $\zeta \frac{a(a+1)}{2}$ bits. The probability that a desired pattern occurs in this space is then $2^{-\zeta \frac{a(a+1)}{2}}$. Thus, on average, the expected number of configurations in the input space that would need to be tested in an exhaustive search before yielding one that results in the desired output, is $2^{\zeta \frac{a(a+1)}{2}}$. For example if $\zeta = 1$ and $a = 5$ we have 2^{15} , that is, approx. 32,000 computations would need to be performed (for each considered major extreme in the window). See Section 6.4 for an experimental analysis.

Given the exponential nature of the increase in required computations for an increasing number of items in the characteristic subset, it is probably not likely to be able to exhaustively handle subsets with more than 8 – 10 items efficiently. While out of the scope of the current paper, the design and use of efficient pruned-space algorithms would be required to significantly reduce these requirements. Alternately, we could deploy a computation-reducing technique that limits the number of m_{ij} averages for which (i) or (ii) needs to hold in the subset bit encoding convention above. In other words, the search process (in the $\{x_1, x_2, \dots, x_a\}$ space) will be stopped once a certain number of the m_{ij} averages feature the desired encoding convention ((i) or (ii)). We call these m_{ij} values “active”. The resulting decrease in required computation time comes at the expense of decreased resilience to transforms.

Also, an (arguably) fast(er) encoding than the use of cryptographic hashes above could be adapted from [1]. The method works by altering the β least significant bits until every one of the longest k pre-fixes of the whole value (most significant bits included), when treated as an integer, becomes a quadratic residue modulo a secret large prime, for embedding a “true” value and a quadratic non-residue modulo the secret prime for embedding a “false” value.

4.4 On-the-fly Quality Assessment

In any watermarking framework, it is important to preserve structural and semantic properties of the watermarked data. Because by its very nature, watermarking alters its input, one has to provide a mech-

anism ensuring that these alterations do not degrade the data beyond usability. In [19] we introduced this essential desiderata in the relational data watermarking framework. We propose to augment our sensor stream marking algorithm with such semantic constraints awareness. Each data property that needs to be preserved is written as a constraint on the allowable change to the dataset, the watermarking process is then applied with these constraints as input and re-evaluates them continuously for each alteration. An “undo” log (quite like the “rollback” log in [19]) is kept to allow undo operations in case certain constraints are violated by the current watermarking step (see Figure 5). The new challenges in this framework are related to the fact that now, due to storage limitations, any data quality preservation constraints can only be formulated in terms of the current available data window. Likely only few window slots can be used to store data aggregates, possibly including some history information to be used in the quality evaluation process but this will all come at the expense of being unable to store and process as much new incoming data. Space considerations prevent us to elaborate further.

5 Analysis

In this Section we analyze the ability of our method to convince in court, survive attacks and transforms.

Court-convincibility can be naturally expressed as follows: given a one bit (e.g. true) watermark, what is the probability of false positives (P_{fp}) for the watermark encoding? In other words, we ask: *What is the probability of a one-bit (true) watermark to be detected in a random data stream?* If this probability is low enough, then a positive detection would constitute a strong proof of rights, with a “confidence” of $1 - P_{fp}$. Here we define confidence as the probability that a given detected watermark was indeed purposefully embedded in the data by the rights owner.

Using the notation in Section 4.3, for each considered extreme η , the occurrence probability of a “good” corresponding m_{ij} (i.e. encoding “true” with respect to the bit encoding convention) in a random stream is naturally $\frac{1}{2}$, because of the cryptographic hash used in the encoding. There are $\frac{a(a+1)}{2}$ possible m_{ij} averages (including all $m_{ii} = x_i$ values). Because for each we consider the last ζ bits of its hash, we effectively have an output space of $\zeta \frac{a(a+1)}{2}$ bits. Thus the probability of the bit “true” being encoded consistently by all of these becomes $2^{-\zeta \frac{a(a+1)}{2}}$ (per extreme). Now, for each $\varepsilon(\chi, \delta)$ items there is a potential major extreme recipient of a one-bit encoding. Out of these how many are actually selected for encoding? As discussed in Section 3.2 only a fraction of $\frac{1}{\phi}$ (because now $b(\text{wm})=1$) of them are actually selected for embedding. Thus if ζ is the stream data rate, we can determine the relationship between the time elapsed since we started reading the incoming stream (t) and the reached level of persuasiveness, as follows.

If $\varepsilon(\chi, \delta)$ models the average number of items that need to be read before a major extreme is encountered, then $\frac{\varepsilon(\chi, \delta)}{\zeta}$ represents the average time-interval “between” major extremes. But only $\frac{1}{\phi}$ of the major extremes are selected for embedding, and so the time-interval between two major extremes that encode the watermark is $\frac{\phi \varepsilon(\chi, \delta)}{\zeta}$. In a time interval of t we are thus likely to see $\frac{t \zeta}{\phi \varepsilon(\chi, \delta)}$ extremes.

As discussed above, each major extreme has an associated probability of false positives of $2^{-\zeta \frac{a(a+1)}{2}}$, thus if we discover a consistent pattern of embedding in a time interval t , the probability of a false-positive becomes $P_{fp}(t) = (2^{-\zeta \frac{a(a+1)}{2}})^{\frac{t \zeta}{\phi \varepsilon(\chi, \delta)}}$. For example if $\zeta = 1$, $a = 5$, $\zeta = 100Hz$, $\phi = 20\%$, $\varepsilon(\chi, \delta) = 50$, after detecting a bit “true” for only $t = 2$ seconds we have $P_{fp}(2) = (2^{-15})^{20} \approx 0$ and an associated proof of rights, with a confidence of close to 100%. Even, at the limit, when due to transforms such as sampling and summarization, for each extreme, only one single m_{ij} average survives and the probability of false positives for each extreme becomes only $\frac{1}{2}$, $P_{fp}(2)$ becomes roughly only “one in a million”. Thus, the persuasion power of our method quickly converges to a comfortable level. In Section 6 we provide experimental results for watermark resilience to various transforms, including random attacks.

Next we explore a theoretical analysis of the vulnerability of our scheme under the following attack model: Mallory starts to modify randomly every a_1 -th ($a_1 > 1$) extreme (η) in such a way as to alter a ratio of $a_2 \in (0, 1)$ of the items in the extreme’s characteristic subset of radius a_3 , $\Xi(\eta, a_3)$. (Thus, on average, Mallory alters only one in every $a'_1 = a_1 \phi$ bit-carrying extremes).

The assumption here is that these alterations do not impact the associated labeling scheme, in other words, they don’t change the “greater than” relationship between extremes used in the labeling process. An extension considering this case is out of the current limited-space scope. Due to space constraints we are going to focus directly on a more challenging, “informed”, Mallory, aware of the characteristic subset radius used at encoding time. This will strengthen our derived bounds. In other words, we assume that $a_3 = \delta$ is known to Mallory, see Section 3.2.

We propose two ways to analyze the vulnerability of the proposed solution: (i) looking at how much an attack “weakens” the encoding, i.e. how many of the *active* m_{ij} values are actually destroyed divided by the total number of active ones (making it thus proportionally harder to detect a watermark in court) and (ii) what is the probability that *all* of the active ones are obliterated? It is easy to see that, for a given extreme η , for which $\Xi(\eta, a_3) = \{x_1, x_2, \dots, x_a\}$ the number of corresponding m_{ij} values altered is $c_m = \frac{1}{2} a a_2 (2a - a a_2 + 1)$.

Now, for (i) the “weakening” of the encoding can be defined as $c_m \times \frac{2}{a(a+1)}$, the ratio of m_{ij} values

that are altered from the total number of potential active ones for each altered extreme. Because one in every $a'_1 = a_1\phi$ bit-carrying extremes gets impacted, the overall “weakening” factor can be defined as $a_1 \times c_m \times \frac{2}{a(a+1)}$. To answer (ii) we first model this scenario by a sampling experiment without replacement. In this experiment, $x + t, t > 0$ balls are randomly removed from a bowl with a total of y balls. The question answered is: if the bowl contained exactly x black balls what is the probability that the $x + t$ removals emptied the bowl of all y black balls. It can be shown that this is $P(x + t, x, y) = \frac{\binom{y-x}{t}}{\binom{y}{x+t}}$. In our model $(x + t) = c_m$, $y = a(a + 1)\frac{1}{2}$ and if $x = a_4y$ (a_4 is the ratio of active m_{ij} values) we can compute the probability that *all* of them are altered.

Thus, for each attacked extreme we have a non-zero probability of altering all active m_{ij} values and removing the corresponding watermark bit. Next we ask, how do these alterations impact our ability to convince in court and detect a watermark bias in the resulting data? Because the alteration is necessarily random (the randomness of the one-way hashes in the encoding in Section 4.3 guarantee this) we can model the attack as essentially a random noise addition attack. Evaluating the resilience of any watermark bias becomes now a matter of asking how many of the embeddings actually survive until detection time. Are there enough of them to actually convincingly reconstruct the multi-bit watermark after error correction? At the beginning of the section we looked at how the watermark bias becomes more convincing in time (and seen data). Losing a fraction of the mark bit encoding extremes can be in fact seen as a reduction of the ϕ value (see Section 3.2). If, for each of the $a'_1 = a_1\phi$ bit carrying extremes that are altered by Mallory, the attack success probability is given by $P(x + t, x, y)$, we can perform a similar reasoning with a new $\phi' = \phi + a'_1 \times P(x + t, x, y)$. What now happens is that the persuasiveness (court-time convince-ability) converges proportionally slower. In other words, we need to see $a_1 \times P(x + t, x, y)$ more stream data to be able to provide an equally convincing proof in court.

For example, for $a_1 = 5$, $a = 6$, $a_4 = 50\%$, $a_2 = 50\%$ we get the average probability $P(15, 10, 21) \approx 0.85\%$ of a complete alteration of all the active m_{ij} values at each extreme. This effectively translates in the need to see only an average of $a_1 \times P(x + t, x, y) \approx 4.25\%$ more data to be equally convincing at detection.

But how does our encoding handle transforms? By construction it certainly survives sampling (A2) up to a degree of $\chi_{max} = |\Xi(\eta, \delta)|$. Indeed this is so if at least one element in the characteristic subset of η is to be found in a sampling of degree χ_{max} . This element can be used in the detection process to recover the corresponding watermark bit for η . Higher degrees of sampling are also quite likely to be survived as there is a non-zero probability of elements in $\Xi(\eta, \delta)$ to be in the sampled stream even for $\chi > \chi_{max}$. Due to

space constraints we do not elaborate further. This is experimentally analyzed in Section 6.

Summarization (A1) up to a degree of $\nu_{max} = |\Xi(\eta, \delta)|$ is also handled well by design, for example due to the use of m_{ij} in the bit-encoding procedure illustrated in Section 4.3. Any summarization of a degree $\nu \leq \nu_{max}$ naturally results in at least one of the m_{ij} averages being in the summarized stream. Even in the initial algorithm, the bit encoding pattern used on the elements in the characteristic subset ensured survival of the pattern in the process of averaging (thus surviving summarization) within the subset. Summarization is experimentally analyzed in Section 6.

But how well is segmentation (A3) survived? More specifically, what is the minimum size of a stream segment from which we are able to recover the watermark? For simplicity let us assume a one-bit watermark, i.e. $b(wm) = 1$. In the following we are trying to determine the minimum required size of a contiguous watermarked stream segment that would enable a proof more “convincing” than a coin-flip stating that a watermark is embedded in the data. This proof would be obtained if we can correctly detect at least two consistent bits (equal to $wm[0]$) from two different extremes found in the segment. In that case, the probability of a false-positive becomes lower than a random coin-flip. But what is the minimum amount of data we need to see to be able to decode two bits? In the best case, the two extremes are adjacent and we need to see enough data to build correct labels for those two extremes. To build the labels correctly, we need to have seen all the previous $\lambda\varrho$ major extremes correctly. Further qualitative analysis must be data dependent, for example if the fluctuating nature of the stream features a major extreme of degree χ and radius δ for every $\varepsilon(\chi, \delta)$ data items, then the minimum required size of a segment enabling watermark detection is $\varepsilon(\chi, \delta)\lambda\varrho$.

6 Experimental Results

We implemented **wms.*** a Java proof-of-concept of the watermarking solution. Our experimental setup included one 1.8GHz CPU Linux box with Sun JDK 1.4 and 384MB RAM.




Figure 5: Overview of proof of concept implementation.

We implemented also a temperature sensor synthetic data stream generator with controllable parameters, including the ability to adjust the data stream

distribution, fluctuating behavior (e.g. $\varepsilon(\chi, \delta)$) and rate (ς). This sensor was used in the initial design phase of some of our experiments because of the ability to produce various fine-tuned data inputs impacting specific strengths of the encoding.

We explored experiment scenarios modeling both the behavior of sub-systems such as the on-the-fly labeling module as well as the overall watermark resilience. Synthetic (temperature sensor model) and real-world data was used in our evaluation.

Because, as discussed in Section 3.3, watermark encoding relies on altering a certain secret statistical bias within the data, when we present resilience results we refer to the ability to detect and reconstruct this bias as an overall measure of encoding performance. In this case, the notion of a “watermark bias” refers to the number of instances of *active* extremes for which the characteristic subset bit encoding (see Section 4.3), survives with a positive true-bit embedding bias⁵.

Unless specified otherwise, the experimental results presented here refer to an underlying normalized stream with values distributed normally with a mean of 0 and a standard deviation of 0.5. The fluctuating behavior of the stream was determined by an average $\varepsilon(\chi, \delta) = 100$ (100 items per each major extreme) and $\varsigma = 100Hz$ (100 items per second). Other parameters include: $\phi = 3$, $\alpha = 16$, $\beta = 16$, $v = 2$, k_1 was chosen by a random number generator. Whenever exact quantitative results are shown, they refer to a data set drawn from about 50 seconds of stream data (i.e. roughly 5000 data values). Additionally, when experiments were performed on real-life test data this is specified in the figure captions. The real life data sets [14] were obtained from the environmental monitors of the NASA Infrared Telescope on the summit of Mauna Kea, at the University of Hawaii. They represent multiple sets of once-every-two-minutes environmental sensor (i.e. temperature) readings at various telescope site locations. The reference data set used refers to 30 days worth of data from the month of September 2003, totaling a number of 21630 temperature readings (with values on the Celsius scale roughly between 0 and 35 degrees).

6.1 Random Alterations

In [19] we defined the *epsilon-attack* in the relational data framework, a transformation that modifies a percentage τ of the input data values within certain bounds defined by two variables ϵ (amplitude of alteration) and μ (mean of alteration). Epsilon-attacks can model any uninformed, random alteration – often the only available attack alternative. A *uniform altering* epsilon-attack (as defined in [19]) modifies τ percent of the input tuples by multiplication with a uniformly distributed value in the $(1 - \epsilon + \mu, 1 + \epsilon + \mu)$

⁵With respect to court-time confidence, for example, a detected watermark bias of 10 yields a false-positive probability of $\frac{1}{2^{10}}$, and an associated proof of rights with a confidence of roughly 99.9%, as discussed in Section 5.

interval. We believe this attack closely resembles (A6), a very likely combination of (A5) and (A2). In Fig-




Figure 6: Label alteration for increasingly aggressive uniform altering epsilon attacks. (a) Different label bit sizes shown. A smaller label size seems to survive better. (b) Different altered data percentages shown.

ures 6 and 7 ($\mu = 0$) we analyze the sensitivity of both our labeling module and overall watermarking scheme to such randomly occurring changes, as direct measures for encoding resilience. In Figure 6 (a), label alteration increases with an increasing degree of data change. Smaller label bit sizes seem to better survive such an attack. In Figure 6 (b), as the percentage of altered data items increases, the labeling scheme naturally degrades.




Figure 7: Watermark survival to epsilon-attacks. (a) Naturally, increasing τ and ϵ values result in a decreasing watermark bias. (b) Same shown for $\epsilon = 10\%$. (**real data**)

In Figure 7, an embedded watermark (bias) is detected in a randomly altered stream. Naturally, an increasing distortion results in a decreasing bias detection. Nevertheless, it is to be noted that the encoding scheme proves to be quite resilient by design, for example for $\tau = 50\%$ of the data altered within $\epsilon = 10\%$ (Figure 7 (b)), the detected bias is still above 25, yielding a false-positive rate of less than “one in thirty million”.

6.2 Sampling and Summarization

The ability to survive summarization (A1) and sampling (A2) is of extreme importance as both are expected attacks. In Figure 8 the labeling algorithm is evaluated with respect to (a) sampling and (b) summarization. Intuitively, a higher label bit-size results in increased fragility to sampling (shown is a sampling degree of 3). Summarization seems to be naturally survived by our design. For example, a summarization of the data down to 5% ($\nu = 20$) still preserves

over 20% of the original label values, thus conferring a strong back-bone to watermark embedding.




Figure 8: (a) Label resilience under sampling conditions. A higher label bit-size naturally yields an increased fragility to sampling. (b) Label alteration for summarization of increasing degree.

The behavior of the watermark encoding algorithm to sampling and summarization is outlined in Figure 9. The natural strength of the bit encoding convention is clearly illustrated here. Both transformations are survived extremely well.




Figure 9: (a) Watermark survival to summarization. An increasing summarization degree results in a decreasing detected watermark bias. (b) Watermark survival to sampling. A bias of 10 ensures a true-positive probability of 99.999%. (real data)

6.3 Segmentation. Combinations

In Section 5 we theoretically assessed the ability of our scheme to survive segmentation (A3), by answering the question: what is the minimum size of a stream segment from which we are able to recover the watermark? In Figure 10 (a) we analyze the impact of actual recovered segment size on the detected watermark bias. From a segment of only 2000 stream values we can detect a watermark bias of 10, corresponding to a very convincing low false positive rate of roughly 0.001. In Figure 10 (b) we outline the impact of a *combined* transformation (sampling and summarization) on the watermark embedding. Because of the nature of both transformations and of the resilience featured in each case, the combination seems to be survived equally well. For example, a 25% sampling, followed by a 25% summarization process still yields a watermark bias of up to 20, corresponding to a low false-positive rate of “one in a million”.




Figure 10: (a) Watermark survival to segmentation. (b) Watermark survival to combined sampling and summarization. (real data)

6.4 Overhead and Impact on Data Quality

The proposed watermarking solution is highly adaptive to both speed and space constraints. By far the most computationally intensive operation is the one-bit encoding operation which alters the characteristic subset data to conform to the bit encoding convention defined in Section 4.3. At the expense of embedding resilience, this operation can be sped up significantly by both pruning of the search space or, more importantly, deployment of a computation-reducing technique as described in Section 4.3. Depending on the actual stream rate, these speed-ups can be gradually deployed to be able to keep up with the incoming data. Additionally, the average amount of computation to be performed per window-load of data is defined also by the actual fraction of extremes “selected” to be bit-carriers. This fraction is determined by $\frac{b(wm)}{\phi}$. If the incoming data rate is too high, ϕ can be increased to reduce the workload.

We performed experiments aimed at evaluating the introduced watermarking computation overhead. Unless specified otherwise, we used the multi-hash encoding discussed in Section 4.3 and parameters set such that the resulting watermark survives 100% any combined sampling and summarization up to a degree of 6. First, we compared the computing times required by the watermarking process with the times spent in a simple read and copy model in which each stream item is read and copied to an output port (with fixed writing time-cost). We obtained consistent value classes clearly identifying each of the separate encoding methods presented. It became clear that, as expected, the majority of time is spent in the actual bit encoding convention routine (and not as much in the labeling module). Not surprising, the encoding convention introduced in Section 3.2 performed fastest with an average of only 5.7% increase in processing times per stream item. The poorest performer was the more complex multi-hash routine in Section 4.3 with an average increase of over 1000%, as expected decreasing almost perfectly exponential with the decrease of the guaranteed resilience (see Figure 11 (a)).

There are two lessons to be learned here. First, different encodings should be used for different scenarios with associated value models. For example for a temperature stream with a likely average reading rate of

under 1Hz, deploying the multi-hash encoding routine for high resilience would be best suited whereas in a very fast streaming scenario the encoding in Section 3.2 would perform much better. Additionally, subject to future research is the issue of better pruning algorithms as discussed in Section 4.3.




Figure 11: (a) Computation overhead (iterations) in multi-hash encoding increases with increasing guaranteed resilience (e.g. sampling degree) levels (logarithmic scale). (b) Decreasing the number of considered bit-encoding extremes (increasing ϕ) decreases the impact on mean and standard deviation in the watermarked data.

We also performed experiments evaluating the impact of our encoding on data quality. More specifically we analyzed the alterations incurred by the mean and standard deviation of the stream data. For the above parameter settings, over a large number (12000+) of runs over the real (and synthetic) data sets, the value of the mean of the watermarked stream varied less than a mere 0.21% average from the original. The alteration to the standard deviation also maintained itself nicely within 0.27% of the original data. There exists a tunable trade-off between attack/transformation resilience and the incurred alterations. A lower level of resilience would definitely require less modifications to the data and have a lower impact on global statistics. In Figure 11 (b) we show how decreasing the number of considered bit-encoding major extremes decreases the impact on the average and standard deviation in the result.

7 Conclusions

In the present paper we introduced the issue of rights protection for sensor streams. We proposed a watermarking solution, based on novel ideas such as on-the-fly labeling and watermark encoding, resilient to important domain-specific transforms. We implemented a proof of concept of the proposed solution and evaluated it experimentally on real data. The method proves to be extremely resilient to all considered transforms, including sampling, summarization, random alterations and combined transforms. In upcoming research we propose to analyze streams of categorical data, to investigate other aggregates (instead of averages) in the summarization process (e.g. min, max, most likely value) and to experiment with alternative resilient and fast(er) bit-encodings.

References

- [1] M. J. Atallah and Jr. S. S. Wagstaff. Watermarking with quadratic residues. In *Proc. of IS-T/SPIE Conf. on Security and Watermarking of Multimedia Contents, SPIE Vol. 3657*, pp. 283–288., 1999.
- [2] M.J. Atallah, V. Raskin, C. F. Hempelmann, M. Karahan, R. Sion, K. E. Triezenberg, and U. Topkara. Natural language watermarking and tamperproofing. In *Lecture Notes in Computer Science, Proc. 5th International Information Hiding Workshop 2002*. Springer Verlag, 2002.
- [3] B. Babcock, S. Babu, M. Datar, and Motwani R. Models and issues in data stream systems. In *Proc. ACM Symp. on Principles of Database Systems (PODS)*, page 1.
- [4] D. Carney, U. Cetintemel, M. Cherniack, C. Convey, S. Lee, G. Seidman, N. Stonebraker, M. and Tatbul, and S. Zdonik. Monitoring streams – a new class of data management applications. In *Proceedings of the Int. Conf. on Very Large Data Bases (VLDB)*, 2002.
- [5] S. Chandrasekaran and M. J. Franklin. Streaming queries over streaming data. In *Proceedings of the Int. Conf. on Very Large Data Bases (VLDB)*, pages 203–214, 2002.
- [6] B. Chen and G. W. Wornell. Quantization index modulation: A class of provably good methods for digital watermarking and information embedding. *IEEE Transactions on Information Theory*, 47(4), 2001.
- [7] Christian Collberg and Clark Thomborson. On the limits of software watermarking, August 1998.
- [8] I. Cox, J. Bloom, and M. Miller. Digital watermarking. In *Digital Watermarking*. Morgan Kaufmann, 2001.
- [9] M. Datar, A. Gionis, P. Indyk, and R. Motwani. Maintaining stream statistics over sliding windows. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 635–644, 2002.
- [10] J. Kang, J. F. Naughton, and S. D. Viglas. Evaluating window joins over unbounded streams. In *Proceedings of ICDE*, 2003.
- [11] J. Kiernan and R. Agrawal. Watermarking relational databases. In *Proceedings of the 28th International Conference on Very Large Databases VLDB*, 2002.
- [12] D. Kirovski and H.S. Malvar. Spread-spectrum watermarking of audio signals. *IEEE Transactions on Signal Processing*, 51(4), 2003.
- [13] F. Korn, S. Muthukrishnan, and D. Srivastava. Reverse nearest neighbor aggregates over streams. In *Proceedings of the Int. Conf. on Very Large Data Bases (VLDB)*, 2002.
- [14] NASA. The Hawaii University Infrared Telescope Facility (<http://irtfweb.ifa.hawaii.edu/>).
- [15] J. Palsberg, S. Krishnaswamy, M. Kwon, D. Ma, Q. Shao, and Y. Zhang. Experience with software watermarking. In *Proceedings of ACSAC, 16th Annual Computer Security Applications Conference*, pages 308–316, 2000.
- [16] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn. Attacks on copyright marking systems. In David Aucsmith, editor, *Information Hiding: Second International Workshop*, volume 1525 of *Lecture Notes in Computer Science*, pages 218–238, Portland, 1998. Springer-Verlag.
- [17] Bruce Schneier. Applied cryptography: Protocols, algorithms and source code in c. In *Applied Cryptography*. John Wiley and Sons, 1996.
- [18] Radu Sion. Proving ownership over categorical data. In *Proceedings of the IEEE International Conference on Data Engineering ICDE*, 2004.
- [19] Radu Sion, Mikhail Atallah, and Sunil Prabhakar. Rights protection for relational data. In *Proceedings of ACM SIGMOD*, 2003.
- [20] M. D. Swanson, B. Zhu, and A. H. Tewfik. Audio watermarking and data embedding – current state of the art, challenges and future directions. In J. Dittmann, P. Wohlman, P. Horster, and R. Steinmetz, editors, *Multimedia and Security Workshop at ACM Multimedia*, volume 41 of *GMD*, Bristol, United Kingdom, September 1998. ACM.