Digital Rights Protection outside Multimedia *

Radu Sion †

Information is probably the most valuable asset of humanity today. By enabling comparably cost-free, fast, and accurate access channels to information in digital form, computers radically changed the way we think and express ourselves. As increasingly more of it is produced, packaged and delivered in digital form in a fast, networked environment, one of its main features threatens to become its worst bug: zero-cost verbatim copies. The inherent ability to produce duplicates of digital Works for virtually no incurred cost can be now misused for illegal profit. This dramatically increases the requirement for an effective rights protection mechanism in the digital world. Different avenues are available, each with its own advantages and drawbacks. Enforcement by legal means is usually ineffective in this new framework, *unless* augmented by a digital counter-part such as Information Hiding. Digital Information Hiding as a method of Rights Protection (also known as Digital Watermarking), hides an indelible "rights witness" (watermark) within the digital Work to be protected, by slightly altering it. The soundness of such a method relies on the assumption that (i) the insertion of the mark does not destroy the value of the Work (i.e. it is still useful for its *intended purpose*); and that (ii) it is difficult for a malicious adversary (Mallory) to remove or alter the mark beyond detection without destroying the value of the Work. Mallory, and the ability to resist his attacks (mostly aiming at removing the embedded watermark) turn out to be one of the major concerns in the design of a watermarking solution.

There exists a multitude of semantic frameworks for information processing and distribution. Each distinct data domain would benefit from the availability of a suitable watermarking solution. The overwhelming majority of research efforts in digital watermarking have been invested in the multimedia data domain (e.g. images, video and audio) [5] [6]. Very recently, other data domains have also been considered, such as natural language [1] and software [4]. In this work we analyze Digital Information Hiding as a method of Rights Protection from a higher level, domain-independent perspective. We propose a theoretical model for Watermarking. We ask: what are the limits of Watermarking? When can these be reached? We then propose, design and analyze watermarking solutions for (i) numeric and categorical relational data (ii) streams and (iii) arbitrary semi-structured content.

Model. In [13] we introduce a model for Watermarking. We define important concepts including: *usability domain* - set of functionals quantifying a digital Work's value in terms of its specific use (see Figure 1), *watermark* - an induced property of a watermarked Work O', so rare, that if we consider any other Work O'', "close-enough" to the original Work O, the probability that O'' exhibits the same property can be upper-bounded, watermark *vulnerability* - the ability of an attack to succeed against a watermarking scheme. One fundamental difference between watermarking and generic data hiding resides in the main applicability and descriptions of the two domains. Data hiding in general and covert communication in particular, aims at enabling Alice and Bob to exchange messages in a manner as resilient and stealthy as possible, through a medium controlled by evil Mallory. Digital watermarking is deployed in court by Alice to prove rights over a given Work, usually in a scenario where Mallory benefits from using/selling that very same Work or maliciously modified versions of it. In digital watermarking, the actual value to be protected lies in the Works themselves whereas information hiding usually makes use of them as simple value "transporters". Rights assessment can be achieved by demonstrating that a particular Work exhibits a rare property (read "hidden message" or "watermark"), usually known only to Alice (with the aid of a "secret" - read "watermarking key"). For court convince-ability purposes this property needs to be so rare that if one considers any other random Work "similar enough" to the one in question, this property is "very improbable" to apply (i.e. bound rate of false-positives). This defines a main difference from steganography: for its purpose, the specifics of the property (e.g. watermark message) are irrelevant as long as Alice can prove "convincingly" it is she who embedded/induced it to the original (non-watermarked) Work. Thus, in watermarking the emphasis is on "detection" rather than "extraction". Extraction of a watermark is usually a part of the detection but just complements the process up to the extent of increasing the ability to convince in court.

Limits. In a seminal paper [9], the main desiderata and features of multimedia watermarking are outlined generically: it should not degrade the perceived quality of the marked Work; the ability to detect the presence/content of a watermark should require the knowledge of a secret (key); different watermarks in the same Work should not interfere with each other; collusion attacks should not be possible; the watermark should survive any value-preserving transformation.

A common un-proved consensus has been implicitly

*Dissertation Outline.

†Computer Sciences, Purdue University, West Lafayette, IN 47907, sion@cs.purdue.edu

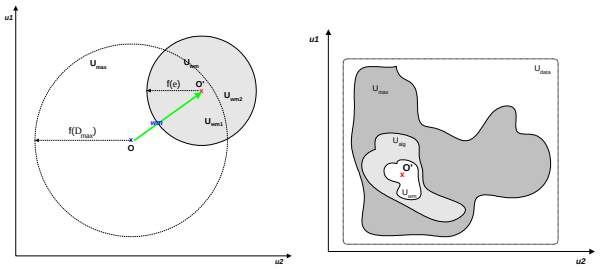

Figure 1: (a) A 2-dimensional view of an usability space. A point uniquely identifies a Work (e.g. coordinates in this space are DCT coefficients). Watermarking is a translation from O to its "watermarked" version, O'. (b) Usability vicinities of a certain Work $O \in \mathbb{D}$ for a given marking algorithm. U_{data} is defined by the actual data type of the usability metrics. U_{max} is the maximal allowable usability vicinity with respect to the associated usability domain(s) (e.g. Human Visual System). U_{wm} is the vicinity in which objects exhibit the watermark.

assumed, namely that watermarking indeed lives up to its claimed features. [5, 6] present excellent area surveys as well as comprehensive examples of algorithms for watermarking (mainly) multi-media Works. We know now that arbitrary large collusion attacks cannot be defeated against [2]. Moreover, while most watermarking algorithms prove to be safe against a considered set of value-preserving transformations (e.g. JPEG compression) they certainly fail with respect to many others. This shortcoming can be directly traced back to the relativity of the "value" and "quality" concepts. Several (mostly experimental) efforts explored the ability to analyze and quantify the "goodness" of watermarking applications, resulting in various watermark benchmarking "suites" (e.g. StirMark: http://www.cl.cam.ac.uk/ ~fapp2/ watermarking/ stirmark, CheckMark: http://watermarking. unige.ch/ Checkmark, OptiMark: http://poseidon.csd.auth.gr/ optimark) mainly for multimedia (i.e. images). Additional research [3, 7, 8] aimed at analyzing concepts such as available bandwidth in the broader area of information hiding from a signal-processing, information-theoretic perspective, focusing mainly on various multimedia techniques. One particular question becomes of interest, namely: *Are there theoretically assessable bounds on watermark vulnerability with respect to an arbitrary watermarking method?* In other words, what is the inherent safety/vulnerability of a generic (i.e. with a minimum amount of assumptions, without considering implementation particularities) watermarking algorithm? An answer to this question might afterward derive real-life recommendations for fine-tuning actual algorithms to increase their marking resilience.

In [11] we explore these and other issues for a broad class of watermarking algorithms. We discover that indeed there exist such limitations. More specifically, we identify an important *convince-ability trade-off*: the more "convincing" in court a watermarking method is, the higher the probability of success of a perfect attack. Moreover we further derive the *watermarking optimality principle* that states that the vulnerability of a watermarking scheme (in our considered class) is likely minimized when it yields watermarked results on the boundary of the maximum allowable usability vicinity of the original un-watermarked Works.

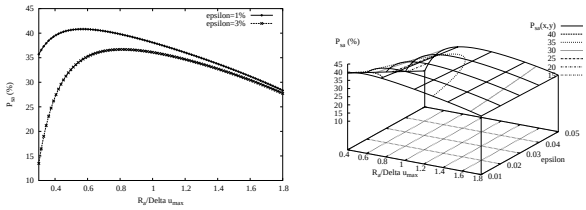

Figure 2: (a) No matter how sophisticated the watermarking method, there exists a random attack with a success probability of 33% and above (although we might not know what the attack is). It can be seen that a lower ϵ value (more convincing in court) yields an even higher upper bound on attack success probability (2D cut through (b)). (b) The 3D evolution of P_{sa} with varying ϵ and $R_a/\Delta u_{max}$. For notations see [11].

From Mallory's perspective this is good news. It turns out that it *is* possible to defeat watermarking algorithms with a surprisingly high success rate, without any additional (insider's) knowledge. This is an inherent limitation of watermarking in general. Any additional knowledge can only improve on this probability. This is the case even if these algorithms conform to the optimality principle. Also, there seems to exist a "sweet spot" in which the probability of a successful attack is maximized. Mallory could make use of this by fine-tuning.

In summary, in this contribution we identified and analyzed inherent limitations of watermarking, including the trade-off between two important watermarking properties: *being enough "convincing" in court* while at the same time *surviving a set of attacks*. In the attempt to become as court convincing as possible, a watermarking application becomes more fragile to attacks aimed at removing the watermark, while preserving the value of the Work. It becomes thus necessary characterized by a significant non-zero probability of being successfully attacked. We discovered an optimality principle (quantified and proved for a broad class of algorithms) that postulates the minimization of vulnerability in specific data points.

Numeric Relational Data. In [15] we introduce a solution for relational database content rights protection through watermarking. Rights protection for relational data is of ever increasing interest, especially considering areas where sensitive, valuable content is to be outsourced. A good example is a data mining application, where data is sold in pieces to parties specialized in mining it. Our solution addresses important attacks, such as subset selection, linear data changes, random alteration attacks and data loss. We introduce wmdb.*, a proof-of-concept implementation and its application to real life data,

namely in watermarking the outsourced Wal-Mart sales data available at our institute.

The main challenges in this new domain derive from the fact that, since the associated data types do not have fixed, well defined semantics (as compared to multimedia) and may be designed for machine ingestion, identifying the available "bandwidth" for watermarking becomes as important as the actual encoding algorithms. Remember that one of the desiderata of watermarking is to insert an indelible mark in the object such that the insertion of the mark does not destroy the value of the object. Clearly, the notion of value or utility of the object is central to the watermarking process. This is closely related to the type of data and its intended use. For example, in the case of software the value may be in ensuring equivalent computation, and for text it may be in conveying the same meaning (i.e. synonym substitution is acceptable). Similarly, for a collection of numbers, the utility of the data may lie in the actual or the relative values of the numbers, or in the distribution (e.g. normal with a certain mean). Because, one can always identify some use of the data that would be affected by even a minor change to any portion of it, it becomes necessary that the intended purpose of the data to be preserved is identified and integrated in the watermarking process.

Our solution starts by receiving as user input a reference to the relational data to be rights-protected, a watermark to be embedded as a copyright proof, a secret key used to protect the embedding and a set of data quality constraints to be preserved in the result. It then proceeds to watermark the data while continuously assessing data quality, potentially backtracking and undo-ing undesirable alterations that do not preserve data quality. Watermark embedding is composed of two main parts: in the first stage, the input data set is securely partitioned into subsets of items; the second stage then encodes one bit of the watermark into each subset. If more subsets (than watermark bits) are available, error correction is deployed to result in an increasingly resilient embedding. The algorithms prove to be resilient to important classes of attacks, including subset selection, linear data changes and random alterations.

The system design, including the ability to evaluate data quality constraints through runtime plugins, is outlined in Figure 3 (a). To exemplify the resilience of our method (e.g. to random alterations), in Figure 3 (b), a comparison is made between the case of uniformly distributed (i.e. values are altered randomly between 100% and 120% of their original value) and fixed alterations (i.e. values are increased by exactly 20%). In the case of fixed alterations the behavior demonstrates the self-healing ability of our method: as more and more of the tuples (past the 50% mark) are altered linearly, the watermark distortion decreases. For example when over 95% of the data is modified consistently and linearly the watermark suffers only 7% alterations.

Another important experiment analyzes the ability to preserve classes in the resulting watermarked object Classification is extremely relevant in areas such as

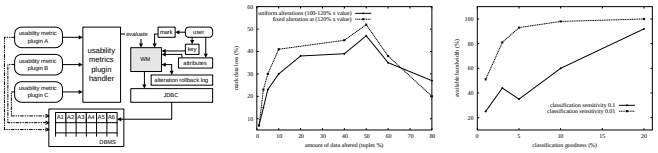

Figure 3: (a) The **wmdb.*** package. (a) Random attack (non-zero average) on a normally distributed data set. (b) Impact of classification preservation on the available watermarking bandwidth.

data mining and we envision that many of the actual deployment scenarios for our relational watermarking application will require classification preservation. Classification preservation deals with the problem of propagation of the classes occurring in the original (input) data in the watermarked (output) version of the data. It provides thus the assurance that the watermarked version still contains most (or within a certain allowed percentage) of the original classes. Figure 3 (c) depicts how classification can be preserved while making optimal use of the available bandwidth. For example, up to 90% of the underlying bandwidth can become available for watermark encoding with a restrictive 6% classification preservation goodness.

These results confirm the adaptability of our watermarking algorithm. As classification tolerance is increased, the application adapts and makes use of an increased available bandwidth for watermark encoding. This also show that classification preservation is compatible with our distribution-based encoding method, an important point to be made, considering the wide range of data-mining applications that could naturally benefit from watermarking ability.

Thus, main contributions of this work include: (i) a resilient watermarking method for relational data, (ii) a technique for enabling user-level run-time control over properties that are to be preserved as well as the degree of change introduced, (iii) a complete, user-friendly implementation for numeric relational data, (iv) the deployment of the implementation on real data, in watermarking the Wal-Mart Sales Database and the analysis thereof.

Categorical Data. While in [15] we propose and analyze the issue of rights protection for numeric relational data, applications handling other types of relational data would certainly benefit from a watermarking solution for these data types. In [10] we introduce a novel method of watermarking categorical data. We discover new watermark embedding channels for relational data with categorical types. We design novel watermark encoding algorithms and analyze important theoretical bounds including mark vulnerability. While fully preserving data quality requirements, our solution survives important attacks, such as subset selection and random alterations. Mark detection is fully "blind" in that it doesn't require the original data, an important characteristic especially in the case of massive data. We propose various

improvements and alternative encoding methods. We perform validation experiments by watermarking the outsourced Wal-Mart sales data available at our institute. We prove (experimentally and by analysis) our solution to be extremely resilient to both alteration and data loss attacks, for example tolerating up to 80% data loss with a watermark alteration of only 25%.

Important new challenges are associated with this domain. One cannot rely on "small" alterations to the data in the embedding process. Any alteration is going to necessarily be significant. The discrete characteristics of the data require discovery of fundamentally new bandwidth channels and associated encoding algorithms. Our method proves to be resilient to important attacks, including subset selection and random alterations.

Our solution starts by discovering two domain-specific watermark embedding channels, namely (i) the *inter-attribute associations* and (ii) the *value occurrence frequency-transform*, (attribute frequency histogram). Next, embedding methods able to resiliently hide information in these channels are designed. The main method starts with an initial user-level assessment step in which a set of attributes to be watermarked are selected. Next, watermark encoding proceeds for each attribute pair (K, A) in the considered attribute set, by selecting a subset of "fit" tuples (determined directly by the association between A and K). These tuples are then considered for mark encoding. Mark encoding alters the tuple's value according to a secret criteria that induces a statistical bias in the distribution for that tuple's altered value. The mark decoding process relies on discovering this induced statistical bias. Yet another embedding method is available to counter extreme vertical partitioning attacks in which only a single attribute A is preserved in the result. If, intuitively, for massive data sets, the number of possible discrete values for A is much smaller than the data set size, then A, contains many duplicate values. There is probably very little value associated with knowing the set of possible values of A. The main value in this scenario (in Mallory's eyes) is (arguably) to be found in one of the only remaining characteristic properties, namely the value occurrence frequency distribution for each possible value of A. If we could devise an alternative watermark encoding method for this set we would be able to associate rights also to this aspect of the data, thus surviving this extreme partitioning attack. In [12] we introduced a watermarking method for numeric sets that is able to minimize the absolute data alteration in terms of distance from the original data set. We propose to apply this method here to embed a mark in the occurrence frequency distribution domain. One concern we should consider is the fact that in the categorical domain we are usually interested in minimizing the *number* of data items altered whereas in the numeric domain we aim to minimize the absolute data change. It is fortunate that, because now we have numeric values modeling occurrence frequency, a solution minimizing absolute data change in this (frequency) domain naturally minimizes the *number* of items altered in the categorical value domain.

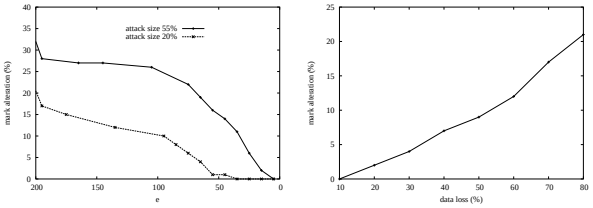

Figure 4: (a) More available bandwidth (decreasing e) results in a higher attack resilience. (b) The watermark degrades almost linearly with increasing data loss.

The experimental results included an analysis of the relationship between the amount of alterations required in the watermarking phase and a minimum guaranteed watermark resilience. It can be seen in Figure 4 (a) that with a decreasing number of encoding alterations (decreasing e) the vulnerability to random alteration attacks increases accordingly. This illustrates the trade-off between the requirement to be resilient and the preservation of data quality (e.g. fewer alterations). An experiment analyzing resilience to data loss is depicted in Figure 4. We observe here the compensating effect of error correction. Compared to data alteration attacks, the watermark survives even better with respect to the attack size (in this case loss of data).

Thus, the main contributions of this effort include: (i) the proposal and definition of the problem of watermarking categorical data, (ii) the discovery and analysis of new watermark embedding channels for relational data with categorical types, (iii) the design of novel associated encoding algorithms.

Streams. Often, streaming information is available 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 [16] we introduce the issue of rights protection for streaming data through watermarking. We propose a solution and analyze its resilience to various types of attacks as well as expected domain-specific alterations, such as sampling and summarization. We implement a proof of concept software (wms.*) and perform experiments to assess these resilience levels in practice. Our method proves to be well suited for this new domain. For example, we can recover an over 97% confidence watermark from a sampled (e.g. less than 8%) stream. Similarly, our encoding ensures survival to stream summarization (e.g. 20%) and random alteration attacks with very high confidence levels, often above 99%.

To the best of our knowledge, the issue of rights

protection for streams has not been addressed. Streaming data sources represent an important class of emerging applications These applications produce a virtually endless stream of data that is too large to be stored in a given system. Recent efforts in the broader area of streaming data, deal with the database challenges of its management. Existing work on discrete data watermarking 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. Moreover, 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. Due to this limitation, earlier work on watermarking relational databases is not applicable to streams. Also, while there seem to be similarities between watermarking multimedia (in particular audio) streams and sensor data, 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 and associated semantic scopes. While in sensor data, 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 process. 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 [1].

A set of novel challenges present themselves in this domain. Any stream processing performed is necessarily both time and space bound. The time bounds derive from the fact that the processing has to keep up with incoming data. The space bounds are referring to the finiteness of any storage mechanism, when compared with the virtually infinite nature of streaming data. At the same time, any quality 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 survive normal domain specific transformations as well as malicious attacks. In this framework we deal with 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) scaling (there might be value in actual *data trends*, that Mallory could still exploit, by scaling the initial values), (A5) addition of stream values and (A6) random alterations.

At an overview level, watermark embedding proceeds as follows: (a) first a set of "major" data extremes (actual stream max/min values) 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 criteria 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 an error correction mechanism. Thus, one of the main ideas 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 associated extremes, more likely to be preserved in value-preserving, domain-specific transforms.

We performed experiments on watermark survival to a variety of transformations, including random alterations and combined sampling and summarization. In Figure 5 (a), random alterations are illustrated. Naturally, an increasing level of distortion results in decreasing detection. Nevertheless, for 50% of the data altered within 10% of the original value, we still detect a watermark bias of roughly 25 bits, yielding a very convincing false- positive rate of less than "one in thirty million". In Figure 5 (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 well. For example, 25% sampling, followed by 25% summarization still yields a watermark bias of up to 20, corresponding to a favorable, low false-positive rate of "one in a million".

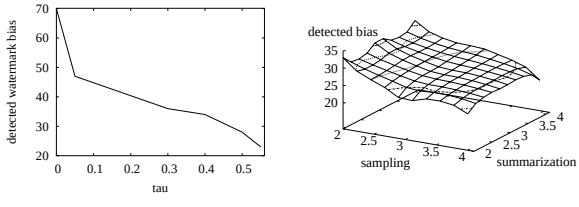

Figure 5: (a) Watermark survival to epsilon-attacks. (b) Watermark survival to combined sampling and summarization.

Thus, the main contributions of this effort include: (i) the proposal and definition of the problem of watermarking streams, (ii) the discovery and analysis

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

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.

Structures. In [14] we discuss the watermarking of abstract structured aggregates of multiple types of content, such as multi-type/media documents. These *semi-structures* can be usually represented as graphs and are characterized by value lying both in the structure *and* in the individual nodes. Example instances include XML documents, complex web content, workflow and planning descriptions. We propose a scheme for watermarking abstract semi-structures and discuss its resilience with respect to attacks. While content specific watermarking deals with the issue of protecting the value in the structure's nodes, protecting the value pertaining to the structure itself is a new, distinct challenge. Nodes in semi-structures are value-carrying, thus a watermarking algorithm could make use of their encoding capacity by using traditional watermarking. For example if a node contains an image then image watermarking algorithms can be deployed for that node to encode parts of the global watermark. Also, given the intrinsic value attached to it, the graph that "glues" these nodes together becomes in itself a central element of the watermarking process that makes use of these two value facets, structural and node-content.

Multiple challenges are encountered in this framework, mostly derived from the requirement to survive domain-specific transformations and likely attacks by Mallory, including: elimination of value-"insignificant" nodes (A1), elimination of inter node relations (A2), value preserving graph partitioning into independent usable partitions (A3), modification of node content, within usability vicinity (A4), addition of value insignificant nodes (A5). Our solution is based on a canonical labeling algorithm that self-adjusts to the specifics of the content. Labeling is tolerant to a significant number of graph attacks ("surgeries") and relies on a complex "training" phase at embedding time in which it reaches an optimal stability point with respect to these attacks. We perform attack experiments on the introduced algorithms under different conditions with very encouraging results. In Figure 6 we show the watermark behavior to data alteration in the case of a random artificially generated structure with 32 nodes and 64 edges. The embedded watermark is 8 bits long. The labeling scheme was trained for 3 surgeries. As the number of attack surgeries increases, the watermark degrades slightly. The results are averaged over 10 runs on the same graph with different random attacks. When 8 attack surgeries are applied to the graph we can still recover 60-65% of the watermark. One has to consider also the fact that an attacker is bound not to modify the structure beyond distortion limits.

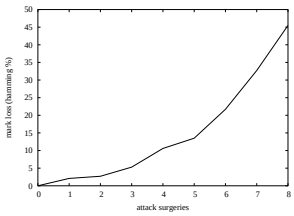

Figure 6: Averaged watermark loss over 10 runs of an 8 bit watermark embedded into an arbitrary 32 node graph with 64 edges. Surgery attacks are applied randomly (node removals 60%, link addition 20%, link removal 20%). The labeling scheme was trained for 3 surgeries.

References

[1] 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.

[2] D. Boneh and J. Shaw. Collusion-secure fingerprinting for digital data. *Lecture Notes in Computer Science*, 963:452–??, 1995.

[3] B. Chen and G. Wornell. An information–theoretic approach to the design of robust digital watermarking systems. In *Proc. Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP), Phoenix, AZ*, 1999.

[4] Christian Collberg and Clark Thomborson. On the limits of software watermarking, August 1998.

[5] I. Cox, J. Bloom, and M. Miller. Digital watermarking. In *Digital Watermarking*. Morgan Kaufmann, 2001.

[6] S. Katzenbeisser and F. Petitcolas (editors). Information hiding techniques for steganography and digital watermarking. Artech House, 2001.

[7] P. Moulin and J. O'Sullivan. Information-theoretic analysis of information hiding. In *manuscript*, 1999.

[8] Pierre Moulin, M. K. Mihcak, and Gen-Iu (Alan) Lin. An information–theoretic model for image watermarking and data hiding. In *(manuscript)*, 2000.

[9] Fabien A. P. Petitcolas, Ross J. Anderson, and Markus G. Kuhn. Information hiding - a survey. *Proceedings of the IEEE*, 87(7):1062–1078, July 1999. Special issue on protection of multimedia content.

[10] Radu Sion. Proving ownership over categorical data. In *Proceedings of the IEEE International Conference on Data Engineering ICDE 2004*, 2004.

[11] Radu Sion and Mikhail Atallah. Attacking digital watermarks. In *Proceedings of the Symposium on Electronic Imaging SPIE*, 2004.

[12] Radu Sion, Mikhail Atallah, and Sunil Prabhakar. On watermarking numeric sets. In *Proceedings of IWDW 2002, Lecture Notes in Computer Science, CERIAS TR 2001-60*. Springer-Verlag, 2002.

[13] Radu Sion, Mikhail Atallah, and Sunil Prabhakar. Power: Metrics for evaluating watermarking algorithms. In *Proceedings of IEEE ITCC 2002, CERIAS TR 2001-55*. IEEE Computer Society Press, 2002.

[14] Radu Sion, Mikhail Atallah, and Sunil Prabhakar. On watermarking abstract semi-structures. In *Proceedings of IWDW 2003, Lecture Notes in Computer Science, CERIAS TR 2001-54*, 2003.

[15] Radu Sion, Mikhail Atallah, and Sunil Prabhakar. Rights protection for relational data. In *Proceedings of ACM SIGMOD*, 2003.

[16] Radu Sion, Mikhail Atallah, and Sunil Prabhakar. Watermarking streams. In *(submitted for review)*, 2004.