

Chapter 2

Related Work

In this chapter, we provide a brief review of related work on privacy preservation and side-channel attacks especially in the three related domains: data publishing, Web applications, and smart metering.

2.1 Privacy Preservation

The privacy preservation issue has received significant attentions in various domains, such as data publishing and data mining [24, 45, 76], networking [9, 11, 44], social networks [30, 43, 71], data outsourcing [17, 82], secure multiparty computation [70], and Web applications [9, 13, 21, 80].

In the context of privacy-preserving data publishing, various generalization techniques and models have been proposed to transform a micro-data table into a safe version that satisfies given privacy properties and retains enough data utility. swapping [29, 34, 38] and cell suppression [28] both aim to protect micro-data released in census tables, but those earlier approaches cannot effectively quantify the degree of privacy. A measurement of information disclosed through tables based on the perfect secrecy notion by Shannon is given in [68]. The authors in [33] address the problem ascribed to the independence assumption made in [68]. The important notion of k -anonymity has been proposed as a model of privacy requirement [76]. The main goal of k -anonymity is to anonymize the data such that each record owner in the resultant data is guaranteed to be indistinguishable from at least $k - 1$ other record owner. That is, each quasi-identifier value in a micro-data should be at least shared by k tuples. Since the data owner modifies the data, some information is distorted. Therefore, it is desirable to find the modified table for k -anonymity with the minimum information loss. However, to achieve optimal k -anonymity with the most data utility is proved to be computationally infeasible [67].

Since the introduction of k -anonymity, privacy-preserving data publishing has received tremendous interest [26, 27, 36, 41, 45, 79, 84]. A model based on the intuition of *blending individuals in a crowd* is proposed in [20]. A personalized requirement for anonymity is studied in [88]. In [16], the authors approach the issue from a different perspective, that is, the privacy property is based on generalization of the protected data and could be customized by users. Much efforts have been made around developing efficient k -anonymity algorithms [3–5, 12, 37, 56, 76], whereas the safety of the algorithms is generally assumed.

Many more advanced models are proposed to address limitations of k -anonymity. Many of these focus on the deficiency of allowing insecure groups with a small number of sensitive values. For instance, l -diversity [65] requires that each equivalence class on the disclosed table should contain at least l well-represented sensitive values; t -closeness [57] requires that the distribution of a sensitive attribute in any equivalence class is close (roughly equal) to the distribution of the attribute in the whole table; (α, k) -anonymity [86] requires that the number of tuples in any equivalence class is at least k and the frequency (in fraction) of each sensitive value is at most α , where k and α are data publisher-specified thresholds. In addition, a generic model called *GBP* was proposed to unify the perspective of privacy guarantees in both generalization-based publishing and view-based publishing [32].

In contrast to micro-data disclosure, aggregation queries are addressed in statistical databases [2, 41, 78]. The main challenge is to answer aggregation queries without allowing inferences of secret individual values. The auditing methods in [22, 35] solve this problem by checking whether each new query can be safely answered based on a history of previously answered queries. The authors of [22, 51, 53] considered the same problem in more specific settings of offline auditing and online auditing, respectively.

Compared with the aforementioned syntactic privacy models, a semantic privacy notation to provide provable resistance to adversaries' background knowledge, differential privacy [39], has been widely accepted as a stronger privacy model mostly for answering statistic queries. Differential privacy aims to achieve the goal that the probability distribution of any disclosed information should be similar enough regardless of whether that disclosed information is obtained using the real database, or using a database without any one of the existing records. However, although differential privacy is extended to privacy preserving data publishing (PPDP) [58, 90], most existing approaches that ensure differential privacy are random noise-based and are suitable for specific types of statistical queries. Further, Kifer et al. [52] disproved some popularized claims about differential privacy and showed that differential privacy cannot always guarantee the privacy in some cases. Moreover, while the qualitative significance of the privacy parameter ϵ is well understood in the literature, the exact quantitative link between this value and the degree of privacy guarantee has received less attention. Actually, more and more works have concluded that both differential privacy and syntactic privacy models have their place, and any one cannot replace the other [25, 59]. With respect to our discussions, differential privacy is less applicable due to the less predictable

but larger sensitivity and the nature of budget sharing among executions of web applications (detailed in Chaps. 5, 6). Due to these reasons, we will focus more on syntactic privacy properties in this book.

2.2 Side-Channel Attacks

Side-channel leaks have been extensively studied in the literature. By measuring the amount of time taken to respond to the queries, an attacker may extract OpenSSL RSA privacy keys [15], and similar timing attacks are proved to be still practical recently [14]. By differentiating the sounds produced by keys, an attacker with the help of the large-length training samples may recognize the key pressed [7]; Zhuang et al. further present an alternative approach to achieving such attack which does not need the training samples [98]. By exploiting queuing side channel in routers by sending probes from a far-off vantage point, an attacker may fingerprint websites remotely against home broadband users [46, 47]. Ristenpart et al. discover cross-VM information leakage on Amazon EC2 based on the sharing of physical infrastructure among users [74]. Search histories may be reconstructed by session hijacking attack [18], while web-browsing histories may be compromised by cache-based timing attacks [42]. Saponas et al. show the transmission characteristics of encrypted video streaming may allow attackers to recognize the title of movies [77].

Meanwhile, many efforts have been made on developing techniques to mitigate the threats of such leakages. Countermeasures based on traffic-shaping mechanisms (such as padding, mimicking, and morphing) are suggested against the exposure of identification of encrypted web traffic in [81]. HTTPOS, a browser-side system, is proposed to prevent information leakages of encrypted HTTP traffic through configurable traffic transformation techniques in [64]. Timing mitigator is introduced to achieve any given bound on timing channel leakage by delaying output events to limit the amount of information in [6]. Zhang et al. present an approach to verifying the VMs' exclusive use of a physical machine. The approach exploits a side-channel in the L2 memory cache as a defensive detection tool rather than a vector of attack [97]. Provider-enforced deterministic execution by eliminating all the internal timing channels has been proposed to combat timing channel attack in cloud [8]. In the rest of this section, we review the work related to the side-channel attacks targeted in each line of our research.

2.3 Side-Channel Leaks in Data Publishing

While most existing work assume the disclosed generalization to be the only source of information available to an adversary, recent work [85, 94] show the limitation of such an assumption. In addition to such information, the adversary may also know about the disclosure algorithm. With such extra knowledge, the adversary

may deduce more information and eventually compromise the privacy property. In the work of [85, 94], the authors discover the above problem and correspondingly introduce models and algorithms to address the issue. However, the method in [85] is still vulnerable to algorithm-based disclosure [49, 50], whereas the one in [94] is more general, but it also incurs a high complexity.

In [94], Zhang et al. presented a theoretical study on how an algorithm should be designed to prevent the adversary from inferring private information when the adversaries know the algorithm itself. The authors proved that it is NP-hard to compute a generalization which ensure privacy while maximizing data utility under such assumptions of adversaries' knowledge. The authors then investigate three special cases of the problem by imposing constraints on the functions and the privacy properties, and propose a polynomial-time algorithm that ensures entropy l -diversity.

Wong et al. in [85] showed that a minimality attack can compromise most existing generalization techniques with the aim of only a small amount of knowledge about the generalization algorithm. The authors assume that the adversaries only have one piece of knowledge that the algorithm discloses a generalization with best data utility. Under this assumption, minimality attacks can be prevented by simply disclosing sub-optimal generalizations. Unfortunately, the adversaries, equipped with knowledge of the algorithm, can still devise other types of attacks to compromise sub-optimal generalizations.

Since the problem was discovered, many solutions have been developed to tackle the problem in the case that l - diversity is the desired privacy property [50, 61, 89, 96]. To improve the efficiency, a so-called exclusive strategy is proposed in [95] to penalize the cases where a recursive process is required to compute the adversarial mental image about the micro-data table. To examine the general case, we have proposed a k -jump strategy [62] (see Chap. 3 for this first line of our research) to penalize such cases where with more control in the sense that only k , instead of all, generalization functions will be skipped. Our proposed family of algorithms is general to handle different privacy properties and different measures of data utility. Despite the improved efficiency, most of those methods are still impractical due to the high complexity.

The concept of l -cover in [96] has been proposed for efficient diversity preservation. However, no concrete methods for building identifier partitions that can satisfy the l -cover property was reported in [96], which is the main focus of the second line of our research(see Chap. 4). The correctness and flexibility of our approach can be further confirmed by the following work in the literature. The authors of [89] introduce algorithms that share the same spirit with our algorithms, and can achieve similar performance (more precisely, their algorithms are slightly less efficient than ours since their time complexity is $O(n^2 \log n)$). In fact, under slight modification, their algorithms, such as ACE algorithm which is originally intended to publish dynamic datasets[91], can be regarded as another instantiation of our model and approach.

2.4 Side-Channel Leaks in Web Applications

In the context of web applications, many side-channel leakages in encrypted web traffic have been identified in the literature which allow to profile the web applications themselves and their internal states [8, 18, 21, 47]. Meanwhile, several approaches have been proposed to analyze and mitigate such leakages, such as [6, 64, 81, 87]. Recently, a black-box approach has been proposed to detect and quantify the side-channel vulnerabilities in web application by extensively crawling a targeted application [19].

Chen et al. in [21] demonstrated through case studies that side-channel problems are pervasive and exacerbated in web applications due to their fundamental features. Then the authors further study approaches to identifying such threats and quantifying the amount of information disclosed in [93]. They show that an application-agnostic approach generally suffers from high overhead and low level of privacy protection, and consequently effective solutions to such threats likely will rely on the in-depth understanding of the applications themselves. Finally, they design a complete development process as a fundamental solution to such side channel attacks. Traffic morphing is proposed in [87] to mitigate the threats by traffic analyzing on packet sizes and sequences through network. Although their proposed system morph classes of traffic to be indistinguishable, traffic morphing pads or splits packets on the fly which may degrade application's performance.

The aforementioned works share an important limitation, that is, they lack formal privacy requirements. In such a case, the degree of privacy breach caused by the transformation of traffic cannot be evaluated during the process of padding. Consequently, it cannot always ensure that privacy is indeed preserved. In contrast, our proposed algorithms described in the third line of our research theoretically guarantee the desired privacy property. Our model and solutions provide finer control over the trade-off between privacy protection and cost, and those solutions can certainly be integrated into the development process of the Web application.

Nonetheless, these solutions assume that adversaries do not possess prior background knowledge about possible user inputs. Our fourth line of research enhances previous works by mitigating the threats of background knowledge. Closest to this research, most recently, a formal framework is proposed to measure security in terms of the amount of information leaked from the observations without the assumption of any particular attacks [10].

2.5 Side-Channel Leaks in Smart Metering

Electrical appliances, even small electric devices, generate detectable electric consumption signatures [48, 55]. Based on such signatures, electric consumption data (collected at a pre-configured granularity) of a household can be decomposed to identify the status of appliances. A domestic electricity demand model based

on occupant time-use data has been presented and its example implementation is made for free download [73]. Worse still, even simple off-the-shelf statistical tools can be used to extract complex usage patterns from the consumption data without a priori knowledge of household activities [69]. Rouf et al. show that real-world automatic meter reading (AMR) systems are vulnerable to spoofing attacks due to the unsecured wireless transmission and continuous broadcast of fine-grained energy data [75].

Many surveys have been conducted to review and discuss the security and privacy requirements and challenges (e.g. [83]). Some efforts have been made to preserve privacy for the load monitoring [40]. In-residence batteries, together with corresponding battery privacy algorithms such as Non-Intrusive Load Leveling (NILL) and stepping approach, used to mask load variance of a household to the grid and consequently avoided recovering of appliance profiles by grid [66, 92]. A distributed Laplacian Perturbation Algorithm (DLPA) has been proposed to achieve provably privacy and optimal utility without the need of a third trusted party [1]. An aggregation protocol is introduced to privately sum readings from many meters without the need of disclose those raw meter readings [54]. A scheme is designed to provide personal enquiry and regional statistics through anonymously-sent readings [23]. EPPA achieves privacy-preserving multi-dimensional data aggregation by using the homomorphic cryptosystem [63].

The other efforts were made to preserve privacy for billing. Rail et al. proposed a set of protocols which allow users themselves to produce a correct provable final bill without disclosing fine-grained consumption data [72], and then the authors combined differentially private mechanisms with oblivious payments to eliminate leakages drawn from the final bill [31]. Recently, Lin et al. used trusted platform module (TPM) and cryptographic primitive to support privacy preserving billing and load monitoring simultaneously [60].

References

1. Gergely Ács and Claude Castelluccia. Dream: Differentially private smart metering. *CoRR*, abs/1201.2531, 2012.
2. N.R. Adam and J.C. Wortmann. Security-control methods for statistical databases: A comparative study. *ACM Comput. Surv.*, 21(4):515–556, 1989.
3. G. Aggarwal, T. Feder, K. Kenthapadi, R. Motwani, R. Panigrahy, D. Thomas, and A. Zhu. k-anonymity: Algorithms and hardness. *Technical report, Stanford University*, 2004.
4. G. Aggarwal, T. Feder, K. Kenthapadi, R. Motwani, R. Panigrahy, D. Thomas, and A. Zhu. Anonymizing tables. In *ICDT'05*, pages 246–258, 2005.
5. G. Aggarwal, T. Feder, K. Kenthapadi, R. Motwani, R. Panigrahy, D. Thomas, and A. Zhu. Approximation algorithms for k-anonymity. *Journal of Privacy Technology*, November 2005.
6. A. Askarov, D. Zhang, and A.C. Myers. Predictive black-box mitigation of timing channels. In *CCS '10*, pages 297–307, 2010.
7. D. Asonov and R. Agrawal. Keyboard acoustic emanations. *Security and Privacy, IEEE Symposium on*, page 3, 2004.

8. A. Aviram, S. Hu, B. Ford, and R. Gummadi. Determinating timing channels in compute clouds. In *CCSW '10*, pages 103–108, 2010.
9. Michael Backes, Goran Doychev, Markus Dürmuth, and Boris Köpf. Speaker recognition in encrypted voice streams. In *ESORICS '10*, pages 508–523, 2010.
10. Michael Backes, Goran Doychev, and Boris Köpf. Preventing Side-Channel Leaks in Web Traffic: A Formal Approach. In *NDSS'13*, 2013.
11. K. Bauer, D. McCoy, B. Greenstein, D. Grunwald, and D. Sicker. Physical layer attacks on unlinkability in wireless lans. In *PETS '09*, pages 108–127, 2009.
12. R.J. Bayardo and R. Agrawal. Data privacy through optimal k-anonymization. In *ICDE*, pages 217–228, 2005.
13. I. Bilogrevic, M. Jadhliwala, K. Kalkan, J.-P. Hubaux, and I. Aad. Privacy in mobile computing for location-sharing-based services. In *PETS*, pages 77–96, 2011.
14. BillyBob Brumley and Nicola Taveri. Remote timing attacks are still practical. In *ESORICS'11*, pages 355–371. 2011.
15. D. Brumley and D. Boneh. Remote timing attacks are practical. In *USENIX*, 2003.
16. J. Byun and E. Bertino. Micro-views, or on how to protect privacy while enhancing data usability: concepts and challenges. *SIGMOD Record*, 35(1):9–13, 2006.
17. N. Cao, Z. Yang, C. Wang, K. Ren, and W. Lou. Privacy-preserving query over encrypted graph-structured data in cloud computing. In *ICDCS'11*, pages 393–402, 2011.
18. C. Castelluccia, E. De Cristofaro, and D. Perito. Private information disclosure from web searches. In *PETS'10*, pages 38–55, 2010.
19. Peter Chapman and David Evans. Automated black-box detection of side-channel vulnerabilities in web applications. In *CCS '11*, pages 263–274, 2011.
20. S. Chawla, C. Dwork, F. McSherry, A. Smith, and H. Wee. Toward privacy in public databases. In *Theory of Cryptography Conference*, 2005.
21. Shuo Chen, Rui Wang, XiaoFeng Wang, and Kehuan Zhang. Side-channel leaks in web applications: A reality today, a challenge tomorrow. In *IEEE Symposium on Security and Privacy '10*, pages 191–206, 2010.
22. F. Chin. Security problems on inference control for sum, max, and min queries. *JACM*, 33(3):451–464, 1986.
23. Cheng-Kang Chu, Joseph K. Liu, Jun Wen Wong, Yunlei Zhao, and Jianying Zhou. Privacy-preserving smart metering with regional statistics and personal enquiry services. In *ASIA CCS '13*, pages 369–380, 2013.
24. V. Ciriani, S. De Capitani di Vimercati, S. Foresti, and P. Samarati. k-anonymous data mining: A survey. In *Privacy-Preserving Data Mining: Models and Algorithms*. 2008.
25. C. Clifton and T. Tassa. On syntactic anonymity and differential privacy. In *ICDEW '13*, pages 88–93, 2013.
26. L.H. Cox. Solving confidentiality protection problems in tabulations using network optimization: A network model for cell suppression in the u.s. economic censuses. In *Proceedings of the Internatinal Seminar on Statistical Confidentiality*, 1982.
27. L.H. Cox. New results in disclosure avoidance for tabulations. In *International Statistical Institute Proceedings*, pages 83–84, 1987.
28. L.H. Cox. Suppression, methodology and statistical disclosure control. *J. of the American Statistical Association*, pages 377–385, 1995.
29. T. Dalenius and S. Reiss. Data swapping: A technique for disclosure control. *Journal of Statistical Planning and Inference*, 6:73–85, 1982.
30. G. Danezis, T. Aura, S. Chen, and E. Kiciman. How to share your favourite search results while preserving privacy and quality. In *PETS'10*, pages 273–290, 2010.
31. George Danezis, Markulf Kohlweiss, and Alfredo Rial. Differentially private billing with rebates. In *IH'11*, pages 148–162, 2011.
32. A. Deutsch. Privacy in database publishing: a bayesian perspective. In *Handbook of Database Security: Applications and Trends*, pages 464–490. Springer, 2007.
33. A. Deutsch and Y. Papakonstantinou. Privacy in database publishing. In *ICDT*, pages 230–245, 2005.

34. P. Diaconis and B. Sturmfels. Algebraic algorithms for sampling from conditional distributions. *Annals of Statistics*, 26:363–397, 1995.
35. D.P. Dobkin, A.K. Jones, and R.J. Lipton. Secure databases: Protection against user influence. *ACM TODS*, 4(1):76–96, 1979.
36. A. Dobra and S.E. Feinberg. Bounding entries in multi-way contingency tables given a set of marginal totals. In *Foundations of Statistical Inference: Proceedings of the Shores Conference 2000*. Springer Verlag, 2003.
37. Y. Du, T. Xia, Y. Tao, D. Zhang, and F. Zhu. On multidimensional k-anonymity with local recoding generalization. In *ICDE*, pages 1422–1424, 2007.
38. G.T. Duncan and S.E. Feinberg. Obtaining information while preserving privacy: A markov perturbation method for tabular data. In *Joint Statistical Meetings*. Anaheim, CA, 1997.
39. C. Dwork. Differential privacy. In *ICALP (2)*, pages 1–12, 2006.
40. Z. Erkin, J.R. Troncoso-Pastoriza, R.L. Lagendijk, and F. Perez-Gonzalez. Privacy-preserving data aggregation in smart metering systems: An overview. *Signal Processing Magazine, IEEE*, 30(2):75–86, 2013.
41. I.P. Fellegi. On the question of statistical confidentiality. *Journal of the American Statistical Association*, 67(337):7–18, 1993.
42. E. W. Felten and M. A. Schneider. Timing attacks on web privacy. In *CCS '00*, pages 25–32, 2000.
43. Philip W. L. Fong, Mohd Anwar, and Zhen Zhao. A privacy preservation model for facebook-style social network systems. In *ESORICS '09*, pages 303–320, 2009.
44. Julien Freudiger, Mohammad Hossein Manshaei, Jean-Pierre Hubaux, and David C. Parkes. On non-cooperative location privacy: a game-theoretic analysis. In *CCS '09*, pages 324–337, 2009.
45. B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu. Privacy-preserving data publishing: A survey of recent developments. *ACM Computing Surveys*, 42(4):14:1–14:53, June 2010.
46. X. Gong, N. Borisov, N. Kiyavash, and N. Schear. Website detection using remote traffic analysis. In *PETS'12*, pages 58–78, 2012.
47. X. Gong, N. Kiyavash, and N. Borisov. Fingerprinting websites using remote traffic analysis. In *CCS*, pages 684–686, 2010.
48. G.W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.
49. X. Jin, N. Zhang, and G. Das. Algorithm-safe privacy-preserving data publishing. In *EDBT '10*, pages 633–644, 2010.
50. X. Jin, N. Zhang, and G. Das. Asap: Eliminating algorithm-based disclosure in privacy-preserving data publishing. *Inf. Syst.*, 36:859–880, July 2011.
51. K. Kenthapadi, N. Mishra, and K. Nissim. Simulatable auditing. In *PODS*, pages 118–127, 2005.
52. D. Kifer and A. Machanavajjhala. No free lunch in data privacy. In *SIGMOD '11*, pages 193–204, 2011.
53. J. Kleinberg, C. Papadimitriou, and P. Raghavan. Auditing boolean attributes. In *PODS*, pages 86–91, 2000.
54. Klaus Kursawe, George Danezis, and Markulf Kohlweiss. Privacy-friendly aggregation for the smart-grid. In *PETS'11*, pages 175–191, 2011.
55. H. Y. Lam, G. S.K. Fung, and W. K. Lee. A novel method to construct taxonomy electrical appliances based on load signatures. *IEEE Trans. on Consum. Electron.*, 53(2):653–660, May 2007.
56. K. LeFevre, D. DeWitt, and R. Ramakrishnan. Incognito: Efficient full-domain k-anonymity. In *SIGMOD*, pages 49–60, 2005.
57. N. Li, T. Li, and S. Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *ICDE*, pages 106–115, 2007.
58. N. Li, W. H. Qardaji, and D. Su. Provably private data anonymization: Or, k-anonymity meets differential privacy. *CoRR*, abs/1101.2604, 2011.

59. Ninghui Li, Wahbeh Qardaji, and Dong Su. On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy. In *ASIACCS '12*, pages 32–33, 2012.
60. Hsiao-Ying Lin, Wen-Guey Tzeng, Shiu-an-Tzuo Shen, and Bao-Shuh P. Lin. A practical smart metering system supporting privacy preserving billing and load monitoring. In *ACNS'12*, pages 544–560, 2012.
61. W. M. Liu and L. Wang. Privacy streamliner: a two-stage approach to improving algorithm efficiency. In *CODASPY*, pages 193–204, 2012.
62. W. M. Liu, L. Wang, and L. Zhang. k-jump strategy for preserving privacy in micro-data disclosure. In *ICDT '10*, pages 104–115, 2010.
63. Rongxing Lu, Xiaohui Liang, Xu Li, Xiaodong Lin, and Xuemin Shen. Eppa: An efficient and privacy-preserving aggregation scheme for secure smart grid communications. *Parallel and Distributed Systems, IEEE Transactions on*, 23(9):1621–1631, 2012.
64. X. Luo, P. Zhou, E. W. W. Chan, W. Lee, R. K. C. Chang, and R. Perdisci. Https: Sealing information leaks with browser-side obfuscation of encrypted flows. In *NDSS '11*.
65. A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian. L-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1):3, 2007.
66. Stephen McLaughlin, Patrick McDaniel, and William Aiello. Protecting consumer privacy from electric load monitoring. In *CCS '11*, pages 87–98, 2011.
67. A. Meyerson and R. Williams. On the complexity of optimal k-anonymity. In *ACM PODS*, pages 223–228, 2004.
68. G. Miklau and D. Suciu. A formal analysis of information disclosure in data exchange. In *SIGMOD*, pages 575–586, 2004.
69. Andrés Molina-Markham, Prashant Shenoy, Kevin Fu, Emmanuel Cecchet, and David Irwin. Private memoirs of a smart meter. In *BuildSys '10*, pages 61–66, 2010.
70. S. Nagaraja, V. Jalaparti, M. Caesar, and N. Borisov. P3ca: private anomaly detection across isp networks. In *PETS'11*, pages 38–56, 2011.
71. Arvind Narayanan and Vitaly Shmatikov. De-anonymizing social networks. In *IEEE Symposium on Security and Privacy '09*, pages 173–187, 2009.
72. Alfredo Rial and George Danezis. Privacy-preserving smart metering. In *WPES '11*, pages 49–60, 2011.
73. Ian Richardson, Murray Thomson, David Infield, and Conor Clifford. Domestic electricity use: A high-resolution energy demand model. *Energy and Buildings*, 42(10):1878–1887, 2010.
74. T. Ristenpart, E. Tromer, H. Shacham, and S. Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *CCS*, pages 199–212, 2009.
75. Ishtiaq Rouf, Hossen Mustafa, Miao Xu, Wenyan Xu, Rob Miller, and Marco Gruteser. Neighborhood watch: security and privacy analysis of automatic meter reading systems. In *CCS '12*, pages 462–473, 2012.
76. P. Samarati. Protecting respondents' identities in microdata release. *IEEE Trans. on Knowl. and Data Eng.*, 13(6):1010–1027, 2001.
77. T. S. Saponas and S. Agarwal. Devices that tell on you: Privacy trends in consumer ubiquitous computing. In *USENIX '07*, pages 5:1–5:16, 2007.
78. J. Schlörer. Identification and retrieval of personal records from a statistical bank. In *Methods Info. Med.*, pages 7–13, 1975.
79. A. Slavkovic and S.E. Feinberg. Bounds for cell entries in two-way tables given conditional relative frequencies. *Privacy in Statistical Databases*, 2004.
80. J. Sun, X. Zhu, C. Zhang, and Y. Fang. Hcnp: Cryptography based secure ehr system for patient privacy and emergency healthcare. In *ICDCS'11*, pages 373–382, 2011.
81. Q. Sun, D. R. Simon, Y. M. Wang, W. Russell, V. N. Padmanabhan, and L. Qiu. Statistical identification of encrypted web browsing traffic. In *IEEE Symposium on Security and Privacy '02*, pages 19–, 2002.
82. C. Wang, N. Cao, J. Li, K. Ren, and W. Lou. Secure ranked keyword search over encrypted cloud data. In *ICDCS'10*, pages 253–262, 2010.
83. Wenye Wang and Zhuo Lu. Cyber security in the smart grid: Survey and challenges. *Computer Networks*, 57(5):1344–1371, 2013.

84. R. C. Wong and A. W. Fu. *Privacy-Preserving Data Publishing: An Overview*. Morgan and Claypool Publishers, 2010.
85. R.C. Wong, A.W. Fu, K. Wang, and J. Pei. Minimality attack in privacy preserving data publishing. In *VLDB*, pages 543–554, 2007.
86. R.C. Wong, J. Li, A. Fu, and K. Wang. alpha-k-anonymity: An enhanced k-anonymity model for privacy-preserving data publishing. In *KDD*, pages 754–759, 2006.
87. C. V. Wright, S. E. Coull, and F. Monrose. Traffic morphing: An efficient defense against statistical traffic analysis. In *NDSS '09*.
88. X. Xiao and Y. Tao. Personalized privacy preservation. In *SIGMOD*, pages 229–240, 2006.
89. X. Xiao, Y. Tao, and N. Koudas. Transparent anonymization: Thwarting adversaries who know the algorithm. *ACM Trans. Database Syst.*, 35(2):1–48, 2010.
90. X. Xiao, G. Wang, and J. Gehrke. Differential privacy via wavelet transforms. In *ICDE '10*, pages 225–236, 2010.
91. Xiaokui Xiao and Yufei Tao. M-invariance: towards privacy preserving re-publication of dynamic datasets. In *SIGMOD '07*, pages 689–700, 2007.
92. Weining Yang, Ninghui Li, Yuan Qi, Wahbeh Qardaji, Stephen McLaughlin, and Patrick McDaniel. Minimizing private data disclosures in the smart grid. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security, CCS '12*, pages 415–427, 2012.
93. K. Zhang, Z. Li, R. Wang, X. Wang, and S. Chen. Sidebuster: automated detection and quantification of side-channel leaks in web application development. In *CCS '10*, pages 595–606, 2010.
94. L. Zhang, S. Jajodia, and A. Brodsky. Information disclosure under realistic assumptions: privacy versus optimality. In *CCS*, pages 573–583, 2007.
95. L. Zhang, L. Wang, S. Jajodia, and A. Brodsky. Exclusive strategy for generalization algorithms in micro-data disclosure. In *Data and Applications Security XXII*, volume 5094 of *Lecture Notes in Computer Science*, pages 190–204. 2008.
96. L. Zhang, L. Wang, S. Jajodia, and A. Brodsky. L-cover: Preserving diversity by anonymity. In *SDM '09*, pages 158–171, 2009.
97. Y. Zhang, A. Juels, A. Oprea, and M. K. Reiter. Homealone: Co-residency detection in the cloud via side-channel analysis. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, pages 313–328, 2011.
98. Li Zhuang, Feng Zhou, and J. D. Tygar. Keyboard acoustic emanations revisited. *ACM Trans. Inf. Syst. Secur.*, 13(1):3:1–3:26, November 2009.

Preserving Privacy Against Side-Channel Leaks
From Data Publishing to Web Applications

Liu, W.M.; Wang, L.

2016, XIII, 142 p. 19 illus., 1 illus. in color., Hardcover

ISBN: 978-3-319-42642-6