

Preface¹

With rapid advancements in information technology, today's organizations routinely collect, store, analyze, and redistribute vast amounts of data about individuals, such as user account information and online activities. In addition, the next generation of smart systems (e.g., smart grids and smart medical devices) will enable organizations to collect personal data about every aspect of our daily life, from real-time power consumption to medical conditions.

Although collecting data may be essential for organizations to conduct their business, indiscriminate collection, retention, and dissemination of personal data represents a serious intrusion to the privacy of individuals. As a fundamental right of all individuals, privacy protection means organizations should only collect and retain sensitive personal information for purposes that have been agreed upon by the individuals and also keep collected information confidential and accessible only to authorized personnel.

Unfortunately, protecting personal information poses serious technical challenges in almost every stage of the data management life cycle, from data collection to data dissemination. A particularly insidious threat in this context is the *side-channel leak* in which an adversary makes inference of confidential data based on some seemingly innocent characteristics of the data, such as data packet sizes or knowledge about public algorithms used to generate the data. While side-channel attacks in specific domains, such as cryptosystems, are well studied, there exist little effort on generalizing side-channel attacks across different domains in order to understand their commonality.

This book studies side-channel leaks and corresponding countermeasures in several domains. First, we focus on privacy-preserving data publishing (PPDP) where side-channel leaks may be caused by adversaries' knowledge about the algorithms used to anonymize the data. For countermeasures, we first study a generic strategy independent of data utility measures and syntactic privacy properties, and then

¹Wen Ming Liu's work on this book was completed during his time as a Ph.D. student at Concordia University.

we propose a more efficient approach by decoupling privacy protection and utility optimization. Second, we examine Web applications where side-channel leaks may be caused by packet sizes and timing. For countermeasures, we first study a privacy-preserving traffic padding method inspired by the aforementioned PPDP solution, and then we further strengthen the approach against adversaries' external knowledge through random padding. Third, we look at smart metering where side-channel leaks may be caused by fine-grained meter readings. Finally, we discuss how those specific instances of side-channel leaks may be modeled using a generic model.

This book provides readers with not only detailed analysis of side-channel leaks and their solutions in each of the aforementioned domains but also a generic model that bridges the gaps between those different threats and solutions. The benefit of such knowledge is twofold. First, it provides readers with sufficient technical background to understand the threat of side-channel leaks in those domains and consequently exposes readers to many challenging and important issues that still remain attractive research topics today. Second, it can also lead readers to look beyond those three domains and apply the insights and ideas to derive novel solutions for dealing with side-channel leaks in other practical applications.

Montreal, QC, Canada

Lingyu Wang

<http://www.springer.com/978-3-319-42642-6>

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