

Preface

In October of 2006, as part of an NSF grant to design the future Internet, the authors got involved in designing a naming and addressing architecture for the future Internet. The Internet's naming and addressing architecture has widely accepted limitations. Its dated and host-centric Domain Name System (DNS) fails to keep up with the current and emerging usage models, and the overloading of the IP address hinders both mobility and security. As we started sifting through the large body of architectural schemes and proposals, two concerns quickly became evident. First, the independent nature of the many contributions to the field has led to the overloading of architectural terms, and to the emergence of a large body of network architecture proposals with no clear understanding of their cross similarities, and their unique properties. Second, there seems to be a growing consensus in the community on the need for designing a “smarter” network. While such evolution is bringing new potentials and service models, the community generally lacks consistent frameworks for thinking about such models and their design implications.

Purpose and Scope This book is a timely attempt at developing a unified model for studying internetwork naming and discovery. Naming is an integral building block within data networks and systems and is becoming even more so as complex data-centric usage models emerge. Several books have been devoted to the broad and general topics of computer networks [104, 154], and distributed systems [53]. We specifically focus on the naming, addressing, and discovery services. We develop the architectural and economic aspects of designing naming and discovery schemes within the broader context of internetwork architecture.

Organization The balance of this book is divided into three parts. Part I aims at framing the naming and discovery services within the broader context of the internetwork architecture, and serves as a reference for the rest of the book. Chapter 1 refreshes the reader on the main events, technologies, and design principles that led to the Internet we know today, and discusses why the very original design that enabled decades of unprecedented evolution and innovation on the Internet is falling short of satisfying the user requirements after three decades of its inception. Global

clean-slate initiatives, such as NSF's FIND [10] and FIA [9], are under way to design the Internet of the future. Chapter 2 surveys some of these prominent proposals highlighting their naming aspects.

Part II presents the formal models for naming and discovery. In terms of naming, Chap. 3 disambiguates the plethora of terms used in the literature—name, address, identifier, locator, binding, routing, discovery, mapping, and resolution. We provide a formalism based on simple relational algebra, we examine the main properties of names and bindings, and we present two illustrative case studies. Chapter 4 presents a simple taxonomy that serves three main purposes. The taxonomy is a tutorial for introduction to the field, it helps in organizing and thinking about the solution space, and more importantly in identifying new avenues for contributing to the topic. In terms of discovery, Chap. 5 presents a general model for exploring the tradeoffs involved in designing Internet name-based discovery systems. Designing a discovery system is commonly challenged with providing three desirable properties: scalability, efficiency, and adaptability. The model highlights these tradeoffs across a wide range of networks. Chapter 6 is an illustrative design exercise. We present an architecture for differentiated discovery services.

Part III develops some of the economic and incentive problems common to naming and discovery services in large scale networks. Chapter 7 presents a broad treatment of the main economic issues that arise in the context of name-based discovery. Providing a discovery service while accounting for the cost and making sure that the incentives of the players are aligned is the general theme of the chapter. We motivate the subject, present a taxonomy of discovery schemes and proposals based on their business model, and we pose several questions that are becoming increasingly important as we proceed to design the internetwork of the future. This sets the stage for Chap. 8 which presents an incentive model for route distribution in the context of path vector routing protocols (mainly BGP). We model BGP route distribution and computation using a game in which a BGP speaker advertises its prefix to its direct neighbors promising them a reward for further distributing the route deeper into the network. The neighbors do the same thing with their neighbors, and so on. The result of this cascaded route distribution is an advertised prefix and hence reachability of the BGP speaker. We first study the convergence of BGP protocol dynamics to a unique outcome tree in the defined game. We then study the existence of equilibria in the full information game considering competition dynamics.

To keep the content simple and intuitive to the reader, we develop the concepts progressively and we provide examples whenever possible. The book has a logical layout to it and we recommend reading the parts sequentially as each part sets the stage for the next. For more advanced networking readers, however, we kept each chapter self-contained to facilitate skipping to particular chapters.

Audience This book is directed at students and professionals seeking a training in advanced networking topics. We believe that it may be specifically appealing for graduate students and researchers specializing in the areas of Internet architecture, naming and discovery, routing, and economics.

Acknowledgements The authors acknowledge the colleagues and the reviewers for their valuable comments, and they gratefully acknowledge the National Science Foundation (NSF) for supporting the early work investigating the field. The authors would like to emphasize that they have done their best to correct any errors that were brought to their attention. Any remaining errors are the responsibility of the authors only. Should you locate any, please email the authors at joudkhoury@gmail.com.

Boston

Joud Khoury
Chaouki T. Abdallah



<http://www.springer.com/978-1-4471-4551-6>

Internet Naming and Discovery

Architecture and Economics

Khoury, J.S.; Abdallah, C.T.

2013, XIV, 142 p., Hardcover

ISBN: 978-1-4471-4551-6