University of Glasgow

Naming and the Tussle for Control

Networked Systems (H)

Lecture 8

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow | This work is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License.
To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/4.0/ or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Lecture Outline

- What is the DNS and DNS resolution
- DNS naming structure and organisation
- Methods for DNS resolution
- The politics of names

- Associated reading: “Tussle in cyberspace: Defining tomorrow’s Internet” discusses the issues of control over the network, and how protocol design influences this – the DNS is one area where we see this tussle



Tussle in Cyberspace: Defining Tomorrow’s Internet

David D. Clark
MIT Lab for Computer Science
ddc@lcs.mit.edu

John Wroclawski
MIT Lab for Computer Science
jtw@lcs.mit.edu

Karen R. Sollins
MIT Lab for Computer Science
sollins@lcs.mit.edu

Robert Braden
USC Information Sciences Institute
braden@isi.edu

Abstract

The architecture of the Internet is based on a number of principles, including the self-describing datagram packet, the end to end arguments, diversity in technology and global addressing. As the Internet has moved from a research curiosity to a recognized component of mainstream society, new requirements have emerged that suggest new design principles, and perhaps suggest that we revisit some old ones. This paper explores one important reality that surrounds the Internet today: different stakeholders that are part of the Internet milieu have interests that may be adverse to each other, and these parties each vie to favor their particular interests. We call this process “the tussle”. Our position is that accommodating this tussle is crucial to the evolution of the network’s technical architecture. We discuss some examples of tussle, and offer some technical design principles that take it into account.

Categories and Subject Descriptors

C.2.1 [**Computer Systems Organization**]: Computer Communications Networks—*Network Architecture and Design*; H.1 [**Information Systems**]: Models and Principles; K.4.1 [**Computing Milieux**]: Computers and Society—*Public Policy Issues*

General Terms

Design, Economics, Legal Aspects, Security, Standardization

Keywords

Tussle, Network Architecture, Trust, Economics, Design Principles, Competition

Work sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-00-2-0553 at MIT, and agreement number F30602-00-1-0540 at ISI. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SIGCOMM’02, August 19-23, 2002, Pittsburgh, Pennsylvania, USA.
Copyright 2002 ACM 1-58113-570-X/02/0008 ...$5.00.

1. INTRODUCTION

The Internet was created in simpler times. Its creators and early users shared a common goal—they wanted to build a network infrastructure to hook all the computers in the world together so that as yet unknown applications could be invented to run there. All the players, whether designers, users or operators, shared a consistent vision and a common sense of purpose.

Perhaps the most important consequence of the Internet’s success is that the common purpose that launched and nurtured it no longer prevails. There are, and have been for some time, important and powerful players that make up the Internet milieu with interests directly at odds with each other.

Some examples are very current. Music lovers of a certain bent want to exchange recordings with each other, but the rights holders want to stop them. People want to talk in private, and the government wants to tap their conversations. Some examples are so obvious that they are almost overlooked. For the Internet to provide universal interconnection, ISPs must interconnect, but ISPs are sometimes fierce competitors. It is not at all clear what interests are being served, to whose advantage, to what degree, when ISPs negotiate terms of connection. It is not a single happy family of people dedicated to universal packet carriage.

We suggest that this development imposes new requirements on the Internet’s technical architecture. These new requirements, in turn, motivate new design strategies to accommodate the growing tussle among and between different Internet players. The purpose of this paper is to explore what these requirements and strategies might be.

We begin by briefly discussing the Internet landscape - some fundamental differences between the mechanisms of engineering and society, and the players that populate our field. We then outline some proposed design principles intended to accommodate within the Internet mechanisms of society as well as those of engineering. We believe this accommodation is central to designing an Internet that is resilient to the challenges of society as well as those of technology. We conclude by discussing some tussle spaces, ways in which our principles might guide the technical response to these spaces, and specific technical research that may be of value in accommodating these tussles.

1.1 The natures of engineering and society

Engineers attempt to solve problems by designing mech-

D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, “Tussle in cyberspace: Defining tomorrow’s Internet”, ACM SIGCOMM Conference, August 2002. https://dx.doi.org/10.1145/633025.633059

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

DNS Name Resolution

- What is the DNS?
- Structure of names
- Name resolution

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

What is the DNS? (1/2)

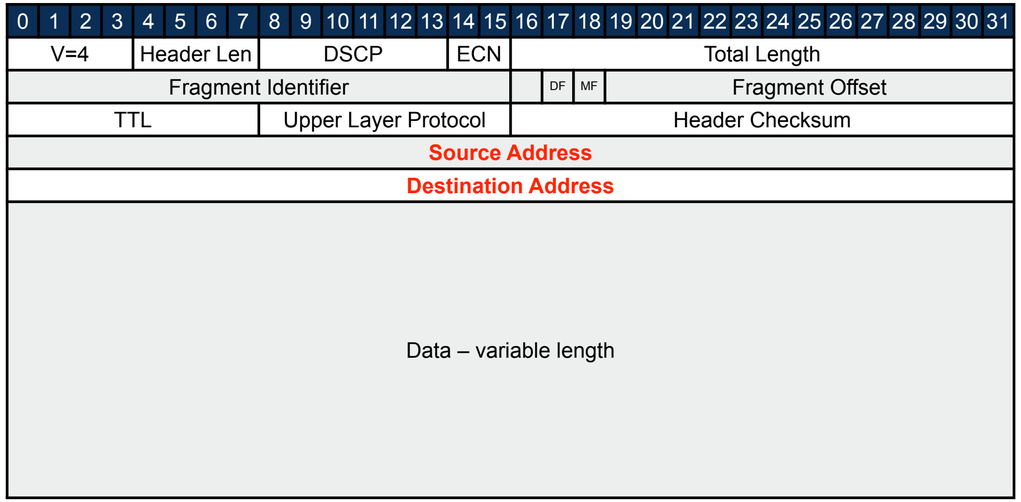

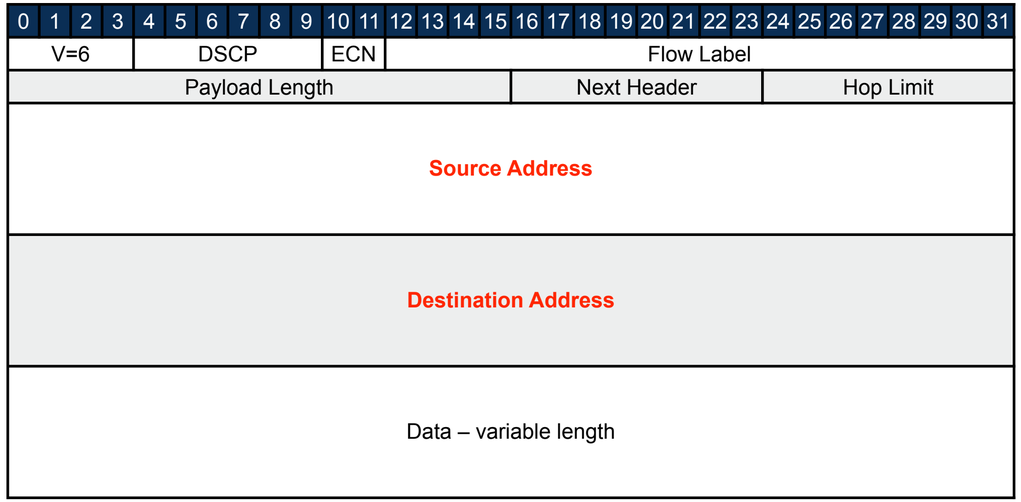

- IP packets contain addresses rather than names
  - Designed for efficient processing by routers determining where to forward the packet
  - Not human readable – people prefer names, not addresses
- Domain name system (DNS) is a distributed database; maps names to IP addresses

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

What is the DNS? (2/2)

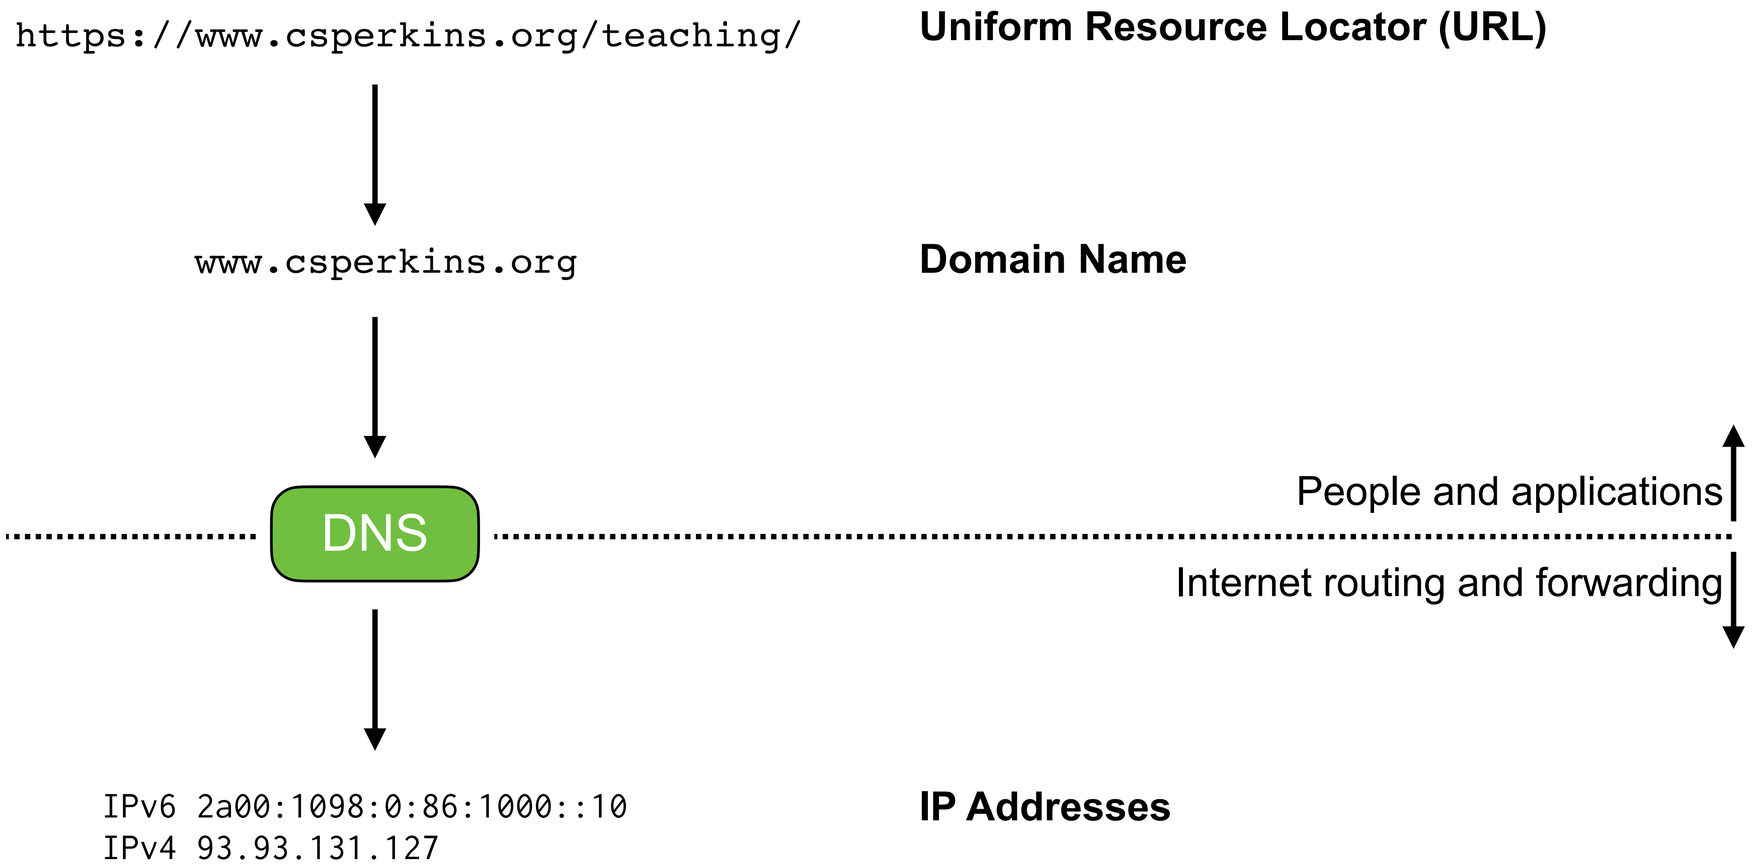

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Structure of DNS Names

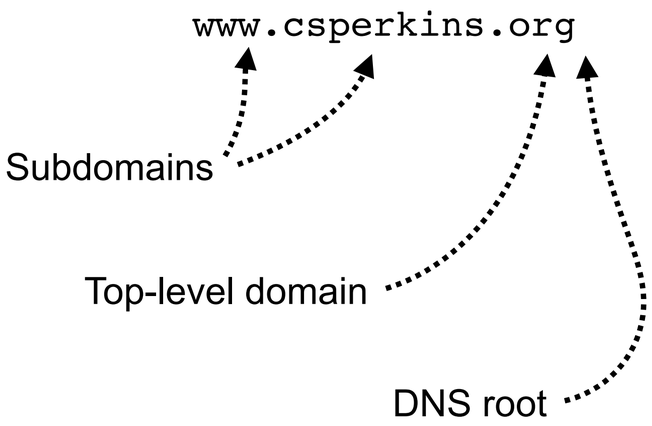

- Naming is hierarchical
  - Sub-domains first
  - Concludes with a **top-level domain** (TLD)
    - Country-code top-level domains (ccTLDs)
      - `.uk`, `.de`, `.cn`, `.io`, `.ly`, …
    - Generic top-level domains (gTLDs)
      - `.com`, `.org`, `.net`, …
  - Top-level domains live within the **DNS root**
    - The **root servers** advertise the top-level domains
    - They have well-known, fixed, IP addresses – new DNS resolvers need to reach them to find the TLDs before they can answer DNS queries
- Each level is independently administered and operated
  - DNS is a distributed database – in authority and implementation
  - Each level in the hierarchy controls its own data

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

DNS Name Resolution (1/3)

- The DNS is used for **name resolution**
  - Given a name, lookup a particular type of **record** relating to that name
  - Many different types of record: `A`, `AAAA`, `CNAME`, `MX`, `NS`, `SRV`, …
  - Most common are `A` and `AAAA` records, that map hostnames to IPv4 and IPv6 addresses, and `NS` records that identify the **name server** for a domain
- A DNS **client** asks its **resolver** to perform the lookup by calling `getaddrinfo()`
  - The resolver could be a process running on the client, it more commonly runs on a machine provided by the network operator

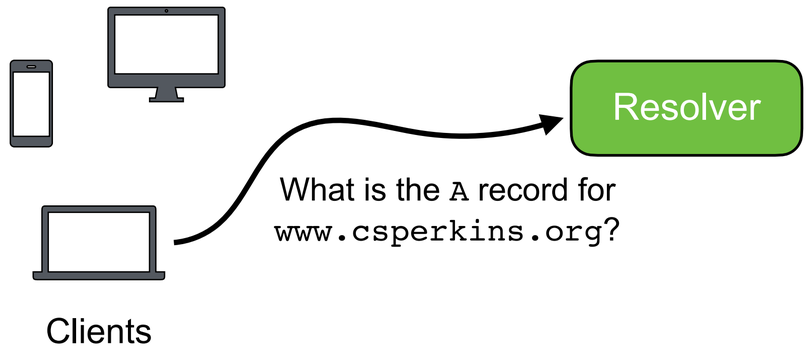

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

DNS Name Resolution (2/3)

- If the resolver has no information, it makes a recursive query via the DNS root servers
  - Query the root to find the TLD
  - Query the TLD to find the subdomain
    - Repeat as necessary
  - Query the subdomain to find the address

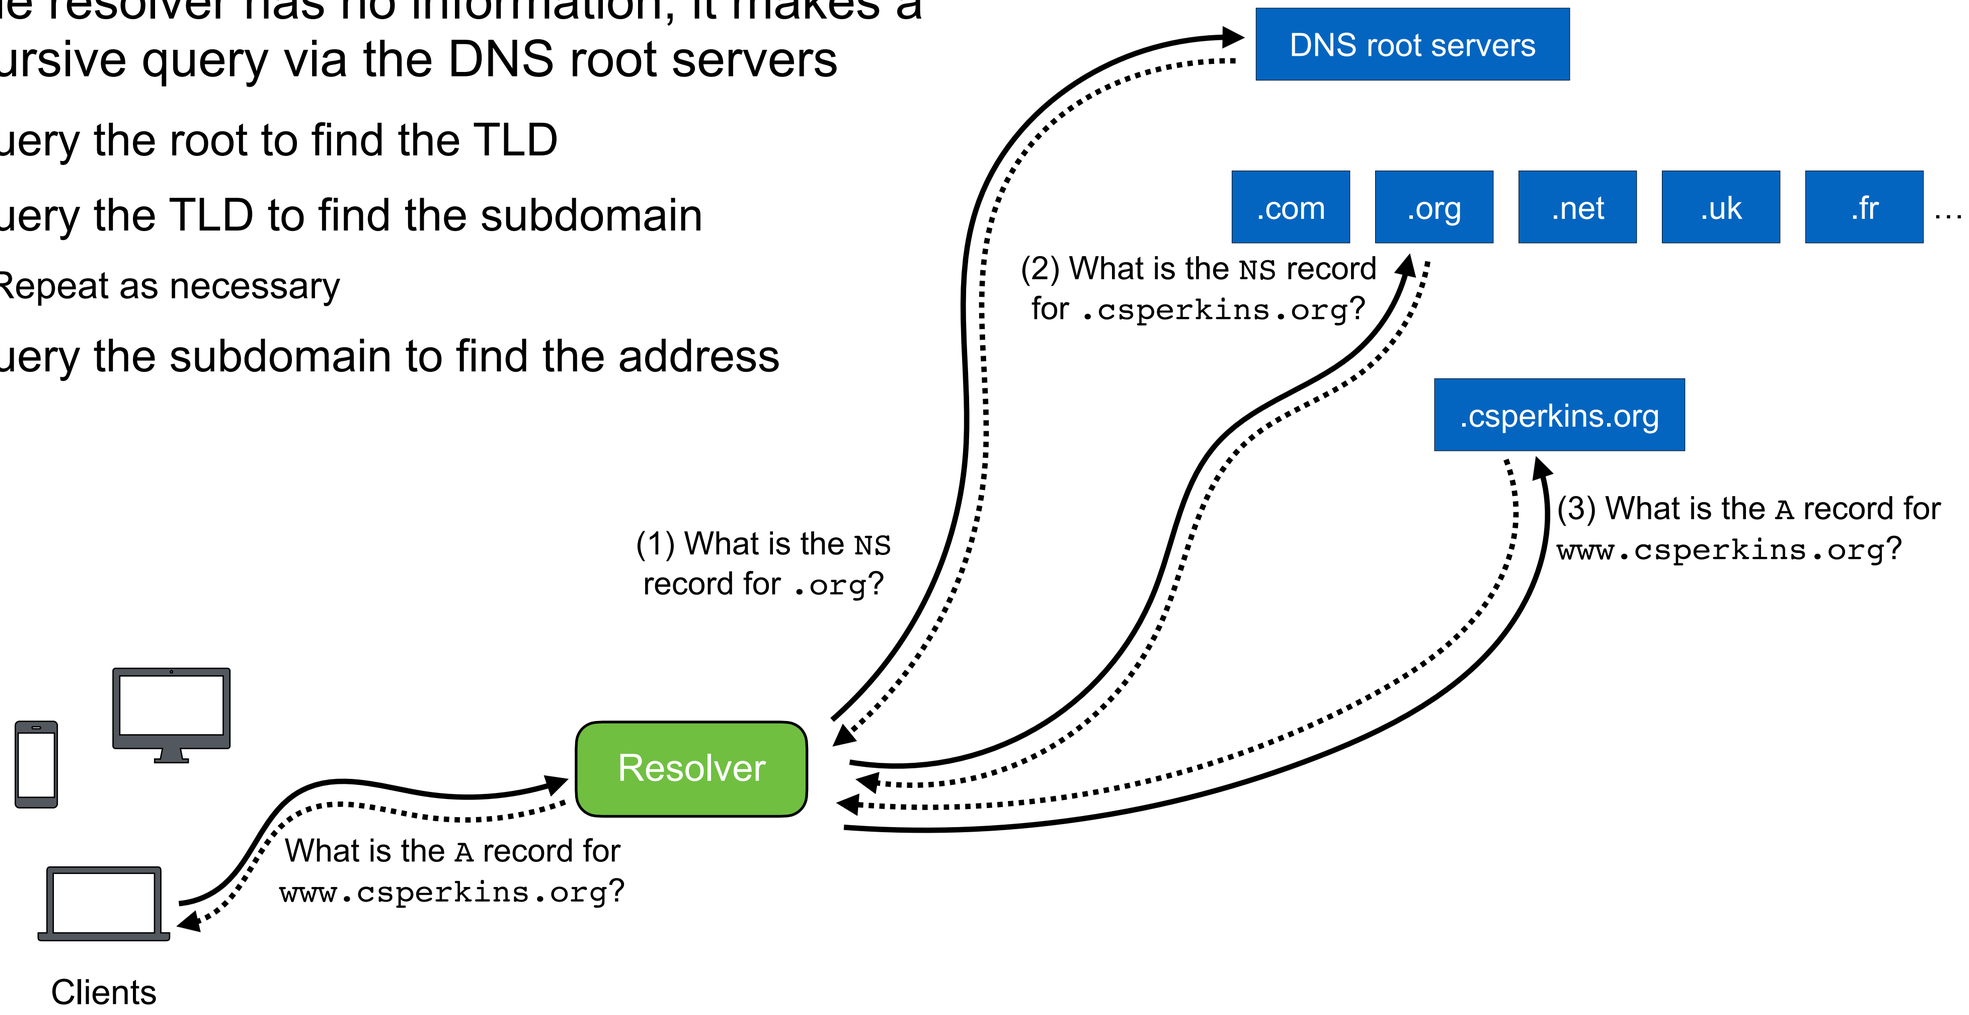

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

DNS Name Resolution (3/3)

- Responses include a **time to live** that allows a resolver to cache the value for a certain time period
  - Subdomains can set a short TTL and give different answers each time a particular name is requested → load balancing; CDNs directing queries to local caches
- Subsequent queries are answered from the cache, where possible
  - If the cached entry times it, it's refreshed from the next level up in the hierarchy
  - Eventually reaching the DNS root
- The IP addresses for the root servers are well known, and never time out

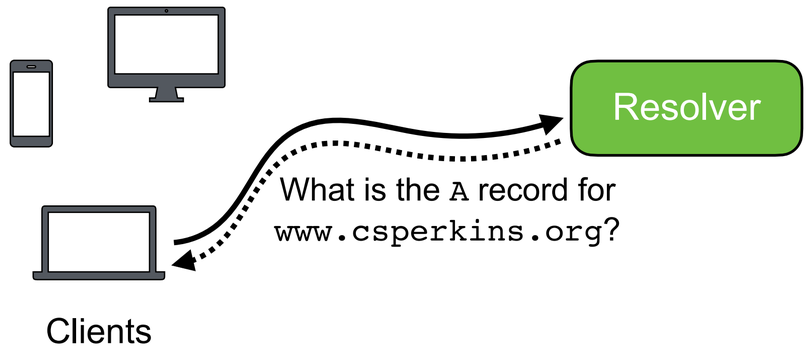

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

DNS Name Resolution

- What is the DNS?
- Structure of names
- Name resolution

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow