

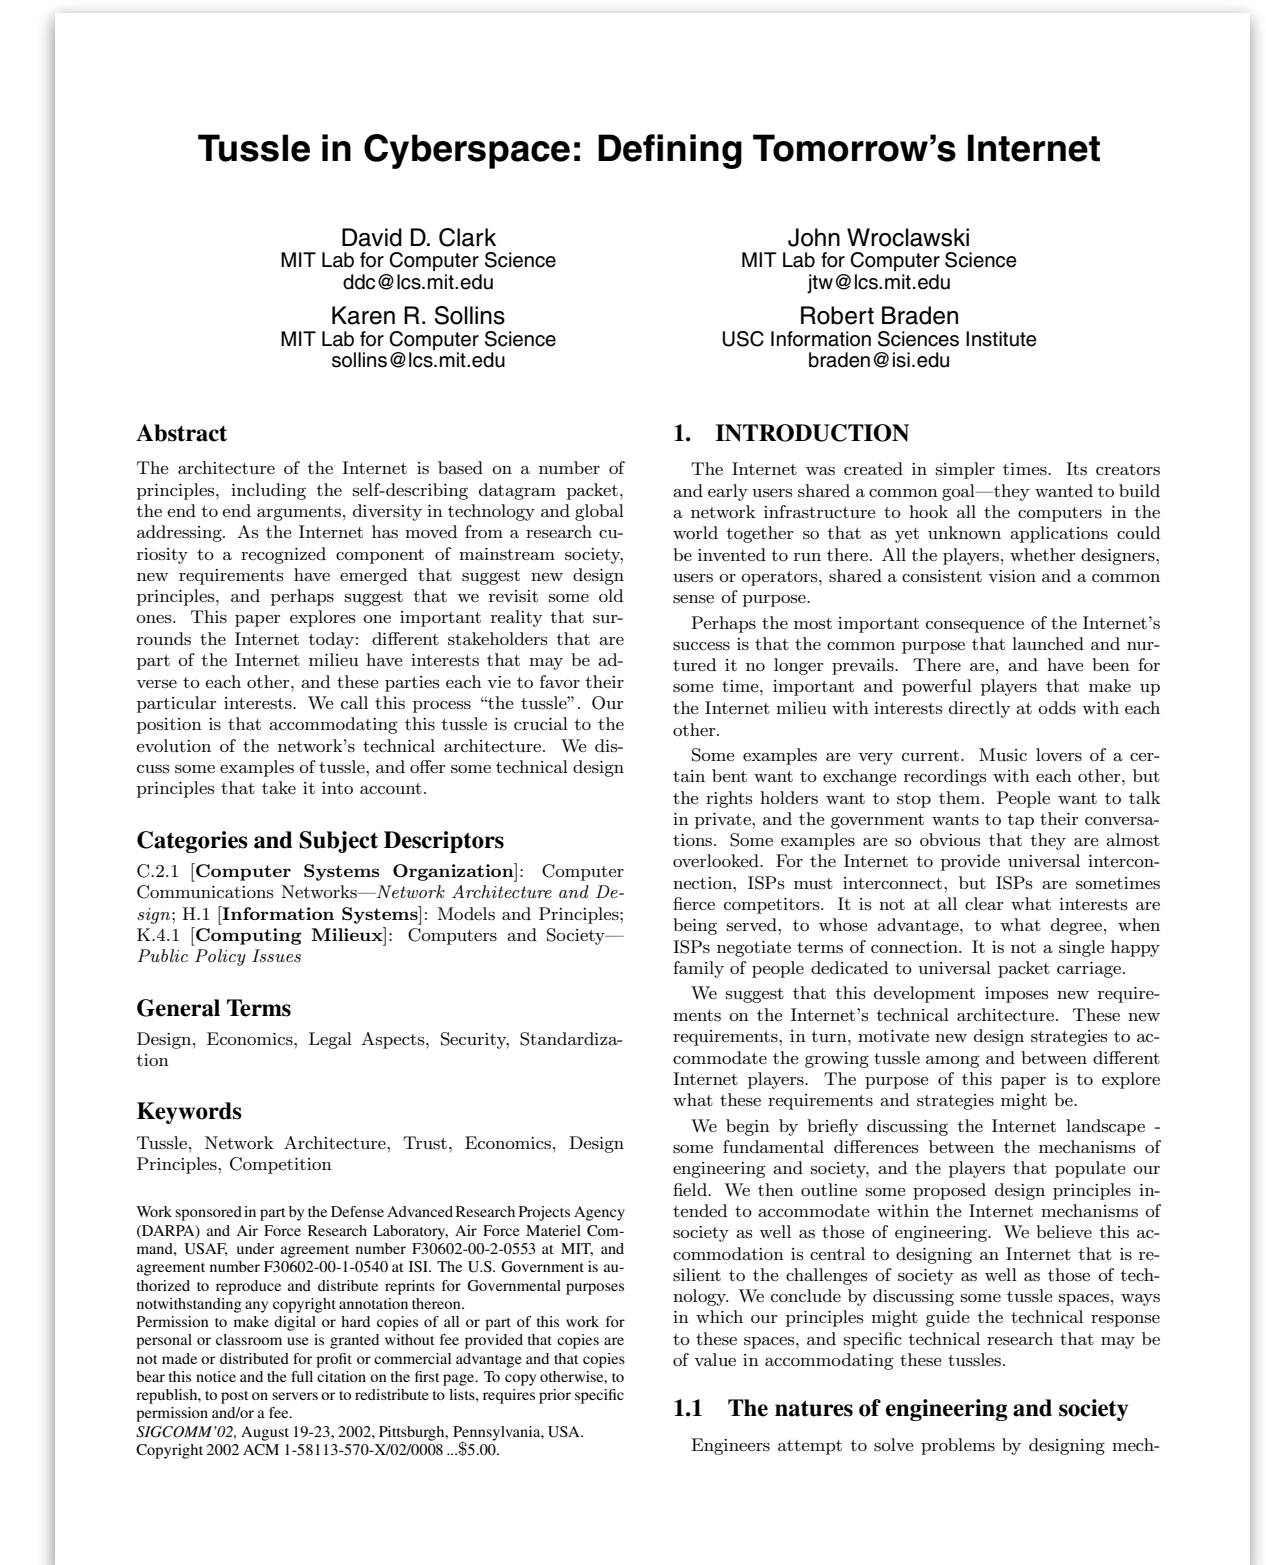
# Naming and the Tussle for Control

Networked Systems (H)

Lecture 8

# 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



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>

# DNS Name Resolution

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

# What is the DNS? (1/2)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
V=4		Header Len		DSCP			ECN		Total Length																						
Fragment Identifier										DF	MF	Fragment Offset																			
TTL			Upper Layer Protocol				Header Checksum																								
<b>Source Address</b>																															
<b>Destination Address</b>																															
Data – variable length																															

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
V=6		DSCP			ECN		Flow Label																								
Payload Length										Next Header					Hop Limit																
<b>Source Address</b>																															
<b>Destination Address</b>																															
Data – variable length																															

- 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

# What is the DNS? (2/2)

`https://www.csperkins.org/teaching/`

**Uniform Resource Locator (URL)**



`www.csperkins.org`

**Domain Name**



People and applications

Internet routing and forwarding

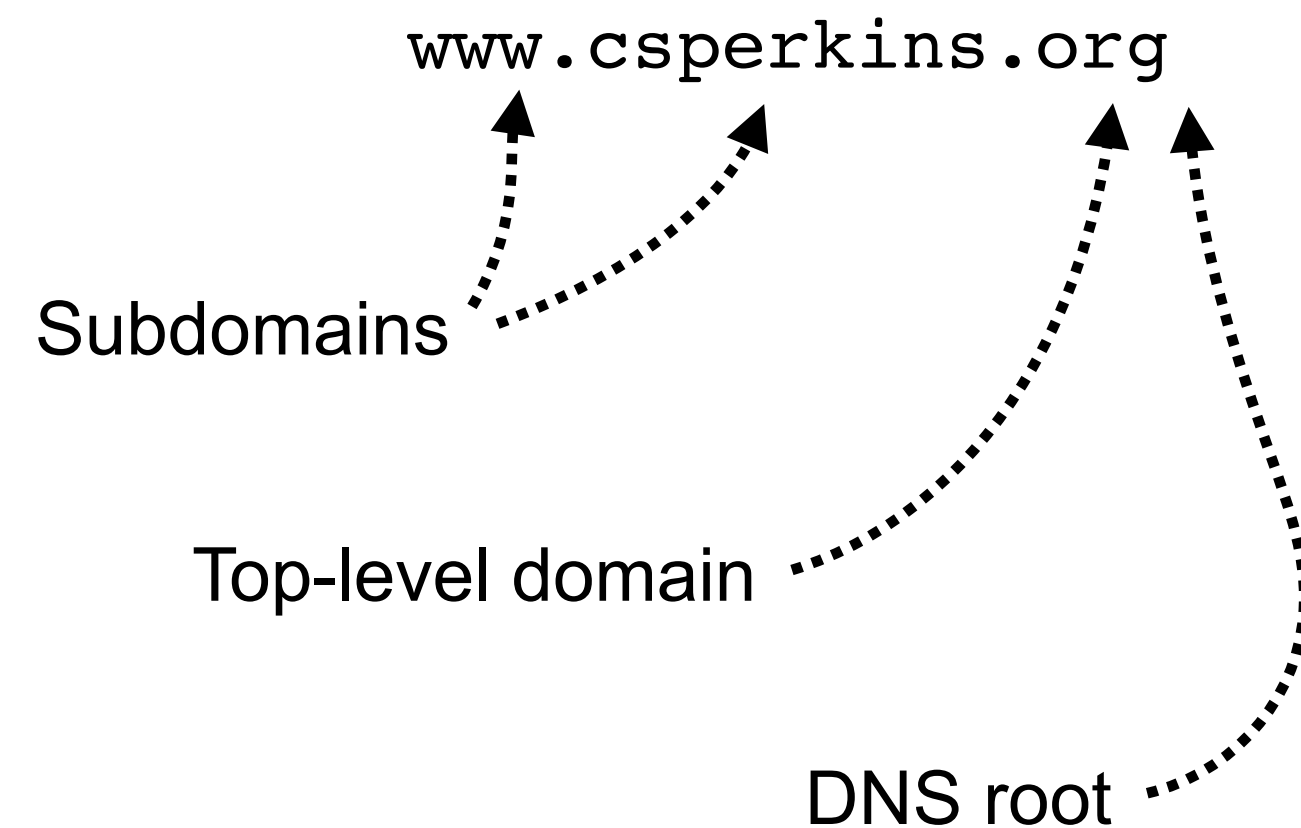


IPv6 `2a00:1098:0:86:1000::10`

IPv4 `93.93.131.127`

**IP Addresses**

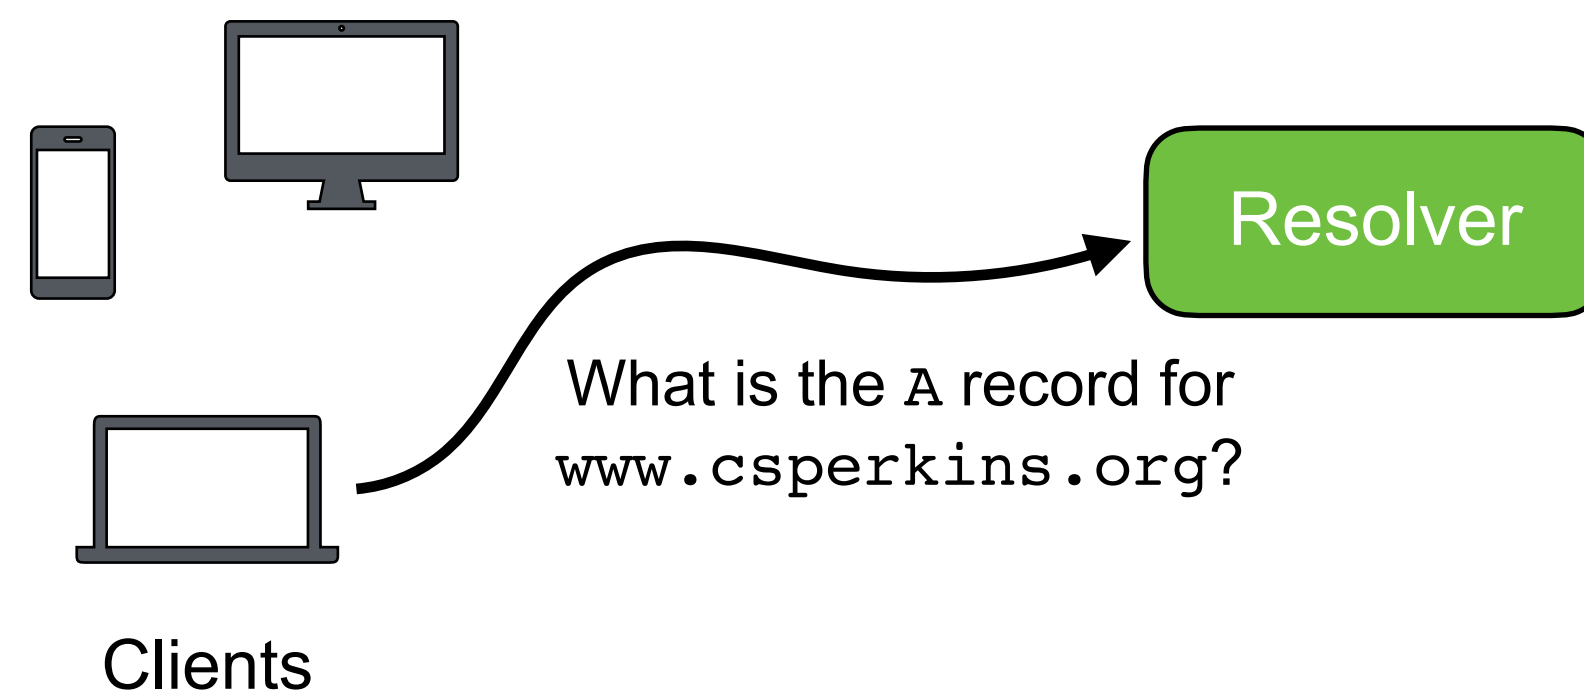
# Structure of DNS Names



- 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

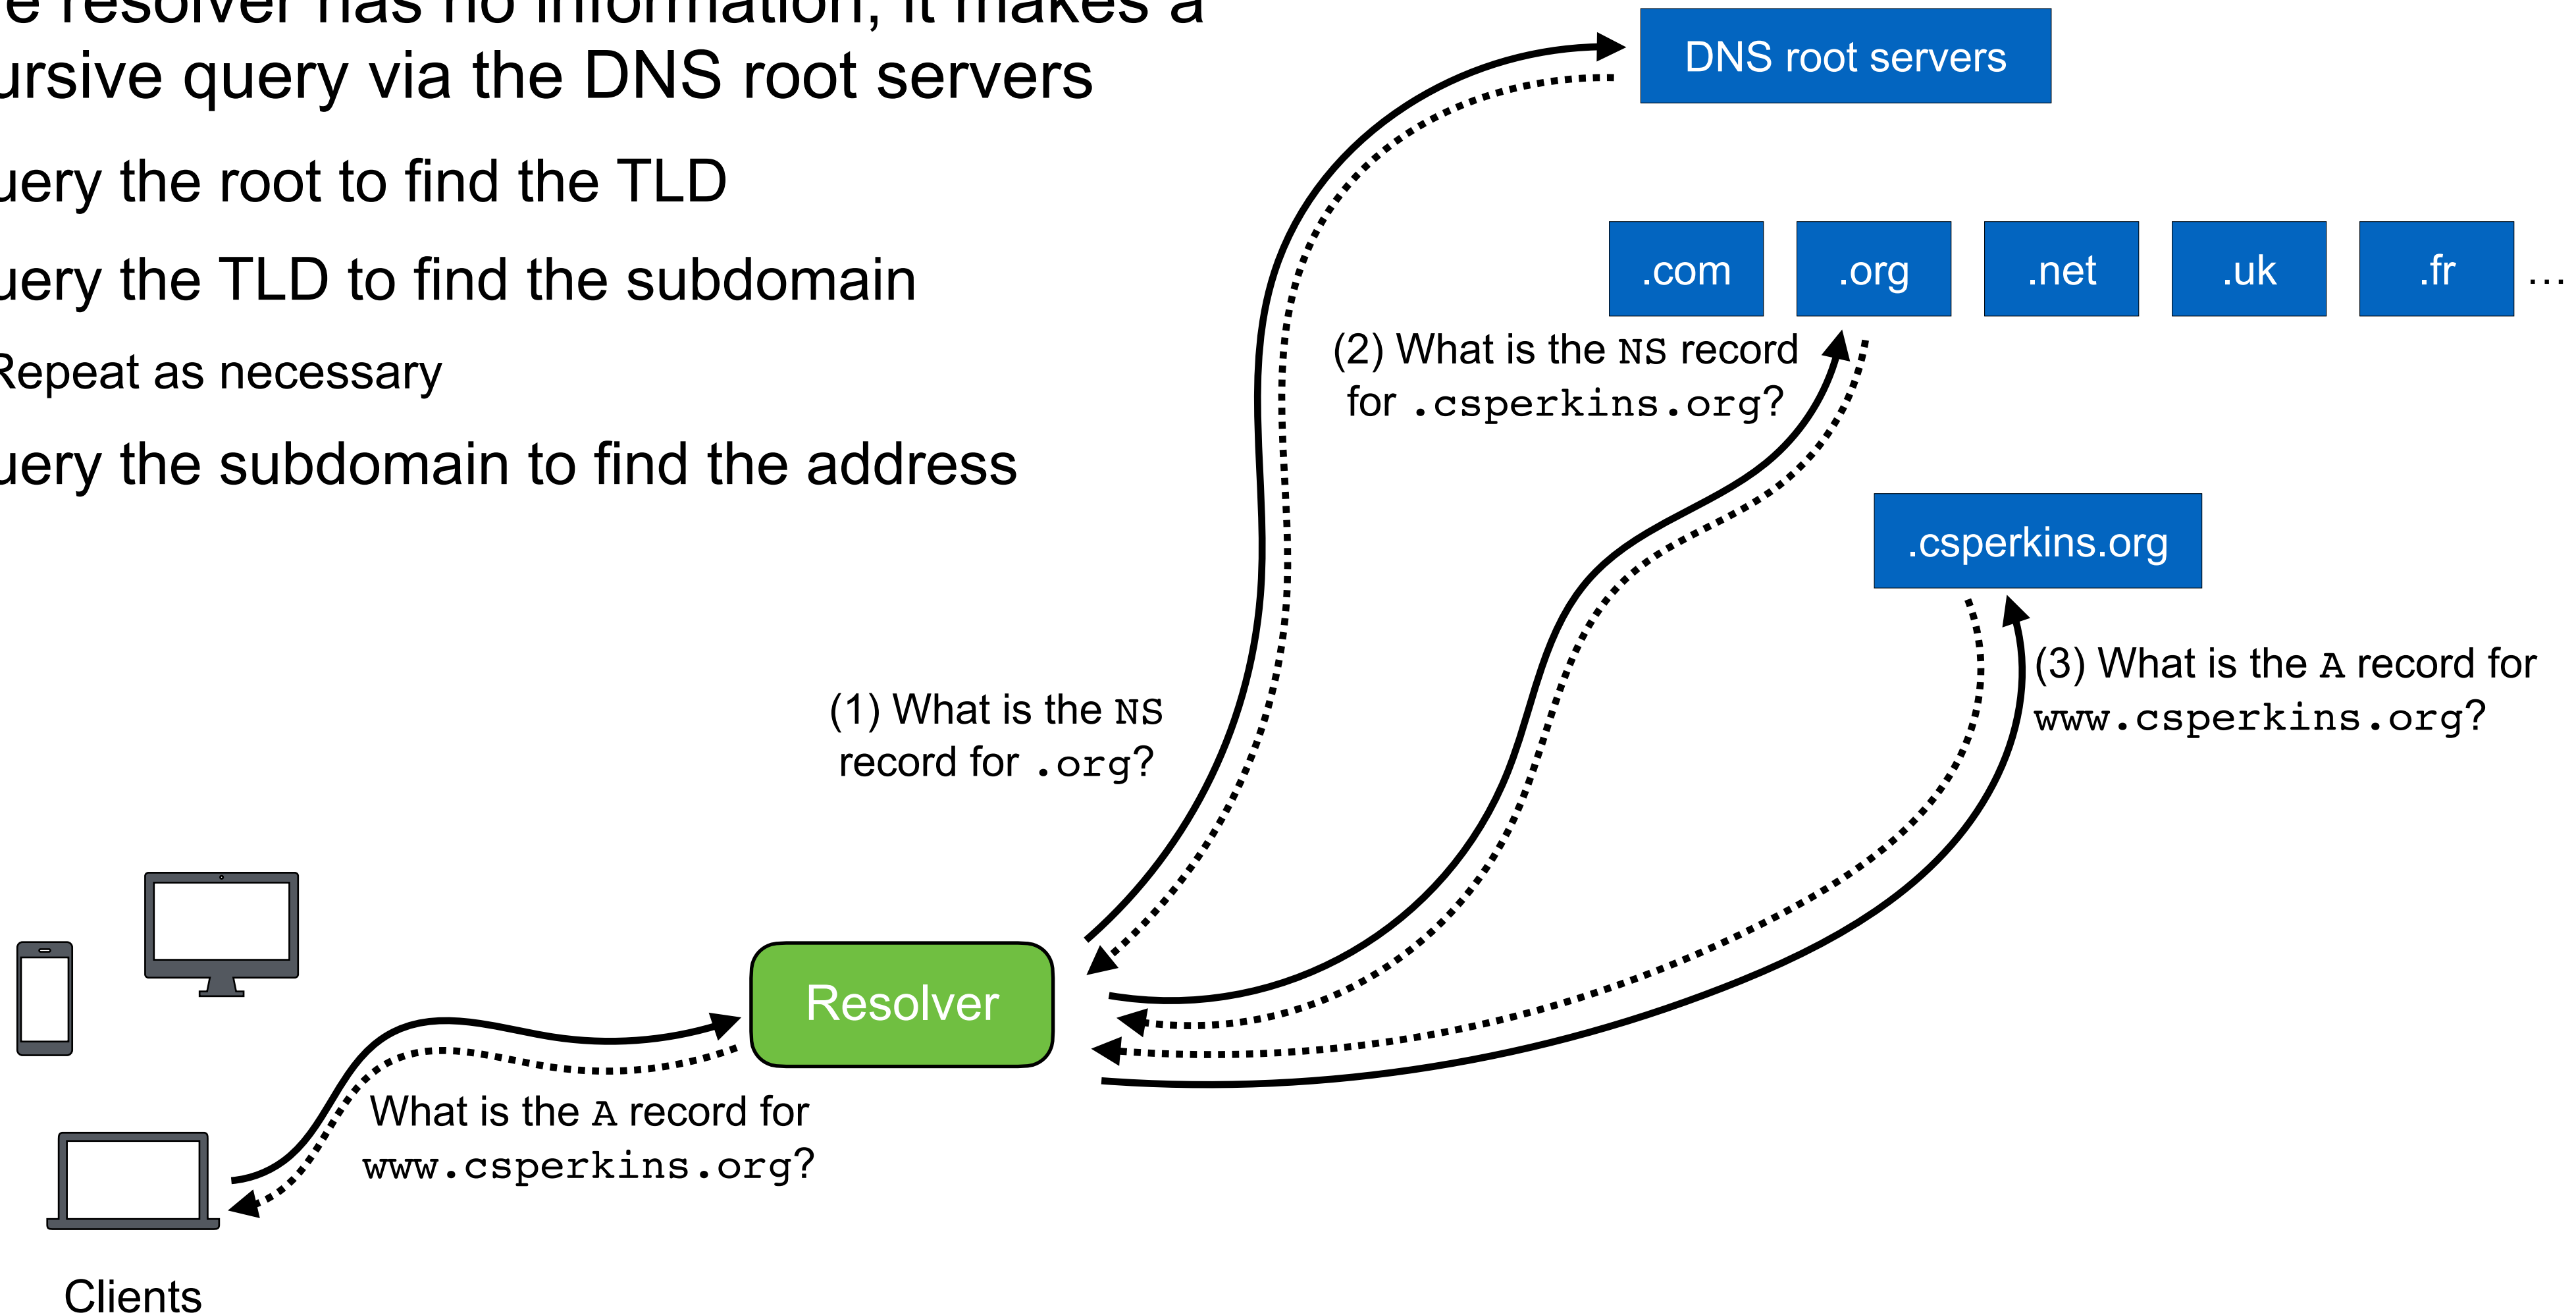
# DNS Name Resolution (1/3)

- 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



# 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





# 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



# DNS Name Resolution

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