

Overview of NDN Security and Privacy Landscape

Gene Tsudik

UCI

Joint work with:

Paolo Gasti (NYIT/UCI), Mauro Conti (Padova), Ersin Uzun (PARC), Jeff Burke (UCLA),
Ashok Narayan (CISCO), Dave Oran (CISCO), Naveen Nathan (UCI), Cesar Ghali (UCI) et al.

Outline

1. Named Data Networking (NDN)
2. Security and privacy Issues in NDN
3. A few topics in more detail

Communication

- For almost 150 years, communication meant:

A wire connecting two devices



- The Web forever changed that:

What matters is content, not the host it came from




Paradigm Shift



1876



2012



Today's Internet

- Tremendous global success story
- Architecture defined in RFC 791/793 (1981)
- Enables any host to talk to any other host
 - Names boxes and interfaces
 - Supports end-to-end conversation
 - Provides unreliable packet delivery via IP datagrams
 - Compensates via complexity of TCP



Inter-networking was originally designed for this world

It was about sharing resources, not data.

Today's Internet

- Helped facilitate today's "world of content" but was never designed for it
- Fundamental communication model: point-to-point conversation between two hosts (IP interfaces)
- The central abstraction is a host identifier corresponding to an IP address



Today's Internet

- Last 15 years – profound change in nature of Internet communication
 - From email/ftp/telnet → → content, content and more content
- Massive amounts of data produced and consumed constantly
 - Web (esp. media sharing and social networking), audio-/video-conferencing, email, etc.
- Devices increasingly mobile and heterogeneous (e.g., IoT)
 - Ad-hoc, disruption-/delay-tolerant networking
 - Broadcast is often intrinsic in comm. medium

DN versus CN

Today we have a
Communication Network
 that we attempt to use a
Distribution Network

CN versus DN


	Communication	Distribution
Naming	Endpoints	Content
Memory	Invisible, Limited	Explicit; Storage = Wires
Security	Communication process	Content

Future Internet Architectures (FIA)

- NSF program, started in Fall 2010
- Five funded 3-year projects:
 - MobilityFirst
 - XIA
 - **NDN**
 - ChoiceNet (funded 2012)
 - Nebula

BTW...

- Security and Privacy in current Internet are NOT a success story
- Retrofitted, incremental band-aid-style solutions
 - E.g., SSH, SSL/TLS, IPSec, IKE
- NSF mandated emphasis on S&P from the outset
- S&P features present in all five architectures



NDN


Named Data Networking

Who is NDN?




Back to CN vs DN

Content Distribution over IP



Content Distribution over NDN



NDN is an instance of CCN / ICN

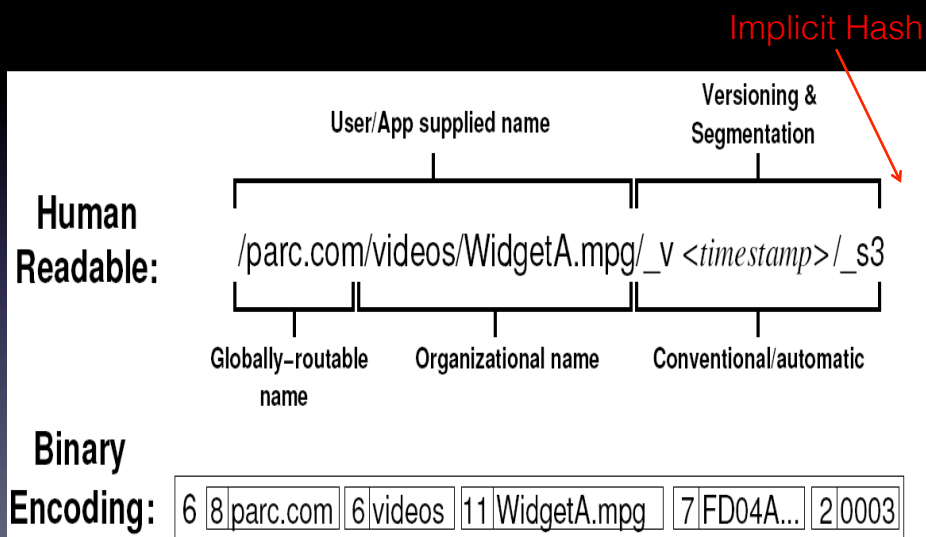
NDN: Basic Concepts

- Consumer
- Producer
- Content
- Interest
- Name
- Router

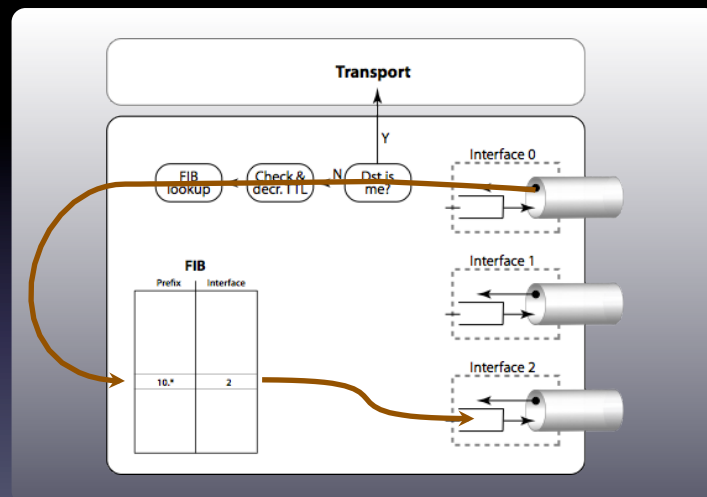
As opposed to IP

- Host
- Interface address (IP address)
- Datagram/Packet
- Router

What's in a Name?



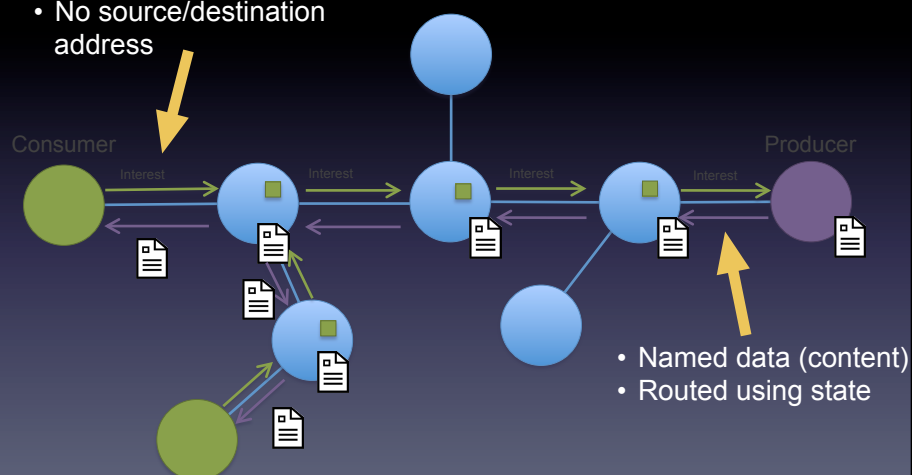
Reminder: IP model



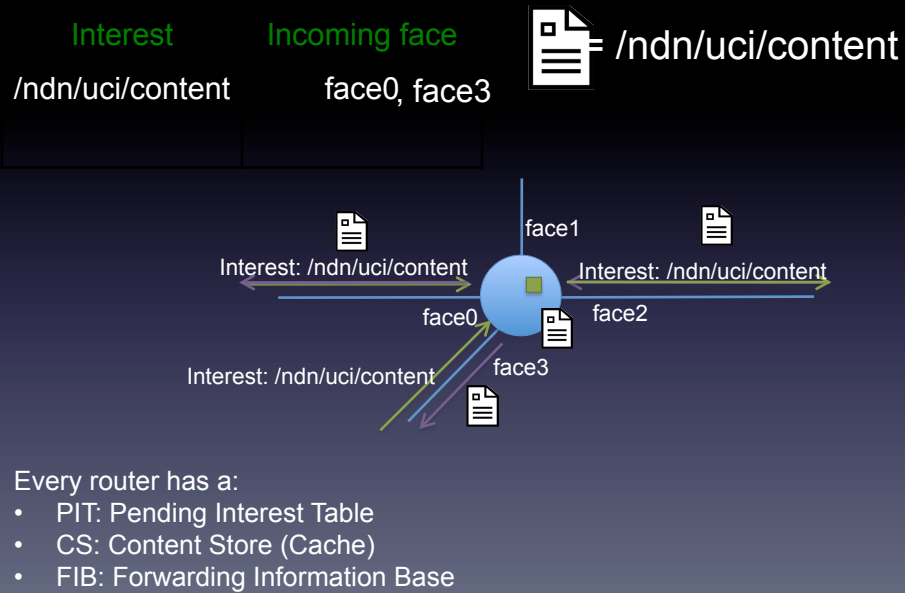
- Each IP packet treated independently of any other
- Stateless routers

NDN Overview: General

- Carries content name
- No source/destination address



NDN Overview: Zooming In





Forwarding and Routing

Forwarding

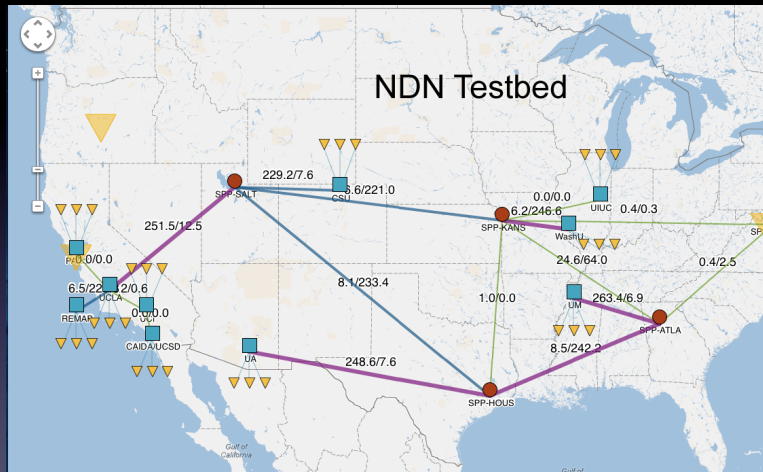
- Key operation is prefix-based longest match lookup, like IP
- Interests forwarded according to routing table, but multipoint forwarding, broadcast, local flooding are all ok
- Data follows interest path in reverse

Routing

- Populating routing tables based on prefix reachability, as in IP
- Can reuse IP routing protocols like IS-IS, BGP



So, is this real?



Current NDN Testbed

Large Scale Demos



www.arl.wustl.edu/~pcrowley/NDN_GEC13_demo.mp4

Outline

1. Named Data Networking (NDN)
2. Security and privacy Issues in NDN
3. A few topics in more detail



Security


- **Now:** secure the pipe
 - Data is authentic because it emanates from the right box (which is an end-point of the right secure pipe)
- **NDN:** Integrity and trust as properties of content
 - Should be inferred from content itself



Authenticity of Content

Content can be retrieved from any place, by any consumer

- How can it be trusted?
- How do we know who produced it?
- How do we know it is the right content?



Securing Content

NDN Content object:

Name
Data
Signature

- **Integrity:** is data intact and complete?
- **Origin:** who asserts this data is an answer?
- **Correctness:** is this an answer to my question?
- **Bonus feature:** routers can choose to verify content (with some caveats)

Securing Content - Performance

- Signing/verifying every content packet is expensive
- Can reduce costs (a little) via techniques like:
 - Merkle hash trees, hash chains, etc, etc.
 - Online/offline signatures
 - Probabilistic verification (spot-checking)



Securing Content

Current SSL/TLS model not a good fit for NDN


- Secures channel, not data
- Authentic content in NDN can come from anywhere
- But, access control (and accounting) is difficult
- After content retrieved from origin, it is served by the network (from router caches)



Private Content

Access to content can be restricted, e.g.:

- Encrypt once with a symmetric key
- Symmetric key distributed using “standard” techniques (pigeons?)
- Access control on key rather than content
 - This can make long-term secrecy problematic



Trust Model

- All content packets are signed
- **Interests are not...**
- NDN is PKI-agnostic
- Application-specific vs network-layer trust



NDN: Privacy Benefits

- Interest has no source address/identifier
- Content can be routed without knowing consumer identity and/or location
- One observed interest may correspond to multiple consumers at various locations
- Router caches reduce effectiveness of observers close to producers

NDN: Privacy Challenges

- Name privacy in interests
 - `/ndn/us/wikipedia/STDs/herpes`
- Name privacy in content
 - `/ndn/zimbabwe/piratebay/XSOQW(#E@UED$%.mp3`
- Signature privacy
 - Leaks content publisher identity
 - Classical privacy vs. security conflict
- Cache privacy
 - Detectable hits/misses

NDN: quick recap

PUBLISHER

- **Announces** name prefixes
- **Names and signs** content packets
- **Injects content** by answering interests

CONSUMER

- **Generates interest packets** referring to content by name
- **Receives content, verifies signature**, decrypts if necessary

ROUTER

- **Routes** interests based on (hierarchical) name prefixes – inherently multicast
- Remembers where Interests came from (PIT), returns content along same path
- Caches content (in CS)
- May verify content signatures



Some Recent & Ongoing Work

- Anonymous content retrieval: ANDaNA
- DoS/DDoS:
 - Content poisoning countermeasures
 - Interest flooding mitigation
- Privacy in Router-Side Caching
- Covert channels
- NDN security in non-distributive settings
 - Instrumented Environments (actuation/control)
 - Sensor Networks
 - Bidirectional low-latency communication
- Trust Management

Why Name Privacy?

NDN names are expressive and meaningful, but...

- Leak information about requested content
- Easy to filter/censor content, e.g., block everything like:

/ndn/cnn/world-news/china/

However:



- NDN names are opaque to the network
- Routers only need to know name component boundaries
- Names can carry binary data

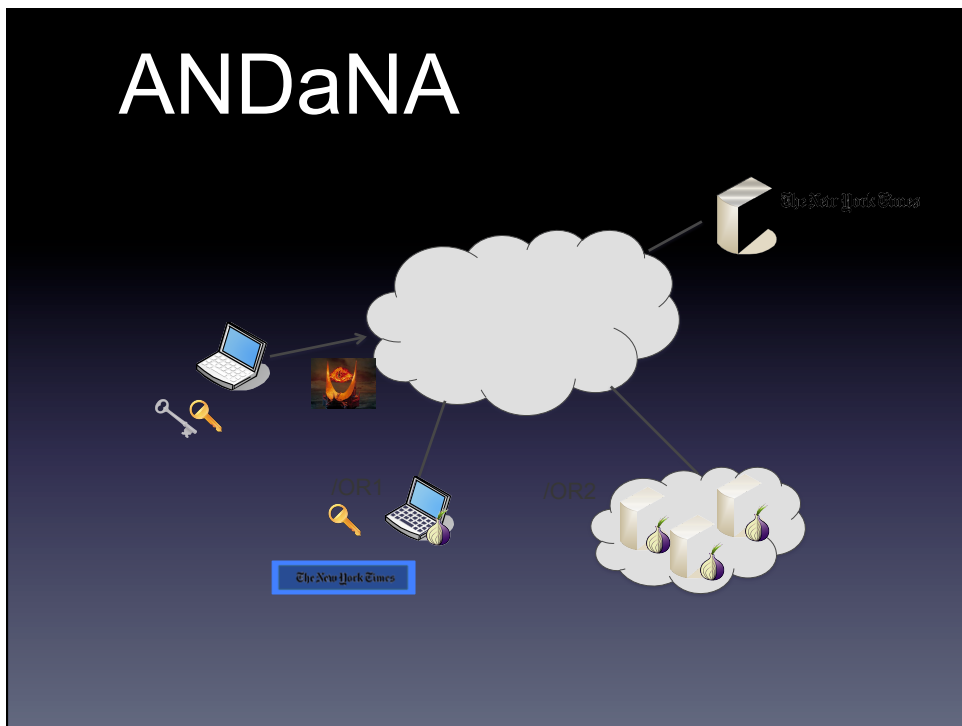

ANDaNA: Anonymous Named Data Networking Application



- Observers close to consumer should not learn what content is being requested
- Low-to-medium-volume interactive communication
- Producers may not be aware of ANDaNA

[DGTU-NDSS2012]













ANDaNA

- Privacy with 2 hops comparable to Tor with 3
 - Why? Lack of source address in interests
 - Anonymizing routers do not learn origin of traffic (only the previous hop)
 - Lower overhead



NDN Cache Privacy

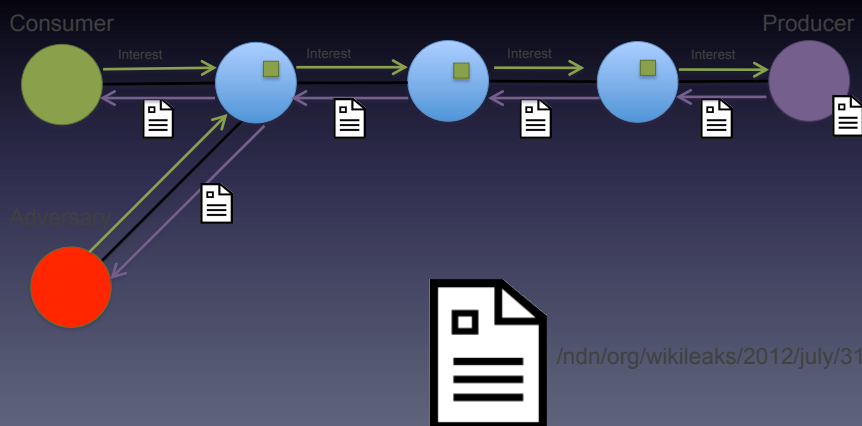
- Router Caching is good for performance
 - Better bandwidth utilization
 - Lower latency
- But... bad for privacy
 - Timing attacks
 - Cache harvesting attacks



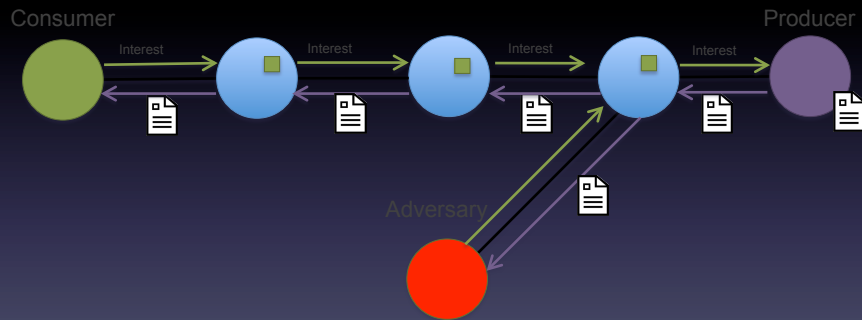
Cache Privacy

- Who could the adversary be?
 - Another host or router
 - A malicious application on victim's device
- Where could the adversary be?
 - Near consumer, e.g., same LAN/WLAN
 - Near producer (opposite sides of first hop router)
 - In both places at once

Scenario 1: Victim=Consumer




Scenario 2: Victim=Producer

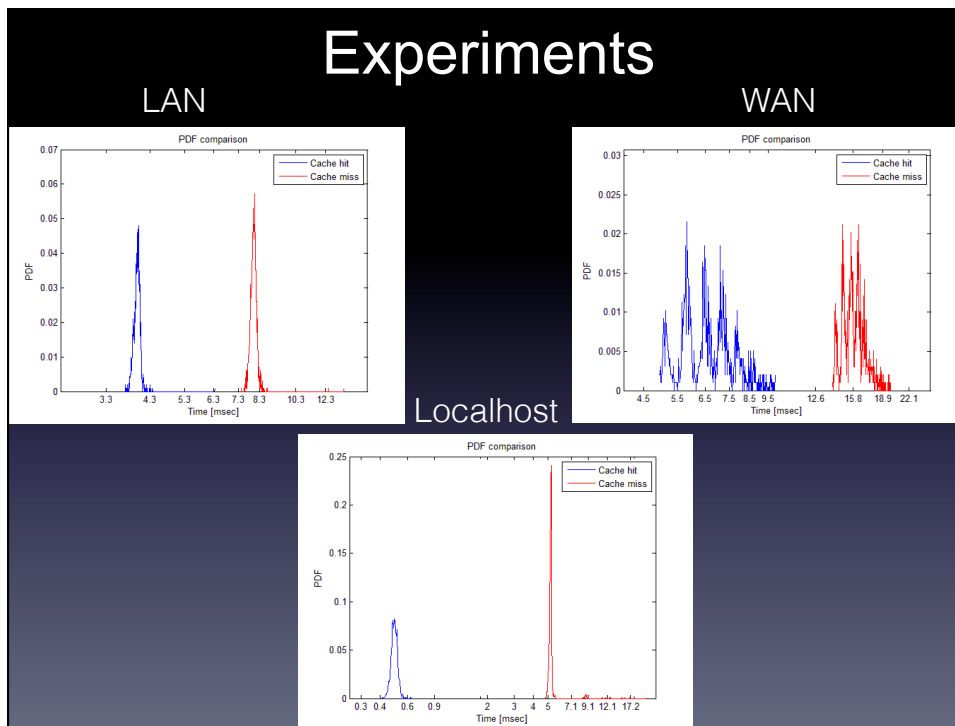


[/ndn.org/wikileaks/2012/july/31](http://ndn.org/wikileaks/2012/july/31)

Scenario 3: Victims=Both



Are Alice and Bob talking?



Countermeasures

- Do not cache content at all
 - Bad idea...
- Cache and delay
 - Which content? Who decides?
 - How long to delay?

Countermeasures

- Two types of traffic:
 - Private
 - Non-private
- Two communication types:
 - Low-latency (interactive) traffic
 - Use unpredictable content names
 - Multicast (distribution) traffic – details in ICDCS'13
 - Random delay
 - Content-specific delay
- Introduce a privacy bit in interests and/or content?

DoS/DDoS in NDN

DoD/DDoS Resistance?

Some current DoS + DDoS attacks become irrelevant because of NDN architecture

- Content caching mitigates targeted DoS
- Content is not forwarded without prior state set up by interests
- Multiple interests for same content are collapsed
- Only one copy of content per “interested” interface is returned

DoS/DDoS

1. Attacks on infrastructure

- Loop-holing/black-holing
- Interest flooding
- Router resource exhaustion

2. Attacks on Consumers + router caches

- Content flooding
- Cache pollution
- Content/cache poisoning

Interest Flooding

Adversary generates numerous non-sensical interests, e.g.:

`/ndn/legitimate-producer/random-string`

- Consumes precious router resources (PIT entries)
- Affects both routers and producers

Interest Flooding

Potential countermeasures:

1. Unilateral rate limiting/throttling

- Resource allocation determined by router state

2. Collaborative rate limiting/throttling

- Routers push back attacks by interacting with neighbors

Content Poisoning

1. Adversary is on the path to producer (e.g., a router)
 - Intercepts genuine interest, replies with fake content
 - Content settles in routers
2. Adversary is NOT on the path to producer
 - Anticipates demand for content
 - Issues own interest(s), replies with fake content
 - Content settles in routers

Content Poisoning

Potential countermeasures:

- Signature verification in routers?
- Consumer feedback?
- Egress router verification only?

BTW: what is "fake" content?

- Bad signature (fails verification),
- Bad signing key

Large Grain of Salt



- NDN represents work-in-progress
- Some security and privacy headaches are gone
- Some new ones are here – they are being explored and (somewhat) mitigated
- The same is true of ALL other FIA projects
- Trust management is the most challenging issue

NDN Security References

"Named data networking project (NDN)",
<http://named-data.org>

"Content centric networking (CCNx) project",
<http://www.ccnx.org>

V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R. Braynard,
 "Networking named content", in ACM CoNEXT, 2009.

S. DiBenedetto, P. Gasti, G. Tsudik and E. Uzun,
 "ANDaNA: Anonymous Named Data Networking Application", NDSS 2012.

J. Burke, P. Gasti, N. Nathan and G. Tsudik,
 "Securing Instrumented Environments over Content-Centric Networking:
 the Case of Lighting Control via Named-Data Networking", IEEE NOMEN 2013.

G. Acs, M. Conti, C. Ghali, P. Gasti and G. Tsudik,
 Cache Privacy in Name-Data Networking , IEEE ICDCS 2013.

P. Gasti, G. Tsudik, E. Uzun, and L. Zhang,
 "DoS & DDoS in Named-Data Networking", IEEE ICCCN 2013.

A. Afanasyev and P. Mahadevan and I. Moiseenko and E. Uzun and L. Zhang,
 Interest Flooding Attack and Countermeasures in Named Data Networking, IFIP Networking 2013.

A. Compagno, M. Conti, P. Gasti and G. Tsudik
 "Poseidon: Mitigating Interest Flooding DDoS Attacks in Named Data Networking", IEEE LCN 2013.

Thank you!

Time for comments/questions