Exploiting OpenFlow resources towards a Content-Centric LAN (CCLAN)

K. Choumas◊, N. Makris◊, T. Korakis◊, L. Tassiulas◊ and M. Ott°

◊ University of Thessaly, UTH
Centre for Research and Technology Hellas, CERTH
° National Informations and Com/tions Technology Australia, NICTA

EWSDN 2013, October 11th, Berlin

Content Centric Network (CCN)

- Content-based search of information, services, etc.

- Contents are retrieved with use of various identifiers:
  - URLs (e.g. Content Delivery Network (CDN))
  - UDP/TCP ports (e.g. http://tools.ietf.org/html/rfc3466)
  - Particular content identifying headers

- CCN could be:
  - An isolated network with novel architecture
  - A LAN that is connected to the Internet
  - A network that relies on the Internet

Content-Centric LAN (CCLAN)

- Content identifier → URL → destination IP address
- Comparing CCLAN with CDN:
  - Both use URL as content identifier
  - CDN could be used on top of the Internet, while CCLAN could be applied on a LAN
  - CDN maps each URL to various IP addresses (for resource availability), changing the operation of the DNS servers
  - CCLAN maps each URL to a unique IP address, like the current DNS protocol does
  - Multiple servers have the same IP address, which now indicates a specific content
- The traditional addressing scheme is maintained, offering backwards compatibility

OpenFlow (OF) networking

- OF switches could replace the Ethernet ones:
  - Extending their capabilities or
  - Modifying their functionality
- In general, OF are able to forward packets not using exclusively their destination MAC addresses
- OF enables the modification of LAN's operation
  - LAN can be any single broadcast domain, including also MAN/WANs that utilize leased lines or VPNs

OF utilization in terms of CCLAN

- The interconnected computers, servers and gateways in a CCLAN do not change their normal operation
- However, some groups of servers/gateways, that offer the same content/service, **share the same IP address**
- As we already mentioned:
  - content is identified by the destination IP address, while
  - OF networking maps each content requester to the most appropriate content provider
  - This is done by modifying the ARP process

CCLAN compared with usual LAN

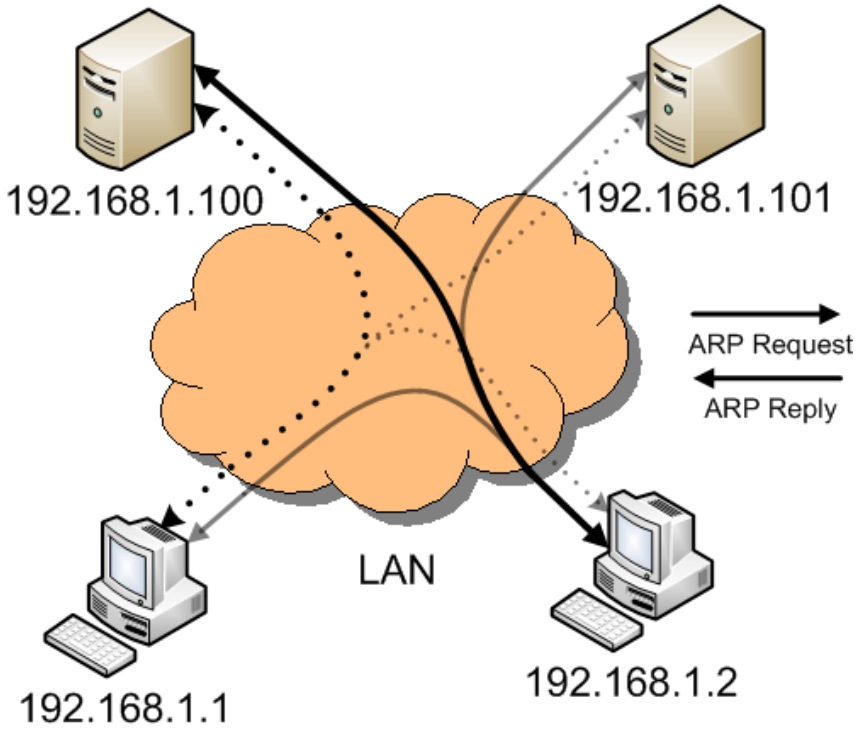

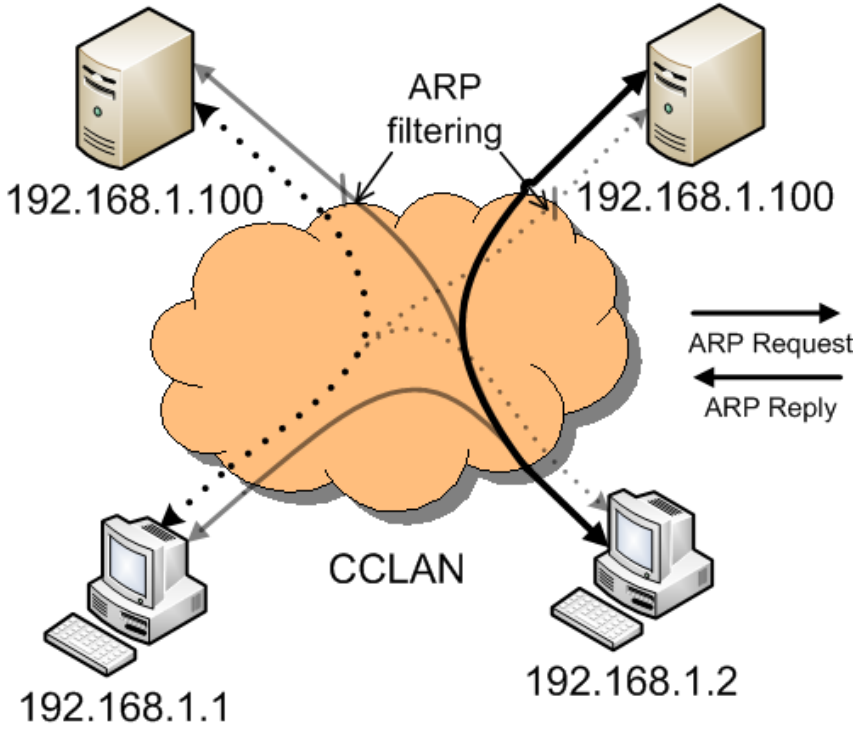

In particular…

- The existing protocol stack remains unmodified
- Content identifying URLs are mapped via the DNS protocol to specific IP addresses, that characterize content as well
- A host machine may feature more than one IP address, depending on the different contents that it provisions
- In this case, an ARP Request for a shared IP address triggers many ARP Replies, which are filtered by the OF switches, imposing various load balancing techniques among the available content servers

Load balancing policies

- **Client-based policy** forwards the ARP Reply of the least loaded server to any new client that initiates an ARP Request, while if the client is not new, then the same server always replies to this user.

- **Load-based policy** tries to dispense traffic among the available servers, stemming from the statistics that OF switches hold.

- **Proximity-based policy** assigns a client to the server that most quickly responded.

Load balancing policies

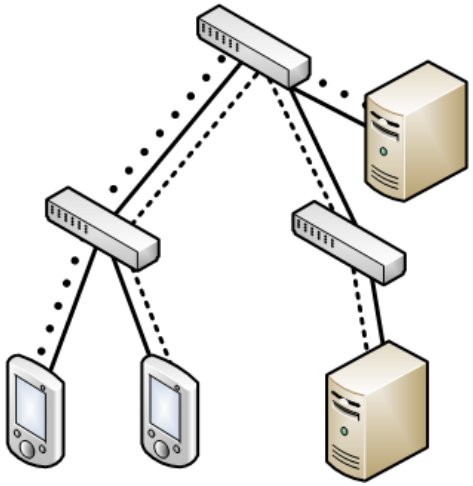

Client-based

Load-based

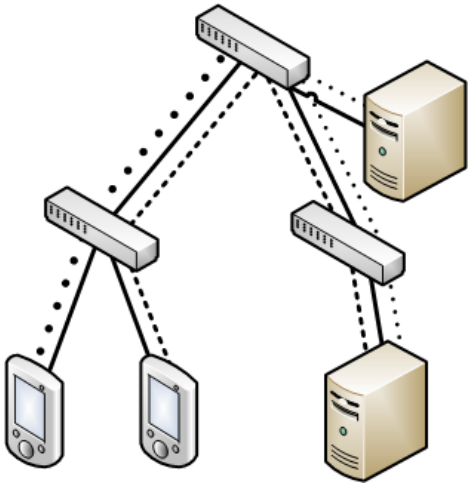

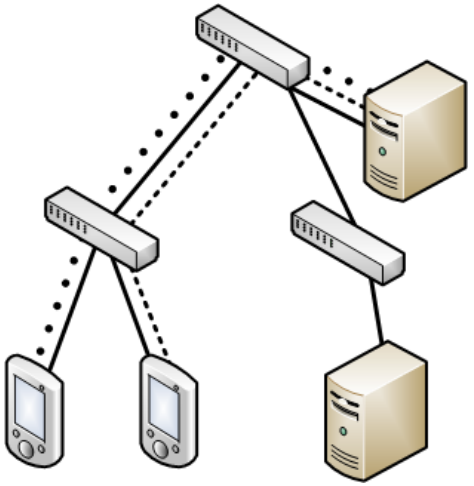

Proximity-based

Experimental Setup

- Experimentation is done in a WAN spanning multiple host machines for various European countries, with use of VPN connections and OF physical and virtual switches
  - 2 OF switches from the NITOS testbed
  - 12 nodes from the NITOS testbed
    - 9 clients, 2 servers and 1 Open vSwitch (OvS)
  - 3 servers from the PLE testbed
- The OF controller is implemented with use of Trema
- The orchestration of the experimentation is based on the OMF framework
- DNS and VPN servers (TAP tunneling) are implemented at the NITOS server

European-wide CCLAN using VPN connections

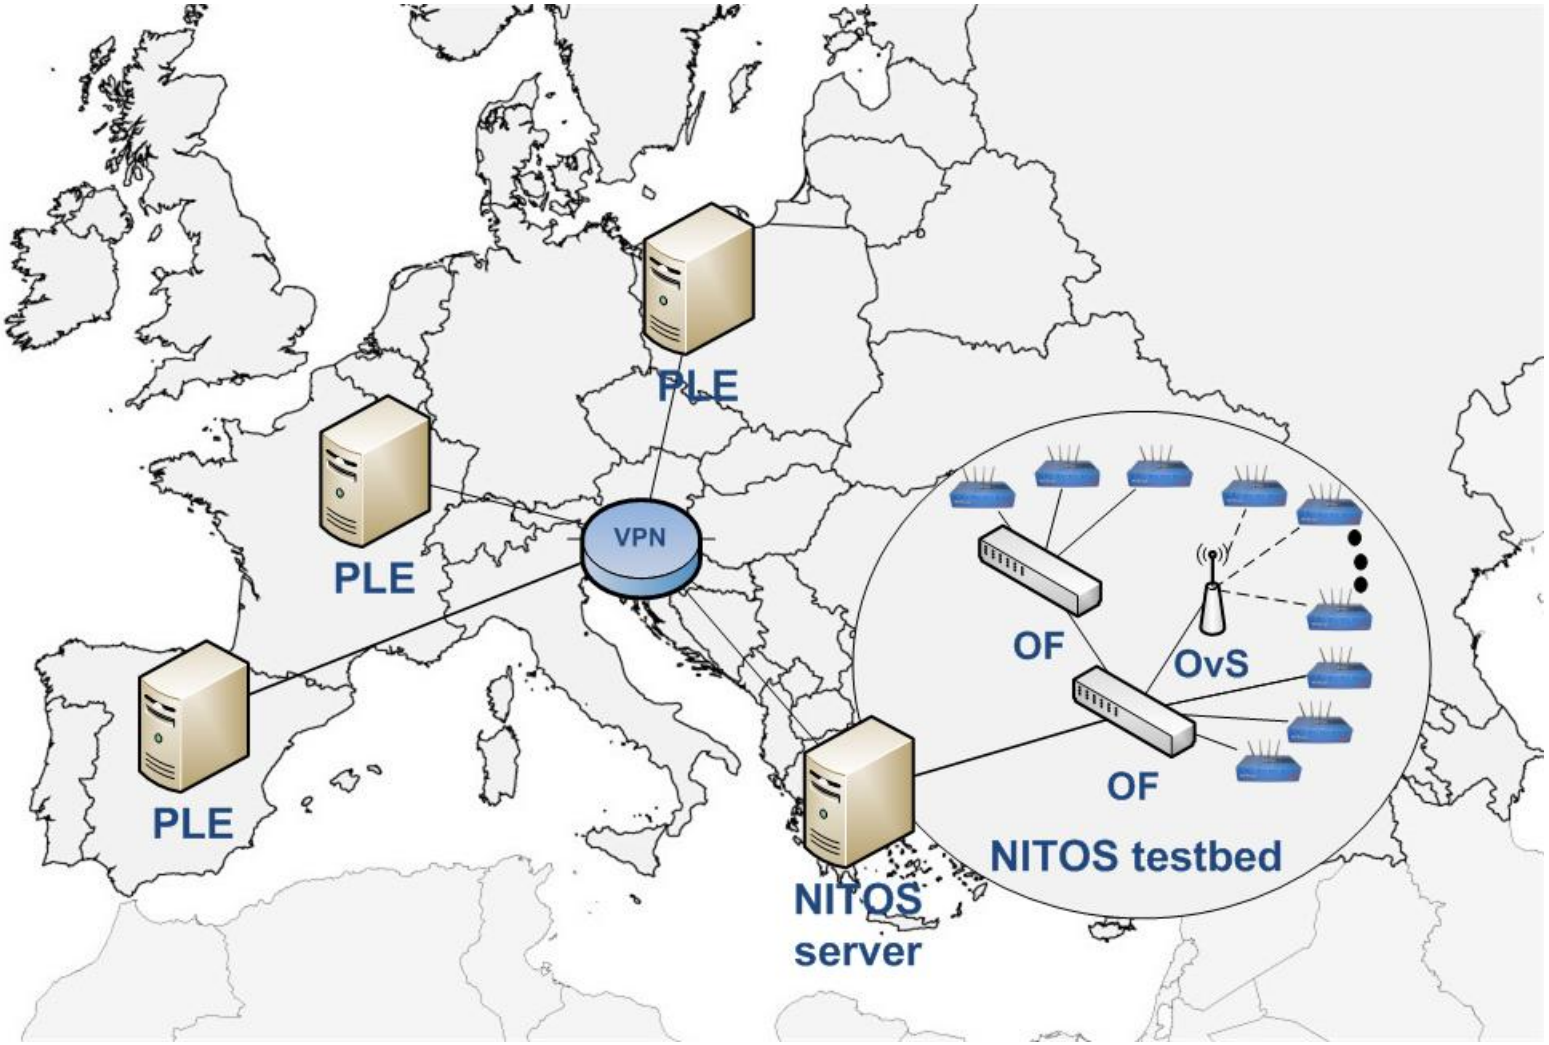

Experimental evaluation of load-balancing policies

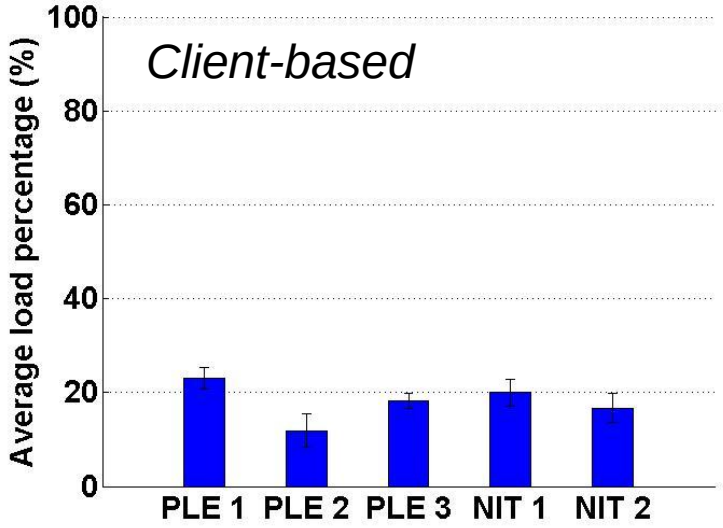

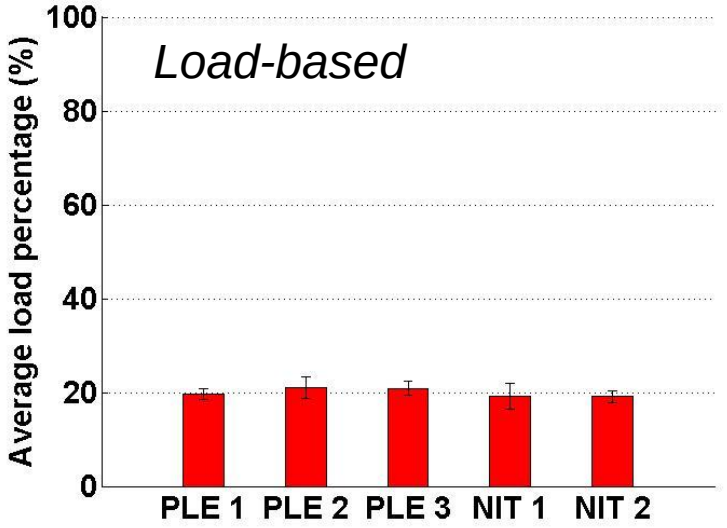

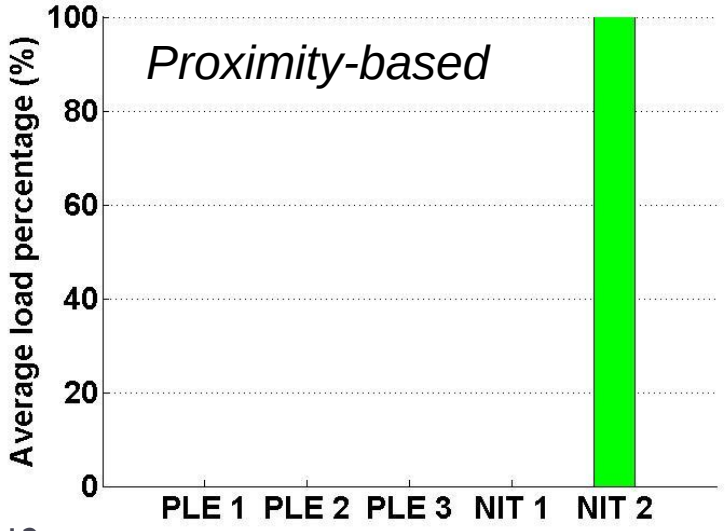

Type of LAN	Average ARP delay
Legal LAN	3.8 msec
Client-based	7.6 msec
Load-based	8.5 msec
Proximity-based	6.9 msec

Conclusions & Future work

- The load-based policy shares almost equally the traffic among the servers
- The proximity-based policy forwards the whole traffic to the closest server
- The proximity-based policy produces minimized ARP delay comparing to the load-based policy
- The policies should be analyzed further, researching deeper in the tuning of their several configuration parameters

The end…

Thank you!!