

<html><head></head><body><pre style="word-wrap: break-word; white-space: pre-wrap;">The present note provides a short summary of the activities of the different IRTF's Research Groups (RG) that held a meeting during the 85th IETF in Atlanta (November 2012).

IRTF Open Meeting (General meeting open to everybody)

=====

Status of the IRTF:

Active Research Groups:

- ASRG - Anti-Spam Research Group
- DTNRG - Delay-Tolerant Networking Research Group
- ICCRG - Internet Congestion Control Research Group
- ICNRG - Information-Centric Networking Research Group
- CFRG - Crypto Forum Research Group

Research Groups with little activity:

- NCRG - Network Complexity Research Group
- NMRG - Network Management Research Group
- RRG - Routing Research Group

Research Groups closing:

- P2PRG - Peer-to-Peer Research Group
- SAMRG - Scalable Adaptive Multicast Research Group

Proposed Research Groups:

- SDNRG - Software Defined Networking Research Group (Proposed)

ANRP - Applied Networking Research Prizes (2 Awards this time):

Srikanth Sundaresan for his measurement study of access link performance on home gateway devices:
Srikanth Sundaresan, Walter de Donato, Nick Feamster, Renata Teixeira, Sam Crawford and Antonio Pescape. Broadband Internet Performance: A View From the Gateway. Proc. ACM SIGCOMM, August 2011, Toronto, Canada.

Slides: <http://www.ietf.org/proceedings/85/slides/slides-85-irtfopen-2.pptx>

Abstract:

We present the first study of network access link performance measured directly from home gateway devices. Policymakers, ISPs, and users are increasingly interested in studying the performance of Internet access links. Because of many confounding factors in a home network or on end hosts, however, thoroughly understanding access network performance requires deploying measurement infrastructure in users' homes as gateway devices. In conjunction with the Federal Communication Commission's study of broadband Internet access in the United States, we study the throughput and latency of network access links using longitudinal measurements from nearly 4,000 gateway devices across 8 ISPs from a deployment of over 4,200 devices. We study the performance users achieve and how various factors ranging from the user's choice of modem to

the ISP's traffic shaping policies can affect performance. Our study yields many important findings about the characteristics of existing access networks. Our findings also provide insights into the ways that access network performance should be measured and presented to users, which can help inform ongoing broader efforts to benchmark the performance of access networks.

Peyman Kazemian for developing a general and protocol-agnostic framework for statically checking network specifications and configurations:

Peyman Kazemian, George Varghese and Nick McKeown. Header Space Analysis: Static Checking For Networks. Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI), April 2012, San Jose, CA, USA.

Slides: <http://www.ietf.org/proceedings/85/slides/slides-85-irtfopen-1.pptx>

Abstract:

Today's networks typically carry or deploy dozens of protocols and mechanisms simultaneously such as MPLS, NAT, ACLs and route redistribution. Even when individual protocols function correctly, failures can arise from the complex interactions of their aggregate, requiring network administrators to be masters of detail. Our goal is to automatically find an important class of failures, regardless of the protocols running, for both operational and experimental networks.

To this end we developed a general and protocol-agnostic framework, called Header Space Analysis (HSA). Our formalism allows us to statically check network specifications and configurations to identify an important class of failures such as Reachability Failures, Forwarding Loops and Traffic Isolation and Leakage problems. In HSA, protocol header fields are not first class entities; instead we look at the entire packet header as a concatenation of bits without any associated meaning. Each packet is a point in the $\{0, 1\}^L$ space where L is the maximum length of a packet header, and networking boxes transform packets from one point in the space to another point or set of points (multicast).

We created a library of tools, called Hassel, to implement our framework, and used it to analyze a variety of networks and protocols. Hassel was used to analyze the Stanford University backbone network, and found all the forwarding loops in less than 10 minutes, and verified reachability constraints between two subnets in 13 seconds. It also found a large and complex loop in an experimental loose source routing protocol in 4 minutes.

NOTE WELL!!!!

Nomination period for 2013 IRTF Applied Networking Research Prizes ends November 30 (More information at <http://irtf.org/anrp>)

DTNRG - Delay-Tolerant Networks Research Group
(<http://irtf.org/dtnrg>)

=====

The Research Group is close to concluding the work it was chartered for. The first part of the meeting was a summary of what has been achieved and what has to be finished. Details on the status can be

found at:

<http://www.ietf.org/proceedings/85/slides/slides-85-dtnrg-0.ppt>

- Marc Blanchet: " LTP, CBHE, and BP Registries"
(<http://www.ietf.org/proceedings/85/slides/slides-85-dtnrg-1.pptx>)
(<https://datatracker.ietf.org/doc/draft-dtnrg-ltp-cbhe-registries/>)

This is just an update on the document requesting IANA to create the registries for LTP, CBHE, and BP protocols.

- Kevin Fall: "Comparing Information Centric and Delay Tolerant Networking"
(<http://www.ietf.org/proceedings/85/slides/slides-85-dtnrg-2.pdf>)

The talk goes through a nice comparison between Information Centric and Delay-Tolerant Networks. The talk highlights both common points (even due to totally different motivation) as well as main differences. The main take away seems to be that despite the fact that they are solving totally different problems there are some design similarities. More specifically, the common research themes are: routing/forwarding scalability; in-network storage management; security and privacy.

ICCRG - Internet Congestion Control Research Group
(<http://irtf.org/iccr>)

=====

- Rong Pan: "A new algorithm for dealing with Bufferbloat"
(<http://www.ietf.org/proceedings/85/slides/slides-85-iccr-2.ppt>)

The core of the proposal, called PIE, is to switch congestion control from a queue length based approach to a delay based approach leveraging on early feedback (opposite to the current approach with feedback only on full queues). The approach is similar to CoDel (by V. Jacobson et al.), however, PIE focuses on enqueue operation whilst CoDel on dequeue. A Linux implementation of the proposed solution has been developed and early measurements show increased throughput and stability (also compared to CoDel).

- Nestor Michael C. Tiglao: "Transport layer caching mechanisms and optimization"
(<http://www.ietf.org/proceedings/85/slides/slides-85-iccr-1.pdf>)

It is a well-known fact that TCP performs poorly in wireless multi-hop networks, hence, in order to improve TCP performance a transport level probabilistic caching model is proposed. Caching helps in reducing transmission cost and loss recovery, as shown both analytically and via simulation (with NS-2). Different cache management policies are explored for different number of flows. Results suggest that is worth to consider a Bandwidth-Delay product congestion control algorithm (instead of AIMD). Nevertheless, caching partitioning seems to be a crucial point to achieve important performance improvement. All results can be found in: Tiglao, N; Grilo, A., Optimal Cache Partitioning in Reliable Data Transport for Wireless Sensor Networks, Proceedings of the 4th Workshop on Network Control and Optimization " NETCOOP"™10, Ghent, Belgium, pp. 90-96, Nov. 29-Dec. 1, 2010.

- Gorry Fairhurst: "Updating TCP to support Rate-Limited Traffic"
(<http://www.ietf.org/proceedings/85/slides/slides-85-iccrgr-0.pdf>)
(<https://datatracker.ietf.org/doc/draft-fairhurst-tcpm-newcwv/>)

The presentation is more an update on the draft-fairhurst-tcpm-newcwv-05.txt. The document proposes a different way to validate the TCP congestion window in order to have better performance for rate-limited applications. The "tuning" of the proposed approach seems not to be finalized since, despite the important number of simulations performed, they are not able to draw a final conclusion.

Further details and papers can be found at:

<http://www.erg.abdn.ac.uk/groups/tsvwg/wiki/30c64/newcwv.html>

ICNRG - Information-Centric Networking Research Group
(<http://irtf.org/icnrg>)

=====

The RG is planning to have an interim meeting in Stockholm the 14/15 February 2013. The meeting will be co-located with the SEAL Project final review, where a demo of the SEAL solution is planned.

Chairs also suggested the following reading list in order to start discussion about deployment.

1. <http://tools.ietf.org/id/draft-fmn-cdni-advanced-use-cases-00.txt>
2. <http://tools.ietf.org/html/draft-li-icnrg-icn-isp-01>
3. http://www.psirp.org/files/Deliverables/FP7-INFISO-ICT-216173-PSIRP-D4.6_FinalReportOnDeplIncBusinessModels.pdf
4. <http://hal.inria.fr/docs/00/68/44/58/PDF/paper.pdf>
5. <http://tools.ietf.org/html/draft-kutscher-icnrg-netinf-proto-00>

- Damien Saucez: "ICNRG Fairness Discussion"
(<http://www.ietf.org/proceedings/85/slides/slides-85-icnrg-0.pdf>)
(<http://trac.tools.ietf.org/group/irtf/trac/wiki/IcnProblemStatement>)

The presentation aims at starting a discussion in the ICNRG concerning the notion of "fairness" in ICN context. During the talk several different definitions of fairness are given (e.g., flow fairness, content fairness, end-user fairness, link fairness, network fairness) and proposes two different type of control: end-to-end and hop-by-hop. The main questions the RG has to answer in order to make progress are: what is a flow? what is a resource? Answering these two questions will allow then to answer the question about what is fairness in the ICN context.

- Will Liu: "On the Content Retrieval in Information Centric Network"
(<http://www.ietf.org/proceedings/85/slides/slides-85-icnrg-1.pdf>)

The talk is about surveying existing work in order to bootstrap the survey document that the RG is willing to write. In particular, some terminology and requirements is presented, based on the existing literature. Further, a summary on the most important key points is presented (e.g., content replication, in-network cache discovery, multicast, k-anycast, content name resolution).

- Kostats Pentikousis: "ICN Baseline Scenarios"
(<http://www.ietf.org/proceedings/85/slides/slides-85-icnrg-2.pptx>)
(<https://datatracker.ietf.org/doc/draft-pentikousis-icn-scenarios/>)

The talk is an overview of the associated draft, which aims at defining common baseline scenarios that can be used by the community as experimental setup and also ground for comparison. The proposed scenarios are taken from real-world use cases (e.g., social networking, real-time communications, mobile networking, content dissemination, energy efficiency, DTNs, infrastructure sharing). Authors are interested to get feedback from implementers on both scenarios definitions and details' level.

- Haiyong Xie: "Coordinated Forwarding and Caching in Content Centric Networks"

(<http://www.ietf.org/proceedings/85/slides/slides-85-icnrg-3.pdf>)

(<http://sites.google.com/site/haiyongxie/draft-xie-icn-coordinated-caching-forwarding.txt>)

The talk focuses on the coordination of forwarding and caching. The basic idea is to extend the caching concept so that not only on-path copies of the content can be delivered to the requester. Rather, content copies are somehow announced in the local domain so that routers are all aware of the content that is locally available. Such information is store in the new Availability Information base (AIB). Finer policies can be imagined, i.e., based on content popularity.

- Dave Oran: "ICN Motivation and Challenges"

(<http://www.ietf.org/proceedings/85/slides/slides-85-icnrg-5.pdf>)

(<http://tools.ietf.org/html/draft-paik-icn-challenges-00>)

(<http://trac.tools.ietf.org/group/irtf/trac/wiki/IcnProblemStatement>)

Dave Oran presented the possible structure of the motivation and challenges document (previously named problem statement) and is asking for people willing to provide contributions.

NCRG - Network Complexity Research Group

(<http://irtf.org/ncrg>)

=====
Michael Behringer (co-chair) provided a short historical summary of the RG activities. The RG started two years ago and had held two meetings so far. A first centerpiece that has to be tackled is the definition of "complexity", for which there is not yet a formal definition. The ultimate goal is to have "practical and objective information on network complexity as an input into the IETF process."

- Xin Sun: "Modeling Complexity of Enterprise Routing Design"

(<http://www.ietf.org/proceedings/85/slides/slides-85-ncrg-0.pdf>)

The work presented (based on a paper with the same title published at ACM CoNEXT 2012) aims at going beyond defining complexity metrics. Focusing on routing, the complexity metric used is configuration dependencies (but could be any other metric), however, the work tries to develop an analytic framework to integrate complexity analysis in the design process.

Using a top-down approach routing design is abstracted and decomposed at several levels of design with respect of the intents.

Starting from the top manager intents, which are realized by the network design, it goes down to device configurations.

The proposed process includes a model of the intent to be sure

that the design meets the intent otherwise it is not even worth talk about complexity.

An example is proposed to show how is possible to they decompose the routing design and how to abstract high-level design intents. They plan to validate the model possibly on new dataset larger than the university network (which was the data set they had access to).

- Michael Behringer: "A Framework for Network Complexity"
(<http://www.ietf.org/proceedings/85/slides/slides-85-ncrg-1.pdf>)
(<https://datatracker.ietf.org/doc/draft-behringer-complexity-framework/>)

The talk presented the content of the related draft document. Such document tries to provide the current understanding of what is a complex network (self-organization, un-predictability, emergence, nonlinearity, fragility), what are the variables involved (state, human operators, templates, dependencies, ownership cost, benchmark cost, churn). The conclusion is a set of possible research directions (definitions and metrics; comparative analysis; containment, control, and reduction of complexity) as well as the proposition of few use cases.

- SDNRG - Software Defined Networking Research Group (Proposed)
(<http://trac.tools.ietf.org/group/irtf/trac/wiki/sdnrg>)

=====

The main focus of the meeting of the SDN RG was about having a panel on the topic "what are the hard open research problems in the SND Space?" The different speakers provide their own view on what is important and worth to carry out research on.

- Dave Meyer (SDNRG Co-Chair): "A Simplified View of the Higher Dimensional the SDN Continuum"
(<http://www.ietf.org/proceedings/85/slides/slides-85-sdnrg-3.pdf>)

A very high level vision of what is the general architecture of SDNs is presented. Besides main technical problems like control/data plane separation; abstraction, and modeling, the speaker points out that there are also sociological and economical issues to be tackled (e.g. flow-based networking, product economics).

- Nick Feamster (SDNRG Co-Chair): "Research Directions in SDN"
(<http://www.ietf.org/proceedings/85/slides/slides-85-sdnrg-6.pptx>)

In a single slide the presenter short-lists what he believes are the most important research directions in SDN. Concerning the data plane the accent is on flexibility and programmability, incorporating a richer set of functions. On the control plane there is a need of higher-level languages to program SDN, which can be used to explore also new policy specifications.

- Diego Lopez: "The Abstraction Track: Bringing the SDN Promise beyond Box Limits"
(<http://www.ietf.org/proceedings/85/slides/slides-85-sdnrg-0.pdf>)

The talk proposes a different view of a network, not anymore as a collection of independent elements (boxes), but rather as a single programmable entity, due to the stronger interaction

between applications and network controllers. With this view is possible to define a Network API providing a higher level of abstraction. Such API will be the interface between SDN applications and the network, which can be considered as a programmable OS. While suggesting the use of OpenFlow as key element, integration with solutions developed in the ALTO WG is also proposed in order to import and export network information.

- Rob Sherwood: "Hard Problems in SDN"
(<http://www.ietf.org/proceedings/85/slides/slides-85-sdnrg-4.pptx>)

The speaker gives his own view on what are hard and important problems in SDN (but not always considered "sexy"). First problem is forwarding memory abstraction. Forwarding equipment use a variety of memory technologies, which have their own trade-offs that need more work to be explored and abstracted. Second problem is the interoperability of OpenFlow and non-OpenFlow networks. For incremental deployment such interoperability is important and there is room for optimization, however, it should be avoided to develop a plethora of new inter-communication protocols. Third problem is about testing the different solutions. It is important to perform test in order to be able to evaluate what solution works better. From a non-technical point of view the speaker claims that standardization of an API between SDN applications and network controllers is a second order problem, since first there is a need of standard applications and network controllers.

- Peyman Kazemian: Troubleshooting SDN
(<http://www.ietf.org/proceedings/85/slides/slides-85-sdnrg-5.pptx>)

Because of the separation between data plane and control plane, in the SDN context, new troubleshooting techniques need to be developed. SDN provides the opportunity to develop richer and more systematic troubleshooting techniques. Eventually this can evolve in automated approaches. An overview of different troubleshooting techniques for different abstraction layers is proposed. The main conclusion is that to develop these new techniques a standard policy expression language is needed.

- Kireeti Kompella: "Research Problems in SDN"
(<http://www.ietf.org/proceedings/85/slides/slides-85-sdnrg-2.pdf>)

The speaker presents his own view on what are the research problems in SDN. A very important issue concerns the different SDN models. There are several models available but no attempt to classify or categorize them. The most widely used SDN models are: Network OS; Broker; Compiler. An overview of the different models is provided. There is no claim on whether a model is better than another, rather that it is important to reduce the problem space by limiting the number of SDN models.

- Ed Crabbe: "A Revaluation of All Values"
(<http://www.ietf.org/proceedings/85/slides/slides-85-sdnrg-7.pdf>)

The presentation asks fundamental questions about network state. What state belongs to distributed protocols? What state must stay in local switches and what state can be centralized? What are the effects of state distribution on different

parameters of the network (e.g., synchronization overhead, stability, efficiency, control loop tightness). The key points are modularity and abstraction, necessary to achieve control, efficiency, and parsimony. Interesting approaches can be identified in FML, Procera, and NetCore.

The speaker is from Google and stated that they are ready to provide funding for interesting academic research in this context.

- Thomas Nadeau: " What are the hard (and interesting) open research problems in the SDN space?
(<http://www.ietf.org/proceedings/85/slides/slides-85-sdnrg-1.pdf>)

The speaker proposed a market-based SDN vision and claims that the most important questions to be answered are:

- Where control plane resides?
- How does the control plane talk to the Data Plane?
- How are the data and control planes programmed?

To this end, the main areas of research are: control plane; programmability; control plane protocols. In this context the points to tackle for control plane distribution and separation are: optimality; resilience; coverage; visibility; scale.

The suggestion is to tighten the collaboration between academia and industry in order to speed up research in this area.

From the room there was the question about what is the difference between the work carried out in the SDN RG and the IRS WG (Interface to the Routing System). Beside pointing out that SDN is a Research Group in the IRTF and IRS a Working Group in the IETF, hence having a different approach and scope, the chairs state that IRS can be considered as a technology enabling SDN but limited only to the routing sub-systems. The SDN RG has a more general and broader approach which (at least for now) does not focus on any specific engineering solution.

</pre></body></html>