Murat Yuksel
yuksem@unr.edu

RESEARCH STATEMENT

My research interests are in ***computer, networked, and wireless systems*** with a recent focus on ***big-data networking, UAV networks, optical wireless, cyber-security and network economics, online network management, cloud-assisted routing, and network architectures***. Major application areas include datacenter management, multi-owner distributed systems (e.g., smart grid), cyber-physical systems, cloud computing, policy sciences, and data-intensive collaborative computations. My work spans both theoretical and experimental aspects of networked systems research. I strongly believe in the power of analytical understanding in research. However, research problems are not always analytically approachable and necessitate hands-on and experimental evaluation. I, therefore, use prototypes and simulations to experimentally evaluate the ideas at hand, and apply analytical modeling methods as long as it is possible. I am a strong believer in inter-disciplinary research and much of my work is collaborative with experts from fields other than networking. Below, I cover some of my current research at University of Nevada, Reno (UNR) and outline future directions. In addition to the followings, my research included wireless routing [1][2], peer-to-peer overlay design [3][8], cross-layer multicast for backbone video delivery [7] (funded by AT&T's VURI Program) and network modeling of thin film growth patterns [10]. For a detailed coverage of my current and past (pre-UNR) research please see my website.

Ongoing Federally Funded Research at UNR

Multi-Transceiver Free-Space-Optical Modules: *https://sites.google.com/site/nsfvlc*

Past Funding: *NSF CNS ($460K, www.cse.unr.edu/~yuksem/fso-manet.htm)*

Current Funding: *NSF CNS ($500K), ARO DURIP ($150K), NASA RIB ($40K)*

Collaborators: *Florida International University, Army Research Lab, and NASA Ames Research Center*

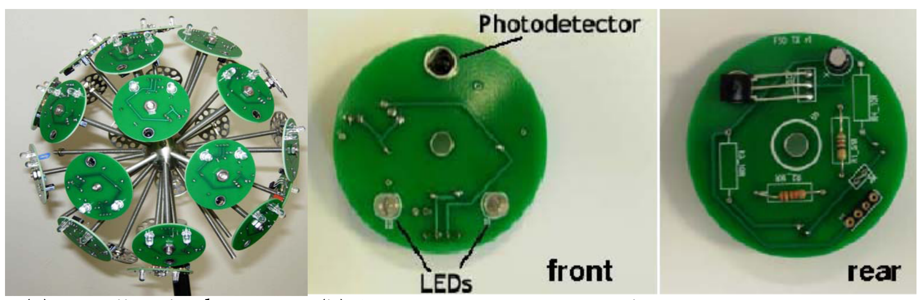

(a) Tessellated sphere (b) One transceiver: Front and rear views

Figure 1: Spherical FSO systems covered with optical transceivers: *Spherical surface provides angular diversity and omni-directional coverage. Spatial reuse becomes possible with multiple directional FSO transceivers.*

Traditional mobile wireless networks have used radio-frequency (RF) communication devices, which require high-power and cause interference. As recent advances in wireless technologies are saturating the RF spectrum, there is an urgent need for new spectrum bands to address the growing wireless demand. I believe that enabling the optical spectrum band in wireless communications is the needed revolution for ultra-high-speed mobile ad-hoc networks (MANETs) of the future. Free-space-optical (FSO) communication has ***great capacity*** (many Gbps) potential, uses 5 times ***lower power*** per bit and is inherently more ***secure*** and ***low-probability-of-intercept*** than RF due to its high directionality and easy containment within walls. On the downside, it is prone to mobility and requires very sensitive line-of-sight (LOS) alignment. In an NSF project earlier, we introduced building blocks for FSO-MANETs. The key idea is to use spherical structures [Figure 1] covered with inexpensive FSO transceivers (e.g., LED and photo-detector pair) solve issues relevant to mobility and LOS management via availability of several transceivers per node. Such structures facilitate electronic LOS tracking (i.e., "electronic steering") methods instead of traditional mechanical steering techniques.

This line of work has been funded by NSF and resulted in several publications [4][5][6]. The concept of multi-transceiver FSO packaging can open a new application for solid-state lighting technology due to the potential integration of illumination and communication on the same devices [12]. We recently got an NSF award on exploring these types of ***illuminication*** (illumination and communication) designs and commenced the collaborative NSF project, for which I am the lead PI. Further, I have received grants from US Army Research Office DURIP and NASA RIB programs to miniaturize our multi-transceiver FSO modules to affix them on Packbot robots and UAVs (that marines use for exploring hostile areas) and K-10 Rovers (that are planned for Mars exploration).

Murat Yuksel
yuksem@unr.edu

Pervasive Spectrum Sharing for Public Safety Communications: *https://sites.google.com/site/pss4psc*

Current Funding: *NSF EARS ($690K)*
Collaborators: *Florida International Univ., Virginia Tech, Univ. of Central Florida, Cisco, and Motorola*

Accommodating the exponentially growing wireless traffic with the existing cellular architectures is not feasible in the long term. The requirements and pressure on getting more from the spectrum is higher for mission-critical applications as in public safety communications (PSC). These trends warranted a new paradigm of increased *heterogeneity* in emerging wireless architectures and increased *sharing* of wireless spectrum. However, sharing of the spectrum has mostly stayed among homogeneous wireless technologies and only at secondary level (i.e., secondary users are allowed only when the primary user is not using the channel). The National Broadband Plan [13] outlined the vision of increasing the efficiency of the spectrum usage via more sharing of it, which we call "*pervasive spectrum sharing*" (PSS) [Figure 2]. The value of such PSS is more pronounced for application scenarios, like PSC, in which the need for spectrum access is vital and infrastructure-less operation is necessary. To realize the PSS vision, models that incentivize users to opportunistically share their spectrum as substrates (e.g., 3G data and WiFi connectivity), and open ***device-to-device*** (D2D) protocols that seamlessly exploit these substrates for wireless applications are much needed. I plan to pursue such game-theoretic D2D protocols that incentivize the users to share their wireless substrates at the primary level while coping with ***jamming*** and ***sniffing*** attacks. I want to develop smartphone apps seamlessly negotiating wireless connectivity using strategies like tit-for-tat, as they proved to be highly successful in peer-to-peer protocols like BitTorrent.

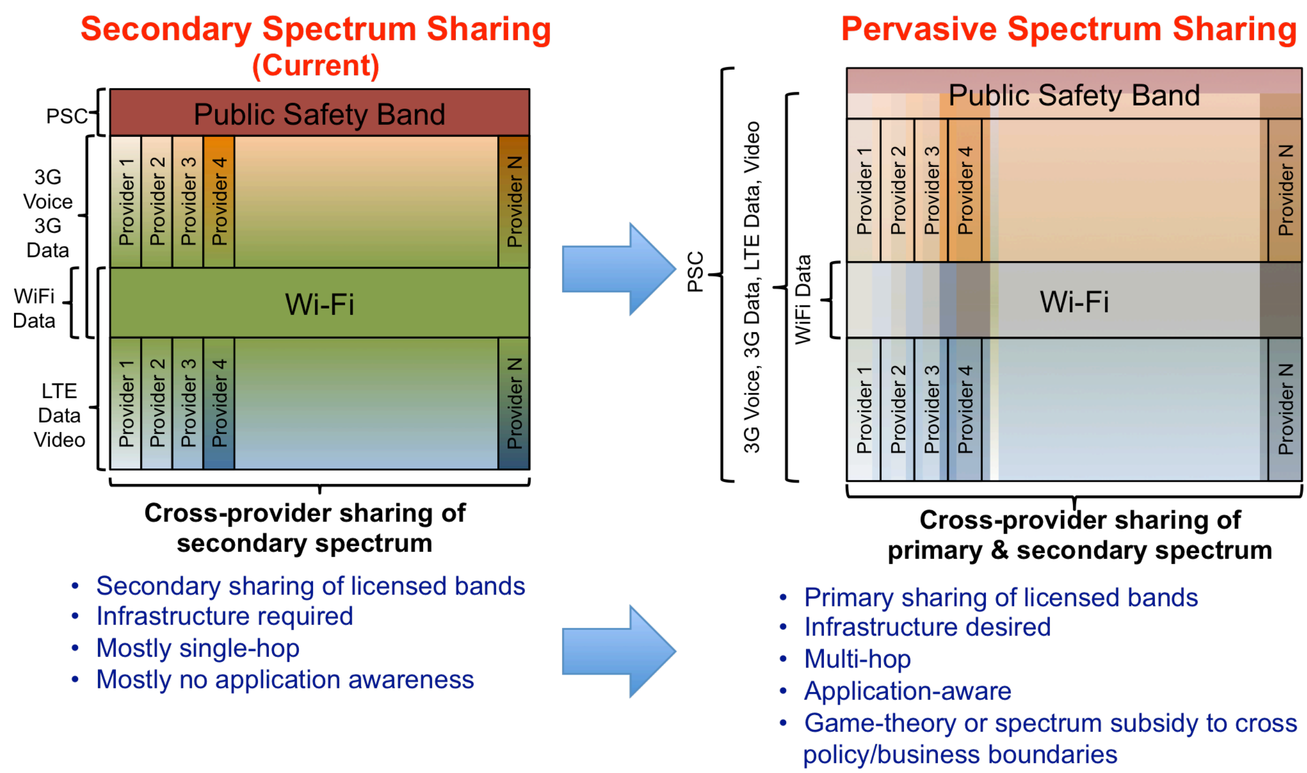

Figure 2: PSS vision: *Evolution of spectrum management towards a smoother sharing across providers and bands.*

A major impediment for PSS is the providers' tendency to protect the bands they earned with a lot of licensing and operating costs. Adopting new technologies to facilitate D2D spectrum sharing is nowadays a key policy consideration for spectrum management. One promising approach to foster more sharing of the spectrum is via the use of regulatory power. We have recently investigated the idea of subsidizing the spectrum to the providers with lower costs in return of ***"proof-of-sharing"*** [14]. That is, the providers will be offered discounted spectrum bands, potentially at different locations, but will be asked to cover users not subscribed to them so as to maintain their subsidy incentives from the government. Our initial results are very promising and In fact, the National Broadband Plan [13] recommends the widespread development of the concept of "spectrum subsidy", particularly for ***securing PSC***. My collaborations with game-theorists, public safety and ***disaster management***, and wireless communication researchers have resulted in an NSF EARS award, for which I am the lead PI.

Online Management, Experimentation, and GAme of Networks: *www.cse.unr.edu/~yuksem/omega.htm*

Current Funding: *NSF CNS ($200K)*
Collaborators: *Behavioral Science Department of UNR*

Online management of a running large-scale network poses many challenges, which have attracted significant research. As critical applications, such as high-definition TV (IPTV) and financial markets, are converging onto the Internet infrastructure, effective response to large-scale network dynamics like failures and demand spikes is gaining more importance. Link or node failures are not rare events for a large-scale network of thousands of devices. Major portion of the time for handling such network dynamics is determining how to respond, mostly performed manually in the current practice. Seeking the optimal response is often impractical, but even settling on a "good" response is very hard as well. Emergence of various networking technologies like 3G wireless and mesh networking is further complicating these management tasks. In most cases, getting the large-scale network to work is the typical target. Experienced human administrators are typically the ones who can quickly find a close-to-optimum response. However, as the networks are getting larger and more diverse, managing and attaining effective responses for an online operational network necessitates ***meta-networking*** tools to swiftly learn and characterize the network. We respond to this need by developing tools to achieve automated ways of managing a running network. Particularly, we frame heuristic optimization, empirical learning, experimental design, and network management with a "game" interface [11]. Similar to ***Flight Simulator*** for trainee pilots, we are developing an online management and experimentation system for large-scale networks in a game-like environment for trainee administrators to play with and explore ***what-if scenarios***, without having to risk the network operation. We also develop algorithms for empirical characterization of network dynamics, and tools for quick and close-to-optimal configuration of numerous network parameters in response to failures or customer traffic trends [9]. I plan to this network management research agenda to ***security management*** for large-scale distributed multi-owner systems, such as the ***smart grid network***.

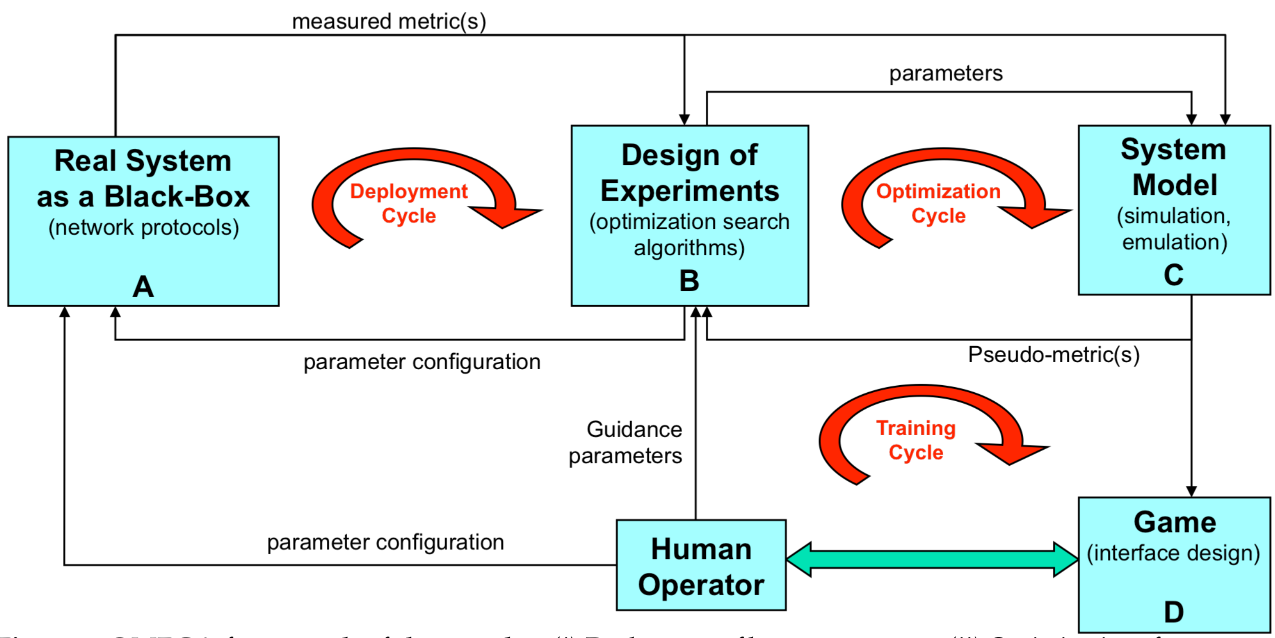

Figure 3: OMEGA framework of three cycles: *(i) Deployment of better parameters, (ii) Optimization of parameters via meta-networking tools, and (iii) Training of human operators via game-like environments.*

Murat Yuksel
yuksem@unr.edu

Future Research Directions

UAVs + Directional Communication = 3D Spatial Reuse: UAVs offer a new dimension to legacy wireless networking by enabling Troposphere altitudes for spatial reuse. As sketched in Figure 4, UAVs can effectively relay (e.g., the red path) traffic from a ground node to another one without causing any major interference to the rest of the ground nodes. If this Tropospheric 3D space is not used, the aggregate throughput of the mobile mesh wireless network on the ground will be limited by the interference in a highly dense and/or saturated spectrum environment at low altitudes, e.g., below 600ft. Further, most of these low altitudes and ground nodes (e.g., class 1 UAVs, UGVs, smartphones) are currently based on omnidirectional RF antennas which limit the spatial reuse. Organized in a hierarchical 3D mesh, UAVs can open up to 1200ft of space with no FAA regulations [17]. This by itself is a 2-fold increase in the spatial reuse gain. Satellites have been attaining this type of gain via microwave or optical bands, but still with large divergence angle beams. Since the main purpose of satellites has been to offer baseline connectivity, the beam angle has been narrow to assure large coverage areas on the ground (see the green beams in Figure 4). Exploitation of the Troposphere (lower altitudes than satellites) is possible if UAVs are organized into a 3D mesh. If such a ***3D mesh of UAVs*** at various altitudes is attained, the total gains can go as high as the upper limits of Troposphere, i.e., 7km. A smart and self-organizing 3D mesh can potentially double the aggregate throughput at every layer going up, and hence, open up ***a super-linear throughput increase possibility for wireless networking***. However, these spatial reuse potentials are possible only if UAVs are equipped with directional antennas with high precision steering and positioning capability. Due to their small form factor and low power consumption, FSO transceivers are great candidates for directional communication on UAVs. I plan to leverage our prior work and prototype multi-transceiver FSO modules conformal to UAV surfaces. Affixing these conformal FSO modules (as in Figure 5(a)) to mechanically steerable parts on UAVs will enable two knobs for managing LOS alignment of an FSO link between two UAVs: ***electronic and mechanical steering***. Our initial studies [18] show that mechanical steering of such highly directional antennas is feasible for UGVs (up to 5m/s speeds) as well as UAVs (up to 25m/s speeds). Opportunistic usage of an FSO link between UAVs and falling back to RF if the FSO link is broken will allow us to construct a ***hybrid RF/FSO 3D mesh of UAVs***, which will open the doors for non-linear increases in 3D spatial reuse and the aggregate capacity of wireless networking. I plan to collaborate with nano-positioning (for high precision steering) and solar power (for longer flying times) experts to undertake an experimental research agenda to prototype UAVs (as in Figure 5(b)) maintaining an FSO link between each other while flying.

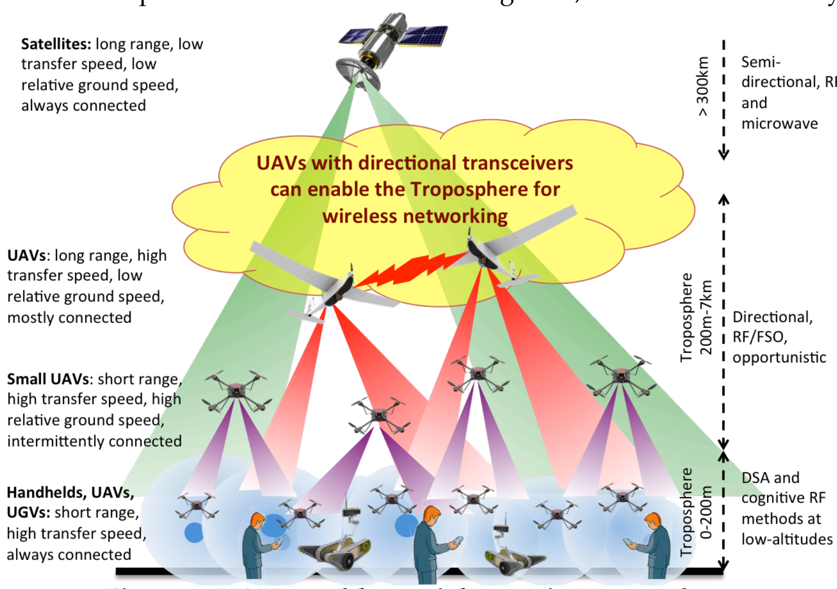

Figure 4: UAVs enable spatial reuse in Troposphere

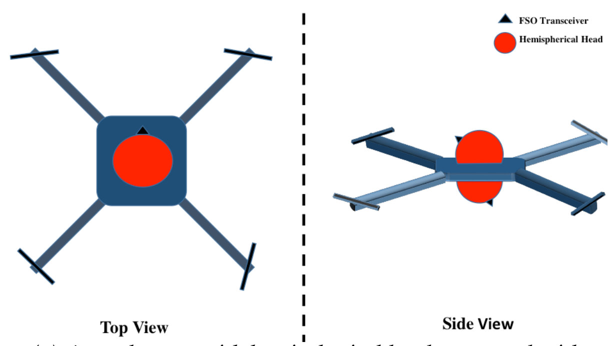

(a) A quadcopter with hemispherical heads mounted with conformal optical transceivers.

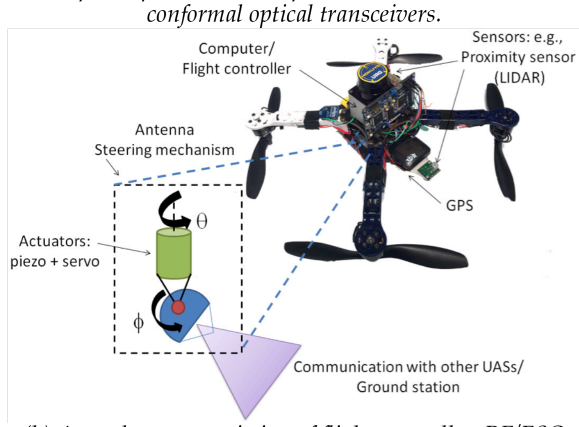

(b) A quadcopter consisting of flight controller, RF/FSO antennas, steering mechanism, and localization sensors.

Figure 5: UAV systems for RF/FSO 3D mesh

Multi-Core Parallel Routing for Big-Data Transfers: As datacenters are gaining more importance in the Internet, the need for big-data intra- and inter-datacenter transfers is more paining. Most of the end-to-end (e2e) sessions are now going through a datacenter, and thus, the performance of large "big-data" intra- and inter-datacenter transfers is very important to the overall Internet experience. More essentially, these big data transfers are crucial to the operation of the data centers due to the needs for large bulks of transfers for maintenance and backup. A successful approach to address the big-data transfers is to use multi-path routing, which spreads the paths available to the e2e transport while satisfying various constraints such as delay or loss. Since the problem is too complex most multi-path routing work boiled down to pre-computed techniques with heavy computations. These multi-path routing techniques, however, proved to be useful for scaling up the e2e reliable transfers. The paths generated by these multi-path routing methods were adopted and TCP sessions were successfully parallelized with effective solutions. Yet, these e2e transfers do not to directly utilize multi-core CPUs available in most routers. We recently experimented with the idea of a "parallel routing" framework that explicitly considers multi-core routers and employs shortest-path calculations only. The key idea is to virtually slice the router topology into "substrate" topologies and assign them to a separate router core, which calculates a shortest path on the assigned substrate. Initial evaluation of parallel routing is very promising [15]. I plan to develop a prototype of it at ESNet via collaborations with Lawrence Berkeley National Lab (LBNL). I have been working on collaborative proposals on the topic with SUNY Buffalo and LBNL researchers, and plan to submit to NSF CSR and Dept. of Energy programs.

Cloud-Assisted Routing: The recent emergence of cloud computing services with high flexibility and virtualization options has made the networking community consider offloading some of the networking services to the cloud. The declining costs and high flexibility of the cloud services make the cloud very attractive for long run. I would like to explore the possibility of circumventing non-trivial complexity and cost of legacy core Internet routers by using cloud services. In particular, we aim to mitigate the increasing routing complexity to the cloud with a new architectural approach, Cloud-Assisted Routing (CAR), and seek an answer to the following question: "*Can techniques leveraging the memory and computation resources of the cloud remedy the routing scalability issues?*" Currently, there are more than 30,000 service providers advertising more than 350K IP prefixes, with a typical edge router receiving hundreds of updates per second. Further, BGP churn can grow prohibitively if topology growth and update dampening are not performed carefully. These concerns become more serious as the Internet topology is becoming more flat, putting more burdens on the core routers. Even if we assume technical scalability is possible with the existing router architectures, the cost of a router became non-trivial economically as the complexity on the routers have increased. The cost of routing unit traffic has not been reducing at a pace similar to the performance improvement of computing capabilities of a router. Given the trends on the state and packet processing capacities expected from a BGP router, the cost of a router that can perform the basic routing functions at the Internet core is unlikely to reduce. These trends clearly point to the urgent need for techniques and architectural approaches reducing or offloading complexities on the routers. As the cloud prices are declining and the cloud services are getting closer (e.g., average delay to the closest cloud provider is now in the order of milliseconds), I believe hybrid and opportunistic routers exploiting cloud services while maintaining the most

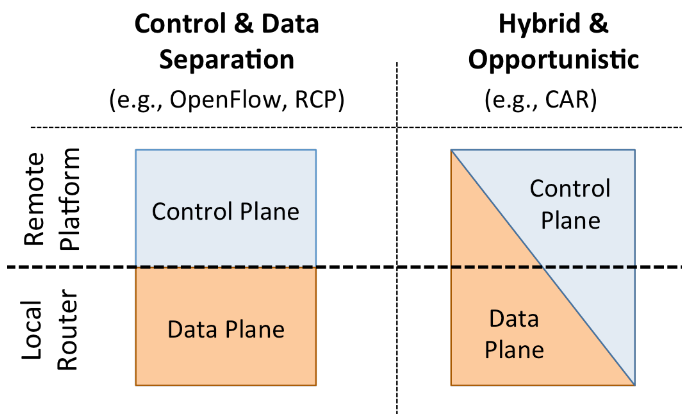

(a) Routing function placement: CAR vs. Legacy SDN

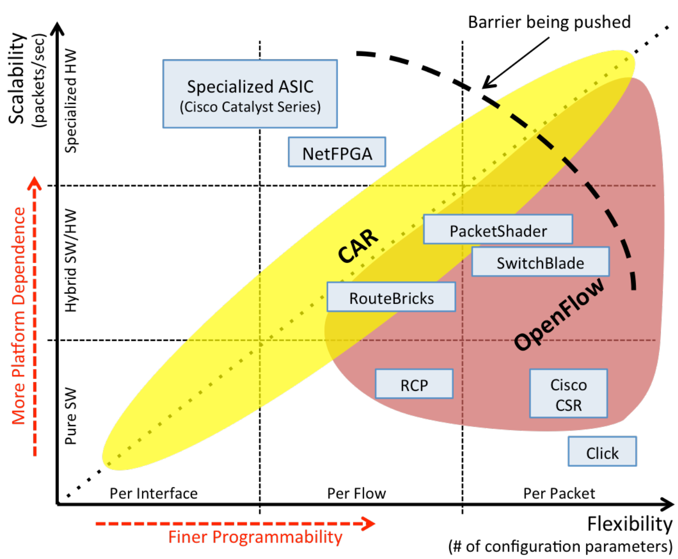

(b) Scalability-flexibility tradeoff for routing

Figure 6: CAR vision

important tasks at the router hardware present a worthy ***software-defined networking*** (SDN) research direction. A key difference of our CAR approach from the recent SDN approaches, like OpenFlow, is its hybrid approach in maintaining both data and control plane routing functions at the router hardware and the cloud. I believe that a hybrid approach [Figure 6] both in data and control planes will likely prove better in the long run. Our initial prototyping efforts [16] showed promising results for offloading CPU complexity to the clouds. I plan to submit proposals to NSF NeTS program and keep exploring CAR in both economically viable and experimental settings.

Economics of Cyber-Security and Cyber-Sharing – Networked Value/Risk and Crypto-Currencies: I find network economics and cyber-security to be highly intertwined areas; but economics of security has typically been ignored or understudied. Arguably, cyber-security problems are a result of cyber-sharing technologies that the networking community has come up with. We share because a cyber-good has more value when shared (e.g., a tweet's value increases only when shared/re-tweeted), or it is utilized more when shared. Going back to Metcalfe's Law, networked value is clearly super-linear as the sharing increases. Yet, sharing comes with a risk of privacy and security attacks. I think there are very interesting research challenges to explore on the dimensions of economics, security and privacy. I give two examples below.

Some of my recent research includes game-theoretic protocols for wireless D2D communication, which is considered to have great potential to respond to the RF spectrum scarcity. It is a challenge to design ***open yet secure*** protocols that can scale to 100s of devices. Obviously, device users will not be willing to use open protocols for sharing wireless connectivity and resources unless they feel comfortable about the security of their devices and confidentiality of their data. Scalable, distributed and plausible incentive mechanisms are needed to attain more spectrum sharing among the devices as well as wireless operators. Using cloud to assist such D2D and other networking protocols pose similar problems as well. Even though the cloud helps in scaling the protocols and provide great economy of scale, users are typically uncomfortable with putting their confidential data in a cloud environment given that several data leakages from the cloud providers were reported.

Multi-owner distributed systems exhibit very similar security-value tradeoffs as well. A very interesting one is crypto-currencies like ***BitCoin*** (BC). BC market has gained a lot of traction recently and crypto-currencies claim a large portion of the "value expression" tools. Many people now think that crypto-currencies will be a very successful mean of expressing and transferring value for the human society. BC, in particular, is maintained by a multi-owner networked system of peer computers. Clearly, the BC value is a networked value, but yet, it is a networked risk as well because of its vulnerability to security attacks. It is intriguing whether or not such crypto-currencies, realized with networked value, will really be successful. It is obvious that bankers (who produce the current real currencies) will be stiff against BC. But, I believe that in the long run BC has a significant edge in lowering the cost of maintaining and expressing value, which will be more attractive to the banking system as well.

Looking at it from an economic view, as is, we are spending time and resources to maintain a one-dollar bill. We spend time to design the dollar bill, all the figures and other things on it. We further spend time and resources (actual paper, the machinery) to print the bill. So, let's say, hypothetically, that we as the whole society are spending 1 cent to maintain a one-dollar bill over a year time frame. So, our dollar's "value expression efficiency" (VEE) is 99%. In BC, the main cost will be CPU, memory and power consumption of computers. Electrical power is used for maintaining the current currencies as well – so we can exclude it from the discussion for simplicity. CPU and memory has been historically getting cheaper over several decades we have observed. It is fair to say that it will continue to do so. From this very high level view, I think BC will have a higher VEE and thus be more attractive to the bankers as well. Major transitions (or spread) of significantly new technology require clear economic motives. I see a lot of relevance in the change from the traditional phone network to the voice-over-IP (VoIP). For many years, the telco companies (the owners of the lines in the traditional phone network) resisted against the move to the VoIP. But, after realizing that transferring the phone traffic over the Internet is much cheaper, they themselves worked for moving the phone on to the Internet. It required them to swallow the hard pill of making a huge change to their infrastructure and innovate new solutions needed for the necessary changes to their networks. I see the bankers against BC in the same spot. They will eventually accept digital currencies and even motivate people to use it.

By means of network economics and game-theoretic models, I want to explore above-mentioned tradeoffs that exist in ***multi-owner*** distributed large-scale systems where value of ***sharing conflicts with security/privacy*** of participants. As was also posted in this year's MURI BAA, US Department of Defense has been very interested in models that can enlighten how groups cooperate even though they may have conflicting goals. I plan to target DoD solicitations on this topic in collaboration with game-theorists and experts in finance.

References

[1] M. O. Kilavuz and M. Yuksel, Path Approximation for Multi-Hop Wireless Routing under Application-Based Accuracy Constraints, *Computer Networks*, Elsevier, 56(1), pp. 345-364, January 2012. (also in IEEE MASS 2008)

[2] B. Cheng, M. Yuksel, and S. Kalyanaraman, Orthogonal Rendezvous Routing Protocol for Wireless Mesh Networks, *IEEE/ACM Transactions on Networking*, 17(2), pp. 542-555, April 2009. (also in IEEE ICNP 2006)

[3] H. Guclu and M. Yuksel, Limited Scale-Free Overlay Topologies for Unstructured Peer-to-Peer Networks, *IEEE Transactions on Parallel and Distributed Systems*, 20(5), pp. 667-679, May 2009. (earlier version in IEEE ICDCS 2007)

[4] A. Sevincer, M. Bilgi, and M. Yuksel, Automatic Realignment with Electronic Steering of Free-Space-Optical Transceivers in MANETs: A Proof-of-Concept Prototype, *Ad Hoc Networks*, Elsevier, 11(1), pp. 585-595, January 2013.

[5] M. Bilgi, A. Sevincer, M. Yuksel, and N. Pala, Optical Wireless Localization, *Wireless Networks*, ACM/Springer, 18(2), pp. 215-226, February 2012.

[6] B. Nakhkoob, M. Bilgi, M. Yuksel, and M. Hella, Multi-Transceiver Optical Wireless Spherical Structures for MANETs, *IEEE Journal on Selected Areas of Communications*, Vol. 27, No. 9, December 2009.

[7] M. Yuksel, K. K. Ramakrishnan, R. D. Doverspike, R. K. Sinha, G. Li, K. N. Oikonomou, and D. Wang, Cross-Layer Failure Restoration of IP Multicast with Applications to IPTV, *Computer Networks*, 55(9), pp. 2329-2351, June 2011.

[8] B. Cheng, M. Yuksel, and S. Kalyanaraman, Virtual Direction Routing for Overlay Networks, *Proceedings of IEEE International Conference on Peer-to-Peer Computing (P2P)*, pp. 61-70, Seattle, WA, September 2009.

[9] B. Gonen, G. Gunduz, and M. Yuksel, Automated Network Management and Configuration Using Probabilistic Trans-Algorithmic Search, to appear in *Computer Networks*, Elsevier Science. (also in IEEE GLOBECOM 2011)

[10] T. Karabacak, H. Guclu, and M. Yuksel, Network Behavior in Thin Film Growth Dynamics, *Physical Review B*, 79(19), May 2009.

[11] E. Arslan, M. Yuksel, and M. H. Gunes, Network Management Game, *ACM CCR*, 43(1), pp. 46-50, January 2013.

[12] A. Sevincer, A. Bhattarai, M. Bilgi, M. Yuksel, and N. Pala, LIGHTNETs: Smart LIGHTing and Mobile Optical Wireless NETworks – A Survey, *IEEE Communications Surveys and Tutorials*, 15(4), pp. 1620-1641, April 2013

[13] "Connecting America: The National Broadband Plan," FCC, 2010. http://download.broadband.gov/plan/national-broadband-plan.pdf

[14] M. Yuksel, T. Quint, I. Guvenc, W. Saad, and N. Kapucu, Fostering Wireless Spectrum Sharing via Subsidization, *Proceedings of IEEE Allerton Conference*, Urbana-Champaign, IL, October 2013. (invited paper)

[15] A. Soran, F. M. Akdemir, and M. Yuksel, Parallel Routing on Multi-Core Routers for Big Data Transfers, *Proceedings of ACM CoNext Student Workshop*, Santa Barbara, CA, December 2013.

[16] H. T. Karaoglu and M. Yuksel, Offloading Routing Complexity to the Cloud(s), *Proceedings of IEEE ICC Workshop on Cloud Convergence: Challenges for Future Infrastructures and Services (WCC)*, Budapest, Hungary, June 2013.

[17] J. J. Roberts, So you want to fly drones? Here's what the law says, GIGAOM, December 2013. https://gigaom.com/2013/12/08/so-you-want-to-fly-drones-heres-what-the-law-says

[18] M. R. Khan and M. Yuksel, Maintaining a Free-Space-Optical Communication Link Between Two Autonomous Mobiles, *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC)*, Istanbul, Turkey, April 2014.