Research Statement

Michalis Faloutsos

My research focuses on computer networks and tries to measure, model and improve the capabilities of current networks. The work spans the three major networking technologies: a) the Internet, which has become an essential technology for people's everyday life, b) wireless ad hoc networks, which is considered the next disruptive technology, and c) wireless cellular networks, which have enjoyed an unprecedented growth over the last five years.

This document provides a quick overview of the research efforts of my research group in each of the three networking technologies and some highlights of the major accomplishments of our work.

- **Publications and Citations:** I have published 39 conference papers, 9 journal papers, and have 8 submitted journal papers. My published papers include 4 conference papers in INFOCOM, 2 in SIGCOMM, 1 in PODC, 1 in ICNP, 2 in IMC and 2 journal papers in IEEE Trans. on Networking and 1 in Distributed Computing.

  I have co-authored one of the top-20 most-cited paper of 1999 in all CS and EE research: "On powerlaw relationships of the Internet topology", with Petros Faloutsos and Christos Faloutsos
  SIGCOMM'99 (Statistics by the Digital Library: www.citeseer.nj.nec.com)

  My work has been referenced by popolular magazine and books:
  Scientific American, in article "Scale-Free Networks", vol. 288, no. 5, 2003
  Book: "Linked: by The New Science of Networks", by Albert-Laszlo Barabasi.

- **Tools and Software:** We have released software that has been highly popular world wide and has attracted the attention of special interest groups.

  - Our SELFIS software release has been downloaded by more than 300 researches across the globe spanning multiple disciplines. The tool, which was developed in collaboration with my PhD student by T. Karagiannis, provides novel capabilities for analyzing network performance.

    SELFIS site: http://www.cs.ucr.edu/∼tkarag/Selfis/Selfis.html
  - We have developed a novel methodology and a tool, NEMECIS, for analyzing and safeguarding the Internet at the BGP level. The methodology and tool were published in INFOCOM 2004, and has attracted the attention of network administrators and practitioners from the the North American Network Operators' Group (NANOG) and Internet Engineering Task Force (IETF) communities. The tool was developed in collaboration with my PhD student G. Siganos.

    NEMECIS site: http://ira.cs.ucr.edu:8080/Nemecis

1 Wired Networks: the Internet

Our work has focused on measuring and modeling the topology, the traffic, and their interactions with the Internet protocols.

Modeling the Internet Topology: The Revolution of Power-laws. (SIGCOMM [8] ToN [9]) We introduce an elegant way to describe the properties of the Internet topology. We observe that the distribution of these properties on measured data is highly skewed, and it can be described fairly accurately by power-laws. The concept of power-laws and skewed distributions has created a mental shift in the way one visualizes the Internet topology. This work has been widely cited and created significant follow up research, as we mentioned above.

Traffic Modeling: Poisson versus Long Range Dependence. (INFOCOM [22]) Most current traffic models seem to view traffic as a long memory (or long range dependent) process. Our work challenges the status quo in traffic modeling by showing that current network traffic can be well represented by the Poisson model (a memoriless process) for sub-second time scales. We find that the traffic characteristics have significantly changed since the first studies that identified long memory processes. Furthermore, we propose a model that attempts to reconcile the seemingly conflicting Poisson and long memory views of the Internet traffic.

In more detail, the identification of long range dependence in network traffic ten years ago has largely discredited Poisson-based models. However, since that original data set was collected, both link speeds and the number of Internet-connected hosts have increased by more than three orders of magnitude. We revisit the Poisson assumption, by studying a combination of historical traces and new measurements obtained from a major backbone link belonging to a Tier 1 ISP. At multi-second scales, we find a distinctive piecewise-linear non-stationarity, together with evidence of long range dependence. Combining our observations across both time scales leads to a non-stationary Poisson characterization of network traffic that, when viewed across very long time scales, exhibits the observed long-range dependence. This traffic characterization reconciles the seemingly contradicting observations of Poisson and long-memory traffic characteristics. Our model also supports recent theoretical models for large-scale traffic aggregation.

Measuring the elusive peer-to-peer traffic. (IMC [21] Globecom [20]) The current peer-to-peer or P2P protocols (think gnutella, napster or kazaa) disguise their traffic (e.g. using random ports), which makes their identification non-trivial. We develop a novel methodology to identify P2P traffic. The challenge and novelty of the work is that we rely only on connection-level patterns of P2P networks, and we do not need to examine the data part of the packet.

To the best of our knowledge, this is the first method for characterizing P2P traffic using only knowledge of network dynamics and not looking at the data of the packet. We rely on intrinsic characteristics of the P2P behavior and thus we were able to detect P2P protocols that we did not know existed, such as a P2P protocol popular in asia. It is important to not have to rely on examining the payload of the packet, since this is not always an option for a variety of reasons (privacy issues, storage capabilities etc). In experiments with real data, we find that our method manages to identify P2P very successfully.

Using our method, we find that the P2P seems to be slightly increasing percentage of the total traffic. This comes in sharp contrast to reports in the popular media that suggest a significant decrease in P2P traffic. The difference lies in our ability to detect the volume of P2P traffic more accurately.

Modeling the policy relationships of the Internet. (INFOCOM [24]) We develop a unique tool, NEMECIS, which can become a crucial aid to network administrators to: a) avoid configuration errors, and b) detect abnormal routing behavior due to error or malice. Our tool operates at the BGP (Border Gateway Protocol) level, namely at the interaction between Autonomous Systems (think Internet Service Providers).

More specifically, the goal of NEMECIS is to automate the validation of the routing policy using all possible sources of information: (a) the Internet Routing Registry databases and (b) BGP routing table information. Note that our tool is the first one to analyze and utilize the IRR database in an exhaustive and comprehensive way. So far, this information has not been used for this purpose, because it is written in a complicated language and it requires non-trivial processing to extract the useful information. Our tool parses, cleans the data (checking consistency and completeness) and develops a logical model with information which cannot be found in any other way currently.

Multicast Routing: Modeling, Quality of Service and Reliability. (SIGCOMM [7] ToN [26] IMC [4] Globecom [14] NGC [15] Network [3]) We have developed protocols to provide one-to-many communication support over the Internet. Our work spans several aspects of multicast routing such as modeling the spatial behavior of its users [4], and mechanisms to provide reliability in a scalable way [3]. It is worth mentioning our multicast routing protocol, QoSMIC, which is a receiver driven decentralized routing protocol that can support multicast application that require high Quality of Service [7] [26]. QoSMIC attracted significant attention from the community as can be seen by its citations in the citeseer site.

Theoretical analysis: multicast and Minimum Spanning Trees. (PODC [10] DC [11] SIROCCO [12] IJFCS [13]) We have also conducted theoretical work on proving bounds for the performance of existing multicast algorithms on directed graphs instead of the undirected that most previous work focused on [12] [13]. In addition, we have identified and corrected an algorithm for identifying the Minimum Spanning Tree of a graph in optimal time and message complexity [10] [11].

2 Wireless Ad hoc Networks

A massively scalable routing protocol for ad hoc networks. (INFOCOM [6] IPTPS [5]) We propose a radical routing architecture for ad hoc networks, which we call **Dynamic Address Routing (DART)**. DART has several desired properties: a) it is self-organizing, b) it requires minimal configuration, and c) it is finely distributed thus eliminating single points of failure and bottlenecks.

DART follows a radically different architecture than current approaches. Its novelty is captured in the following three characteristics: (a) separation of address from identity, (b) dynamic addressing, and (c) proactive routing.

The strength of DART lies in the efficient integration of these three elements. As nodes move around, their addresses change so that nodes with "similar" addresses are topologically close Note that proximity here is from a routing point of view, in number of hops, and not geographically. The immediate effect of this is that routes can be aggregated, reducing the amount of routing state necessary at each node. Preliminary results indicate that DART requires per-node state on the order of $O(\log_2 N)$ where N is the number of nodes in the network, in contrast to $O(N)$ state required by previous proactive protocols.

The Grand Vision. The DART approach has the potential to revolutionize the architecture of future networks. We claim that the *address equals identity* assumption in current ad hoc routing protocols is most likely inherited from the wireline world. By freeing our design from this constraint, we are able to start from scratch and design an architecture that considers and addresses mobility in scalable way.

Group Communications in ad hoc networks. (ICNP [2] MedHocNet [16]) We propose protocols and analyze the performance of group communications. We study both broadcasting (one to all) and multicasting (one to many) communication paradigms. More specifically, we propose a power-efficient broadcasting protocol [2], where the power level is adjusted in order to minimize the total energy consumed. We also propose a receiver-driven overlay multicast protocol, which outperforms previously proposed IP-layer and overlay multicast schemes [16].

TCP performance and Denial of Service Attacks. (Networking [17] MASS [1] Globecom [23] Milcom [25]) We propose mechanisms and study the capability of ad hoc networks to deliver high performance in the face of mobility, long multi-hop paths or Denial of Service attacks.

3 Cellular Networks

Modeling hotspots and heterogeneity. (INFOCOM [18] WCMC [19]) We study and quantify the effects of hotspots in wireless cellular networks. Most studies so far assume homogeneous networks, which do not necessarily reflect real scenarios. Our work provides the first theoretical and practical basis for a systematic approach to heterogeneous and more realistic models and simulations.

A hotspot is an area of the network where the bandwidth resources available are not enough to sustain the needs of the users. As a result, there is a higher probability that users are blocked or dropped in a hotspot. Hotspots can be the result of different causes, and we identify three distinct types of hotspots: (a) capacity based, created by uneven capacity of the cells, (b) delay based, created by uneven user delay across network areas, and (c) preferential mobility based, created by the biased mobility of users. We model and study all three types of hotspots. We show that different types create diverse effects on the network performance. Thus, not only introducing hotspots is necessary, but the way we model hotspots is important.

We develop a fluid flow model and an analytical model to study hotspots. The fluid flow model is surprisingly simple yet effective in helping us understand hotspots and their properties. We also describe an analytical model in which we consider a cell as an M/M/B/B queue. We conduct extensive simulations which validate the accuracy and effectiveness of our theoretical models.

References

[1] Dan Berger, Zhenqiang Ye, Prasun Sinha, Srikanth Krishnamurthy, Michalis Faloutsos, and Satish Tripathi. Tcp friendly medium access control for ad-hoc wireless networks: Alleviating self contention. In *IEEE International Conference on Network Protocols (ICNP)*, 2003.

[2] X. Chen, M. Faloutsos, and S. V. Krishnamurthy. Power adaptive broadcasting with local information in ad hoc networks. In *IEEE International Conference on Network Protocols (ICNP)*, 2003.

[3] Jun-Hong Cui, Michalis Faloutsos, and Mario Gerla. An architecture for scalable, efficient and fast fault-tolerant multicast provisioning. *IEEE Network special issue on Protection, Restoration, and Disaster Recovery*, 18(2):26–34, March 2004.

[4] Jun-Hong Cui, Michalis Faloutsos, Dario Maggiorini, Mario Gerla, and Khaled Boussetta. Measuring and modelling the group membership in the internet. In *ACM SIGCOMM/USENIX Internet Measurement Conference (IMC 2004)*, October, 2003.

[5] J. Eriksson, M. Faloutsos, and S. V. Krishnamurthy. Peernet: Pushing peer-2-peer down the stack. In *International Peer-To-Peer Symposium (IPTPS)*, 2003.

[6] J. Eriksson, M. Faloutsos, and S. V. Krishnamurthy. Scalable ad hoc routing: The case for dynamic addressing. In *IEEE INFOCOM, Hong Kong*, 2004.

[7] M. Faloutsos, A. Banerjea, and R. Pankaj. QoSMIC: a QoS Multicast Internet protoCol. *ACM SIGCOMM*, Sep 2-4, Vancouver BC 1998.

[8] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. *ACM SIGCOMM*, pages 251–262, Sep 1-3, Cambridge MA, 1999.

[9] M. Faloutsos, P. Faloutsos, C. Faloutsos, and G. Siganos. Power-laws of the Internet topology. *IEEE/ACM Trans. on Networking*, 1(4):514–524, 2003.

[10] M. Faloutsos and M. Molle. Optimal distributed algorithm for minimum spanning trees revisited. *Proceedings of 1995 Principles Of Distributed Computing (PODC)*, 1995.

[11] M. Faloutsos and M. Molle. A linear-time optimal-message distributed algorithm for minimum spanning trees. *Distributed Computing*, 17(2):151–170, August 2004.

[12] M. Faloutsos, R. Pankaj, and K. C. Sevcik. Bounds for the on-line multicast problem in directed graphs. *Proceedings of 4th International Colloquium on Structural Information and Communication Complexity (SIROCCO '97), Monte Verita', Ascona, Switzerland July 24-26*, pages 81–98, 1997.

[13] M. Faloutsos, R. Pankaj, and K. C. Sevcik. The effect of asymmetry on the on-line multicast routing problem. *International Journal on Foundations in Computers Science*, 13(6):889–910, 2002.

[14] A. Fei, J. Cui, M. Gerla, and M. Faloutsos. Aggregated multicast: an approach to reduce multicast state. In *IEEE GLOBECOM Global Internet Symp., San Antonio*, 2001.

[15] A. Fei, J. Cui, M. Gerla, and M. Faloutsos. Aggregated multicast with inter-group sharing. In *Networked Group Communications (NGC), London*, November, 2001.

[16] M. Ge, S. V. Krishnamurthy, and M. Faloutsos. Overlay multicasting for ad hoc networks. In *The Third Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net 2004), Turkey*, 2004.

[17] V. Gupta, S. V. Krishnamurthy, and M. Faloutsos. Improving the performance of tcp in the presence of interacting udp flows in ad hoc networks. In *IFIP Networking 2004, Athens, Greece*, 2004.

[18] J. Jobin, M. Faloutsos, S. K. Tripathi, and S. V. Krishnamurthy. Modeling hotspots in wireless cellular network simulation. In *IEEE INFOCOM, Hong Kong*, 2004.

[19] J. Jobin, S. Tripathi, M. Faloutsos, and S. Gokhale. Using statistical data for reliable mobile communications. *The Journal of Wireless Communications and Mobile Computing*, Special Issue on Reliable Transport Protocols for Mobile Computing, 2001.

[20] T. Karagiannis, A. Broido, N. Brownlee, kc claffy, and M. Faloutsos. Is p2p dying or just hiding? In *IEEE Globecom 2004 - Global Internet and Next Generation Networks*, Dallas, November, 2004.

[21] T. Karagiannis, A. Broido, M. Faloutsos, and kc claffy. Transport layer identification of p2p traffic. In *ACM SIGCOMM/USENIX Internet Measurement Conference (IMC 2004)*, Italy, October, 2004.

[22] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido. A Nonstationary Poisson View of Internet Traffic. In *IEEE INFOCOM*, Hong Kong, 2004.

[23] S. Kopparty, S. V. Krishnamurthy, M. Faloutsos, and S. Tripathi. Split-tcp for mobile ad hoc networks. In *IEEE GLOBECOM, Taipei 2002.*, 2002.

[24] Georgos Siganos and Michalis Faloutsos. Analyzing BGP policies: Methodology and tool. *IEEE INFOCOM*, 2004.

[25] S.V. Krishnamurthy V. Gupta and M. Faloutsos. Denial of Service Attacks at the MAC Layer in Wireless Ad Hoc Networks. In *Proc. IEEE MILCOM*, 2002.

[26] S. Yan, M. Faloutsos, and A. Banerjea. Qos-aware multicast routing for the Internet: The design and evaluation of qosmic. *IEEE/ACM Transactions on Networking*, 10(1):54–66, Feb 2002.