Instructor : Prof. Kihong Park
Class : TR 1:30–2:45 (CS G066)
Office Hours : TBA
E-mail : `park@cs.purdue.edu` (Tel. : 494–7821)
Teaching Assistants : TBA
Course Homepage : `http://www.cs.purdue.edu/~park/cs690N.html`

Course Content : Advanced graduate course in computer networks with focus on multimedia communication and high-speed network issues.
Prerequisites : CS 536 (Computer Networks), CS 503 (Operating Systems), or equivalents.
Text Books : TBA
Work Load and Grading Policy : Homework assignments (30 %), midterm (20 %), final (25 %), project (25 %).

Lab : Our main programming/implementation platform will be the Xinu Lab. A secondary platform will be specialized equipment in the Network Systems Lab (`http://www.cs.purdue.edu/~park/nsl.html`) for real-time MPEG video/audio compression, high quality video/audio capture, ATM switches and IP-over-ATM networking, programmable high-end IP routers, and a testbed for quality of service (QoS) experiments.

Course Outline :

Part 0: Overview

- Introduction to multimedia communication and high-speed networks

Part I: Payload Characteristics and Requirements

- Real-time traffic—video, audio, voice, interactive data—Web workload/user profiles, their characteristics and QoS requirements
- MPEG I, II, IV, VII standards, multi-resolution compression/representation, forward error correction, streaming and synchronization

Part II: Basic Traffic Control and Scheduling

- Fundamentals of closed-loop/open-loop traffic control, congestion and stability, end-to-end and per-hop implementations, impact of burstiness
- Real-time end system scheduling, EDF/RMS, packet GPS and weighted fair queueing, priority scheduling, RED

Part III: Network QoS Scheduling and WAN Architecture

- Provisioning guaranteed services: traffic shaping/policing, admission control, Int-Serv and RSVP signalling, statistical multiplexing, application to real-time video/audio transport
- End-to-end QoS control for pseudo real-time services: dynamic prefetching and buffer management,

static and adaptive forward error correction, Internet radio, RTP framing and synchronization

• Provisioning elastic services: aggregate flow control, distributed QoS control, Diff-Serv, aggregate RSVP signalling, SBS, scalable next generation Internet architectures

• QoS routing, reliable multicast, receiver-oriented control

Part IV: New Traffic Control Dimensions—User Behavior, Pricing, Time Scale

• Congestion pricing, usage/incentive pricing, selfish user behavior, noncooperative protocols, greedy TCP, network games, resource accounting, and Internet economics

• Multiple time scale congestion control, long-range/short-range dependence, protocol layering and interactions, heavy-tails, long/short duration connection management

Part V: Application Layer Control and Systems

• Web caching, Web server scheduling, user profile/workload modeling, bidding systems, QoS translation, integrated scheduling

• Scalable network management, reactive network security, secure QoS, fault-tolerance

• Testbeds and experimental systems: INDEX (Internet Demand Experiment), SBS (Stratified Best-effort Service), IP telephony, vBNS, prototype Diff-Serv systems

Synopsis of the Course :

This is an advanced graduate-level networking course which aims to expose the student to well-developed, yet "post-textbook," material in multimedia communication and high-speed networks. The course is organized into five parts.

In Part I, we will study the diverse traffic characteristics and requirements of both real-time and non-real-time payloads such as video, audio, voice, interactive data, and Web workloads and their access profiles as facilitated by a wealth of recent Internet traffic measurements. We will look at the widely deployed ISO compression/data representation standards and their error resilience, streaming and synchronization properties.

In Part II, we will study the fundamentals of traffic control and network resource scheduling, which includes real-time end system scheduling needed for integrated end-to-end quality of service (QoS) control.

Part III has two principal components, one, the mechanisms and algorithms underlying a wide area network (WAN) architecture which is able to provide guaranteed services with stringent QoS properties, and two, a scalable architecture for provisioning QoS-sensitive services, albeit with soft QoS properties in place of imputable ones. These two components, also referred to as Int-Serv and Diff-Serv under the IETF umbrella, are expected to determine the architecture of the next generation Internet. Part III is augmented by end-to-end QoS control techniques that can be applied over the current best-effort IP Internet as well as next generation Int-Serv/Diff-Serv internetworks, followed by QoS routing and reliable multicast.

In Part IV, we study the impact of three new traffic control dimensions—user behavior and its reflection on transport protocols, pricing, and time scale—which are expected to dominate the design and implementation of current and future traffic control protocols. Pricing, for example, is being investigated as a congestion control mechanism and tested by ISPs in private testbeds, and traffic control at multiple time scales—existing and newly proposed ones—has surfaced as an important dimension to effective traffic management.

Part V will study application layer controls falling under the "middleware" umbrella which includes Web caching, Web server design and task scheduling, inter-layer QoS translation, network management and security with emphasis on scalability, and fault-tolerance/reliability. We will conclude with a discussion of testbeds, new architectures, and experimental systems in deployment.