CIS 500: SOFTWARE FOUNDATIONS

Prerequisite: Undergraduate-level course in programming languages or compilers; significant programming experience (CIT 591 or equivalent).
This course introduces basic concepts and techniques in the foundational study of programming languages. The central theme is the view of individual programs and whole languages as mathematical objects about which precise claims may be made and proved. Particular topics include operational techniques for formal definition of language features, type systems and type safety properties, polymorphism and subtyping, foundations of object-oriented programming, and mechanisms supporting information hiding and programming in the large.

CIS 501: COMPUTER ARCHITECTURE

Prerequisite: Knowledge of computer organization and basic programming skills.
This course is an introductory graduate course on computer architecture with an emphasis on a quantitative approach to cost/performance design tradeoffs. The course covers the fundamentals of classical and modern uniprocessor design: performance and cost issues, instruction sets, pipelining, superscalar, out-of-order, and speculative execution mechanisms, caches, physical memory, virtual memory, and I/O. Other topics include: static scheduling, VLIW and EPIC, software speculation, long (SIMD) and short (multimedia) vector execution, multithreading, and an introduction to shared memory multiprocessors.

CIS 502: ANALYSIS OF ALGORITHMS

Prerequisite: CIT 594 or equivalent.
An investigation of several major algorithms and their uses in areas including list manipulation, sorting, searching, and graph manipulation. Efficiency and complexity of algorithms. Complexity classes.

CIS 505: SOFTWARE SYSTEMS

Prerequisite: Undergraduate-level knowledge of Operating Systems and Networking, programming experience (CIT 594 or equivalent).
This course introduces basic concepts and techniques in advanced software systems for first year graduate students in computer science. It provides the students with a background in the design, the implementation and the analysis of experimental systems. The course will focus on distributed systems - systems that distribute state and computation across networked elements. The first part of the course introduces the basics of network and kernel support for building distributed systems. The second part of the course covers the key concepts of interprocess communication and coordination, such as logical clocks and remote procedure call. The third part of the course covers case studies of distributed systems. Students will be expected to design, program and analyze software systems.

CIS 511: THEORY OF COMPUTATION

Prerequisite: Basic notions of discrete algebra.
Finite automata (deterministic and nondeterministic) regular graphs, regular expressions, regular grammars, (Nerode congruence), the ``pumping lemma'', closure properties. Context-free languages. Standard forms: removal of e-rules, chain rules, reduced grammars. Chomsky Normal Form. Context-free languages as fixed points (Ginsburg and Rose's Theorem). Greibach Normal Form (using Rosenkrantz matrix method). Ogden's Lemma and the``pumping lemma.'' Pushdown automata (PDA's). Equivalence of PDA's and context-free grammars. Brief sketch of top-down and bottom-up (nondeterministic) parsing. Deterministic PDA's. Closure properties. Partial recursive functions. Turing machines and RAM programs. Primitive recursion. Minimization. Equivalence of the models. Church/Turing's thesis. Acceptable Codings. A Universal RAM program. Undecidability of the halting problem. Recursively enumerable sets (RE sets).

CIS 520: INTRODUCTION TO ARTIFICIAL INTELLIGENCE

*Prerequisite*s: Elementary probability, calculus, and linear algebra. Basic programming experience.
This course will provide a survey of mathematical methods in artificial intelligence and machine learning. Topics will include: probabilistic reasoning with graphical models; autonomous agents, reinforcement learning, and game theory; statistical methods for prediction, clustering, and dimensionality reduction; and applications to vision, robotics, speech, and natural language processing.

CIS 550: DATABASE AND INFORMATION SYSTEMS

Prerequisite: CIT 591 or equivalent.
Introduction to the theory and practice of database management systems. The Entity-Relationship approach as a modeling tool. The relational model, algebra and calculus. Commercial systems: SQL, Quel and Ingres. Database design and relational normalization. Physical data organization and indexing structures. Updates and integrity: transaction management, concurrency control and recovery techniques. Logic as a data model: Datalog and evaluation techniques. The network model and object oriented approaches.

CIS 555: INTERNET AND WEB SYSTEMS

Prerequisites: CSE 380 (Operating Systems), CIS 505 (Software systems), or equivalent; CSE 330 (Database Management Systems), CIS 550 (Database and Information Systems), or equivalent; proficiency in Java programming.
This course will require a significant amount of programming and will require the ability to work with your classmates in teams. This course focuses on the issues encountered in building Internet and web systems: scalability, interoperability (of data and code), atomicity and consistency models, replication, and location of resources, services, and data. We will examine how XML standards enable information exchange; how web services and other communications schemes support cross-platform interoperability; how caching, replication, and hierarchy are used in distributed environments; and how agreement and transactions are

addressed in the distributed context. We will study techniques for locating machines, resources, and data (including ranked web search, publish/subscribe systems, directories, and peer-to-peer protocols). We will also examine the ideas that have been proposed for tomorrow's Web, including the "Semantic Web," and see some of the challenges, research directions, and potential pitfalls. This project has a significant project-based component, in order to provide hands-on experience with the ideas and algorithms discussed. Students will construct and validate a large-scale distributed system.

OPIM-661: SYSTEMS ANALYSIS, DESIGN, AND IMPLEMENTATION

This course is about the problem of designing, developing, and managing large, complex systems. This problem is regularly faced by nearly all managers and other professionals throughout their careers. The purpose of the course is twofold. First, we aim to familiarize students with the issues involved in conceiving, designing, building, and maintaining the kinds of large-scale, complex information systems now required for commercial (and governmental) purposes. Second, we aim to provide students with experience working with the main tools and techniques in systems analysis, design, and implementation. Special focus will be given to modern object-oriented design methodologies, and methodologies appropriate for Internet-based electronic commerce (e.g., UML and modern CASE tools).

OPIM-672: DECISION SUPPORT SYSTEMS

With the advances in computers, databases, communications and Internet technologies, modern organizations nowadays collect massive amounts of data on about everything we do – what we buy, the phone calls we make, payment records, financial transactions, loan applications and others. Analyzing data on this scale and converting it into knowledge to affect decision making, presents exciting new challenges. This course discusses DSS for business applications with a focus on data mining modeling and analysis. Data mining is a new

generation of computerized technologies for discovering knowledge hidden in large amounts of data. It is an interdisciplinary field lying at the interface of Statistics, Computer Science and Artificial Intelligence. The application areas of data mining cut across many industries, including marketing, financial, retail, telecommunication, and in fact any data-rich industry. Business applications include direct and database marketing, CRM, risk management, fraud detection, web mining, recommendation systems and many others. The goal of this course is to bridge the gap between theory and practice and empower the students to apply data mining solutions in real world applications. This course is practically-oriented, emphasizing the knowledge discovery process, economical considerations and implementation issues. Real world assignments, involving realistic data, provide the students with hands-on experience in using data mining in practice.

In today's complex business environment and tough competition, understanding the power of data mining is an asset that is certain to open up more challenging career opportunities for the Wharton graduates.

ENM 503. INTRODUCTION TO PROBABILITY AND STATISTICS

Prerequisite: MATH 240 or equivalent.

Introduction to probability. Expectation. Variance. Covariance. Joint probability. Moment generating functions. Stochastic models and applications. Markov chains. Renewal processes. Queuing models. Statistical inference. Linear regression. Computational probability. Discrete-event simulation.

- - - - - - -

Complete CIS course descriptions can be found at http://www.cis.upenn.edu/grad/course-descrip.shtml.

More information about the Operations Information Management Program courses can be found at http://opim.wharton.upenn.edu/index.php.