EECS Area II (Computer Science) Open House
Wednesday, September 27, 2006

3:00pm: Seth Teller, Area II Chair: Welcome, Introductory remarks

3:05pm-4:15pm: Student Presentations:

Alex Gruenstein: Web-Based Multimodal Dialogue Systems
Abstract: Spoken dialogue systems allow users to interact with computers using spoken natural language. In this presentation, I'll describe recent work on developing a scalable and portable framework for enabling spoken dialogue interaction as a web-based application. Specifically, I'll demonstrate a web-based spoken dialogue system which allows users to search a database of restaurants in several major U.S. metropolitan areas. The system can easily be expanded to new metropolitan areas, since all proper nouns are stored in dynamic language model classes. Finally, I'll show how speech recognition accuracy can be improved by incorporating contextual information about the current dialogue state into the language model, especially when there is limited data available to train the language model.

Ben Snyder: Extracting Information from Text
Abstract: With the advent of the internet, textual data has become abundant in a wide variety of domains, including news reports, blogs, and customer reviews. One goal of Natural Language Processing is to automatically analyze the information content of such bodies of text. In this talk, I will mention two attempts to perform this sort of content extraction and analysis: (a) automatically constructing sports databases from news accounts and (b) building opinion databases from product reviews.

Yajun Fang: Time To Contact Estimation for Intelligent Vehicle
Abstract: Automatically detecting potential collision is an important function for intelligent transportation system, which makes it necessary to estimate the "time to contact(TTC)" fast and accurately. Time to contact is the time until two objects or surfaces moving relative to one another will touch or collide if they continue in their current trajectory. In typical applications, the time to contact equals the ratio of the distance between two objects or their surfaces to the relative velocity of motion of one with respect to the other. We propose a method to determine the time to contact using spatial and temporal derivatives of brightness in time-varying images or image sequences obtained from an imaging system. The method of this invention uses the time-varying image or an image sequence from an imaging system, does not require careful calibration of the optical system, and is computationally efficient, which is vital for fast response of intelligent vehicle applications. The current experiments show promising initial results.

Calvin Newport: Virtual Infrastructure for Mobile Ad Hoc Networks
Abstract: The nature of ad hoc networks makes it challenging to solve the standard problems encountered in mobile computing, such as location management, or global coordination, using classical tools. The difficulties arise from the lack of a fixed infrastructure to serve as the backbone of the network. In this project, we are developing several new approaches that allow existing distributed algorithms to be adapted for highly dynamic ad hoc environments. These approaches take advantage of geographic information to implement high-level abstract objects

that facilitate the design of algorithms for this challenge context. For example, one approach is to allow the mobile nodes to coordinate in the emulation of arbitrary virtual state machines rooted at known locations.

Our work on this project spans from low-level algorithms that tame communication uncertainty at the network layer, to the middle ware emulation algorithms that provide our abstraction layer, to the high-level applications that run on top of this layer. We have also developed a real world implementation of these systems which can be deployed on wireless devices or within a custom-built packet-level simulator. More at **http://theory.csail.mit.edu/tds/vi-project/index.html .**

Ce Liu: Motion -- What Computers Can See and What We Can See
Abstract: In this 5-minute talk I will demonstrate motion magnification (Siggraph2005) and contour motion analysis (NIPS 2006). The human visual system overlooks some subtle motions in video sequence. We designed a motion magnification system to magnify and render small motions so that they can be easily perceived and diagnosed. On the other hand, it is difficult for classical computer vision systems to correctly analyze the motion of textureless objects under occlusions. We designed a contour motion analysis system to group boundaries with both spatial and temporal cues, obtaining motion as humans perceive it.

Chris Batten: The Vector-Thread Architecture
Abstract: Embedded digital systems have traditionally been characterized by their fixed functional requirements and their stringent power and price constraints. To reduce power or cost, designers were often willing to accept lower performance or decreased flexibility. The embedded landscape, however, has changed drastically over the last decade as embedded devices have become more pervasive and embedded applications have become more sophisticated. While there remains a significant number of simple micro-controller based embedded systems, there is a new breed of high-performance embedded devices which must flexibly perform several different functions while still meeting strict power and price constraints.

To address these needs, we have developed a novel architectural paradigm called vector-threading (VT) which unifies the vector and multi-threaded compute models. The VT abstraction provides the programmer with a control processor and a vector of virtual processors (VPs). The control processor can use vector-fetch commands to broadcast instructions to all the VPs or each VP can use thread-fetches to direct its own control flow. A seamless intermixing of the vector and threaded control mechanisms allows a VT architecture to flexibly and compactly encode application parallelism and locality, and a VT machine exploits these to improve performance and efficiency. We are developing a prototype processor, Scale, which is an instantiation of the vector-thread architectural paradigm optimized for low-power and high-performance embedded systems. We have evaluated the Scale processor using detailed simulation of a broad range of embedded applications and have shown that its performance is competitive with larger and more complex embedded processors. We are now in the last stages of taping out a VLSI implementation of the Scale processor in a TSMC 0.18um process.

Emma Brunskill: LittleDog: Robust Robot Locomotion for Rough Terrain
Abstract: Wheeled robots such as the Mars rovers have already proved their use in many applications. But there remain many places which are easy for humans to reach but where it is

extremely difficult for wheeled robots to go. For such places walking robots are necessary. The Tedrake and Roy labs have been collaborating to create a robust walking quadruped that can quickly navigate rough terrain as part of the DARPA LittleDog project. We hope the results of this research will lead to better walking robots that can even run and jump.

Evdokia Nikolova: From Stochastic Shortest Paths to Nonconvex optimization and Games
Abstract: For the last couple of years I have been doing research on the complexity and algorithms for stochastic shortest paths, also known as route planning under uncertainty. The general problem area has a wide variety of applications (such as route planning in transportation networks as well as more abstractly task planning in robotics, etc) and has been extensively studied yet a natural and very useful version of it has received little attention. In particular, I employ a decision theoretic framework for defining the optimal route: for a given source S and destination T in the graph, I seek an ST-path of lowest expected cost where the edge travel times are random variables and the cost is a nonlinear function of total travel time. Although this is a natural model for route-planning on real-world road networks, results are sparse due to the analytic difficulty of finding closed form expressions for the expected cost, as well as the computational/combinatorial difficulty of efficiently finding an optimal path which minimizes the expected cost.

I will briefly mention a recent result on stochastic shortest paths for the objective of arriving on time, and its implications to the much more general field of nonconvex optimization. If time permits, I will also mention some work on shortest path auctions and prediction markets.

Jake Beal: Programming Spatial Computers
Abstract: Space-filling computers are an emerging problem across a wide range of disciplines---sensor networks, biofilms, FPGAs, morphogenesis, swarm robotics, etc. A typical spatial computer is composed of a vast number of unreliable parts, and we need programming tools that can marshal them effectively to produce robust aggregate behavior. The amorphous medium abstraction enables this by breaking the task into three loosely coupled subproblems: global programs for a continuous space, a global to local compiler, and emulation of continuous space by a discrete network. We have implemented these ideas with our language Proto, and have begun applying it in the areas of sensor networks and reconfigurable robotics, both in simulation and on Mica2 Motes.

http://people.csail.mit.edu/jrb/stp/stpg.htm
http://www-swiss.ai.mit.edu/projects/amorphous/

Jeremy Fineman: Cache-Oblivious Dictionaries with Fast Insertions
Abstract: Many modern applications stream huge amounts of data to a searchable database (or "dictionary") on disk, later performing searches to random records. Since the number N of data items is so large that they cannot simultaneously fit into main memory, we use the external-memory model to analyze data structures. In this model, we attempt to minimize the number of "block transfers" or I/O operations from disk to memory as opposed to the number of CPU instructions performed. The B-tree, a commonly used dictionary, achieves provably good performance under the assumption that the number of searches is large. In particular, B-trees have good (but not optimal) insertion cost and asymptotically optimal search cost. Many of these streaming applications, however, perform far more insertions than searches, making the B-

tree suboptimal as a whole. Moreover, the efficiency of B-trees and other "cache-aware" data structures depends on a priori knowledge about the "block size," which is nontrivial to determine in complex memory hierarchies. We like cache-oblivious data structures as a provably good alternative.

We have designed a data structure, called the cache-oblivious lookahead array, which has provably good performance, especially when the number of insertions far exceeds the number of searches. In particularly, the data structure supports insertions in O((1/B) log_2(N)) block transfers and searches in O(log_2(N)) block transfers, where B is the number of records in a block. Note that the insert cost of O((1/B) log_2(N)) is less than 1 for reasonable values of B, so this bound is only achievable in an amortized sense. The cache-oblivious lookahead array achieves these bounds even though the block size B is not known or used by the data structure. In contrast, B-trees support insertions and searches in O(log_B(N)) block transfers, and the B-tree does rely on knowledge of the block size B.

In a cache-aware setting, it is possible to tune for a tradeoff between insertion and search costs. In particular, for any constant $0 < c < 1$, it is possible to achieve O((1/B^c) log_B(N))-cost insertions and O(log_B(N))-cost searches, where c affects the constants hidden by the big-O notation. We are currently investigating whether there is a cache-oblivious data structure that achieves similar tradeoffs.

Olivier Koch: Wide-Area 3D Tracking from Omnidirectional Video and 3D Structure
Abstract: We present a method for real-time vision-based localization in indoor environments. Given a rough model of the structure and a video sequence captured from an omnivision camera, the system computes the camera pose (translation and rotation) in the structure coordinate frame. The system has several novel aspects: it performs localization in 3D; it handles highly cluttered and dynamic environments; it scales well over wide-space areas made of several buildings and it runs for long periods without breaking. We demonstrate that the localization problem can be split into two distinct problems: an initialization phase and a maintenance phase. In the initialization phase, the system determines the camera pose with no other information than the estimate provided by the user (building, floor, area, room). In the maintenance phase, the system keeps track of the camera pose from frame to frame without any user interaction.

Stephen McCamant: Quantitative Information-Flow Tracking for Real Systems
Abstract: I'll describe a technique that instruments software to track, at the level of individual bits, whether the public outputs it computes depend on confidential inputs. The technique is quantitative, computing a bound on the number of bits of information revealed so far, and is compatible with type-unsafe languages such as C and C++. Using an implementation based on the Valgrind debugging framework, I'll give some results from applying the tool to a real bug.

Steven Bauer: Applied Economic Engineering of Large-Scale Network Architectures
Abstract: Our research deals with building a better Internet. Specifically, we explore how the architectural and protocol design process of networking research can accommodate conflicting economic incentives of network participants. We consider both economic theory, such as mechanism design and game theory, and architectural and engineering methodologies that address the problems posed by economic tussles. We advance an analytic and design

methodology for incorporating economic incentives into the engineering of large-scale networked systems.

We apply our methodology to a problem steeped in conflicting incentives -- improving Internet connectivity. This problem goes by many other names as well in 2006 -- the network neutrality debates, the broadband incentive problem, etc. -- but the essence of the problem is one of conflicting incentives between the network participants that carry the traffic and network participants that send and receive the traffic.

We apply the methodology we are advocating to analyze the origins and make specific predictions about evolution of the incentive tussles over Internet connectivity. To validate our own predictions and inform both the research process and larger societal debate about the evolution of Internet connectivity, we design, build, and deploy a system for monitoring the evolution of Internet connectivity over time. This system provides empirical evidence that informs both the research and regulatory communities that are shaping the evolution and improvements in Internet connectivity**.**

Yeuhi Abe: Animation of Dynamic Object Manipulation
Abstract: In interactive animations, it is important that the animated characters respond realistically to the objects they manipulate. Characters should react to unexpected collisions and strain when wielding heavy or fast moving objects. Standard kinematic techniques employed to animate characters are generally incapable of producing such responses. Our approach is to physically simulate and control characters, thereby automatically generating physical effects. So far, we have shown how a control theory from robotics can be used to coordinate simple motion descriptions with recorded posture data to produce lifelike animations of characters manipulating objects. In time, we hope to expand our technique to encompass an entire repertoire of fully interactive controllers for characters within physical simulation.

Soonmin Bae: Two-scale Tone Management for Photographic Look
Abstract: We introduce a new approach to tone management for photographs. Whereas traditional tone-mapping operators target a neutral and faithful rendition of the input image, we explore pictorial looks by controlling visual qualities such as the tonal balance and the amount of detail. Our method is based on a two-scale non-linear decomposition of an image. We modify the different layers based on their histograms and introduce a technique that controls the spatial variation of detail. We introduce a Poisson correction
that prevents potential gradient reversal and preserves detail. In addition to directly controlling the parameters, the user can transfer the look of a model photograph to the picture being edited.

Jennifer Carlisle: Tavarua: A Mobile Telemedicine System
Abstract: Tavarua is a multimedia streaming system that leverages network-striping to deliver relatively high bit rate video over present-day cellular wireless wide-area networks. The Tavarua system achieves this by building on our previously developed flexible network-striping middleware. This paper describes a motivating mobile telemedicine application, and the design of the Tavarua system. It also describes experiments in which our initial Tavarua implementation was used to stripe video over multiple 3G cellular-phones from different providers.

~4:15pm: Informal Discussion

Reminder: new Area II students should create an account on the Area II website:

http://area2.csail.mit.edu/ !

Thanks to **Joanne Talbot Hanley, Area II secretrary,** for assistance assembling this program.