ANNUAL REPORT

FOUNDATIONS OF HYBRID AND EMBEDDED SYSTEMS AND SOFTWARE

NSF/ITR PROJECT – AWARD NUMBER: CCR-0225610

UNIVERSITY OF CALIFORNIA, BERKELEY

November 27, 2010

PERIOD OF PERFORMANCE COVERED: JUNE 1, 2009 – August 31, 2010

Contents

1 Participants

1.1 People

PRINCIPAL INVESTIGATORS:
- THOMAS HENZINGER (UC BERKELEY, EECS)
- EDWARD A. LEE (UC BERKELEY, EECS)
- ALBERTO SANGIOVANNI-VINCENTELLI (UC BERKELEY, EECS)
- SHANKAR SASTRY (UC BERKELEY, EECS)
- CARLO SEQUIN (UC BERKELEY, EECS)
- CLAIRE TOMLIN (UC BERKELEY, EECS)

POST DOCTORAL RESEARCHER:
- ALVARO CÁRDENAS (UC BERKELEY)

GRADUATE STUDENTS:
- ANIL ASWANI (UC BERKELEY, PROF. TOMLIN)
- JAN BIERMEYER (UC BERKELEY, PROF. SASTRY)
- YOUNG-HWAN CHANG (SUMMER) (UC BERKELEY, PROF. TOMLIN)
- MILOS DREZGIC (UC BERKELEY, PROF. SASTRY)
- SLEIMAN ITANI (UC BERKELEY, PROF. TOMLIN)
- TODD TEMPLETON (UC BERKELEY, PROF. SASTRY)
- INSOON YANG (UC BERKELEY, PROF. TOMLIN)

STAFF:
- NEAL MASTER (UC BERKELEY, ASSISTANT II - TOMLIN)
- ZHAOQIAN XIA (UC BERKELEY, ASSISTANT IV - TOMLIN)
- JESSICA GAMBLE (UC BERKELEY, ASSISTANT II - SASTRY & TOMLIN)
- MARY P STEWART (UC BERKELEY, PROGRAMMER - SASTRY & TOMLIN)

BUSINESS ADMINISTRATOR:
- GLADYS KHOURY (UC BERKELEY)

EXECUTIVE DIRECTOR:
- CHRISTOPHER BROOKS (UC BERKELEY)

1.2 Partner Organizations:

UNIVERSITY OF CALIFORNIA, BERKELEY

1.3 Collaborators:

SAURABH AMIN (UNIVERSITY OF CALIFORNIA, BERKELEY)
CHRISTIAN BERGER (SOFTWARE ENGINEERING, RWTH AACHEN)
PETER BICKEL (UC BERKELEY)
MARK BIGGIN (LAWRENCE BERKELEY NATIONAL LABORATORY)
JAMES BROWN (UC BERKELEY)
JOHN CHUANG (UC BERKELEY)
CHARLES C. FOWLKES (UC IRVINE)
CHRISTOPHER GEYER (UC BERKELEY, CMU, IROBOT)
JOEL GRAY (LAWRENCE BERKELEY NATIONAL LABORATORY)
JENS GROSSKLAGS (PRINCETON UNIVERSITY)
JAY GROVES (UC BERKELEY, CHEMISTRY)
CHRIS HOOFNAGLE (UC BERKELEY)
YU-LUN HUANG (NATIONAL CHIAO TONG UNIVERSITY, TAIWAN
SOILE VE KERANEN (LBNL)
DAVID W. KNOWLES (LBNL)
JAMES KORKOLA (LAWRENCE BERKELEY NATIONAL LABORATORY)
ZONG-SYUN LIN (NATIONAL CHIAO TONG UNIVERSITY, TAIWAN
GARRY P. NOLAN (STANFORD)
TANYA ROOSTA (UC BERKELEY)
BERNHARD RUMPE (SOFTWARE ENGINEERING, RWTH AACHEN)
KAREN SACHS (STANFORD)
PAUL SPELLMAN (LAWRENCE BERKELEY NATIONAL LABORATORY)
HSIN-YI TSAI (NATIONAL CHIAO TONG UNIVERSITY, TAIWAN
NICKOLAS WANG (LAWRENCE BERKELEY NATIONAL LABORATORY)
QIAN XU (UC BERKELEY, CHEMISTRY)

2 Activities and Findings

2.1 Project Activities

This is the eighth Annual Report for the NSF Large ITR on "Foundations of Hybrid and Embedded Systems and Software." This year was a no-cost extension for certain researchers at the University of California, Berkeley (Center for Hybrid and Embedded Systems and Software (CHESS), http://chess.eecs.berkeley.edu. Research at the other CHESS partners: ISIS at Vanderbilt University (Institute for Software Integrated Systems, and the Department of Mathematical Sciences, at the University of Memphis ended before the period covered by this report.

The web address for the overall ITR project is:

http://chess.eecs.berkeley.edu/projects/ITR/main.htm

This web site has links to the proposal and statement of work for the project.

2.1.1 ITR Events

Main events for the ITR project in its eighth year were:

- A weekly Chess seminar was held at Berkeley. The speakers and topics are listed in Section 4.6.1, presentations for the seminar are available at http://chess.eecs.berkeley.edu/seminar.htm

We organize this section by thrust areas that we established in the statement of work. As year eight was a no-cost extension, we include only thrust areas funded by the no-cost extension.

2.1.2 Hybrid Systems Theory

Locally-linear system identification: a new statistical method based on the estimation of the exterior derivative

As part of our system biology efforts, we have developed models for pattern formation in Drosophila embryos, based on significant prior knowledge about the system being modeled. This prior knowledge is used to formulate a specific functional form of the equations with free tuning parameters that are picked using regression procedures in order to match experimental data. The advantage of this approach is that it provides models that are highly interpretable, with specific terms of the equations corresponding to specific phenomenon in the system. The disadvantage is that it does require extensive knowledge about the system.

An alternative form of system identification is the class of nonparametric techniques; they are sometimes used in the robotics and control communities. The advantage of these techniques is that they are well-suited for situations in which there is little or no prior knowledge about the system being modeled. The disadvantage of nonparametric techniques is that they are statistically not as well behaved as parametric techniques. We have developed a new statistical technique for nonparametric system identification that has better statistical performance than existing techniques.

The advantage of requiring limited prior knowledge is important in biological applications, because there is often little or unreliable knowledge about the large networks being modeled. This makes nonparametric methods particularly useful for network identification problems in biology. In such problems, we are interested in learning not only parameter values of the model, but also the interconnections between different parts of the network. In fact, for some of the drug multitarget selection tools discussed in the second part of this thesis, it is more important to know the interconnections than to know the parameter values. From an informal perspective, systems described well by ordinary differential equations (ODEs) can be nonparametrically identified by matching trajectory data of the system with an ODE whose functional form is given by a series expansion.

Local linearization techniques are an important class of nonparametric system identification. The prevalence of control tools which utilize piecewise-affine models or linear models makes this class of techniques important. Nonparametric identification techniques involve solving regression problems of the form $Y = X\beta$ where Y is a vector of noisy response variables, X is a matrix of noisy predictor variables, and β is a vector of regression coefficients. Identifying β can be challenging due to collinearity of the predictors (which can come about if the system evolves on a manifold with dimension lower than that of the state space). Of techniques that explicitly consider manifold structure, none of these can take advantage of any sparsity in the system, that is, systems which have an upper bound on the number of non-zero entries in the linearization. Dynamical systems found in engineering problems are typically sparse because they are an interconnection of preexisting, engineered components. Biological systems (especially regulatory networks) are often sparse as well because the organization of such networks is known to be hierarchical. It is important to develop identification techniques that take advantage of this special structure.

We interpret collinearities in the language of manifolds, and this provides the two contributions of our work. This interpretation allows us to develop a new method to do regression in the presence of collinearities or near-collinearities. This insight also allows us to provide a novel interpretation of regression coefficients when there is significant collinearity of the predictors. On a statistical level, our idea is to learn the manifold formed by the predictors and then use this to regularize the regression problem. This form of regularization is informed by the ideas of manifold geometry and the exterior derivative. Our idea is to learn the manifold either locally (in the case of a local, nonlinear manifold) or globally (in the case of a global, linear manifold). The regression estimator is posed as a least-squares problem with an additional term which penalizes for the regression vector lying in directions perpendicular to the manifold. Our manifold interpretation provides a new interpretation of the regression coefficients. The gradient describes how the function changes as each predictor is changed independently of other predictors. This is impossible to do when there is collinearity of the predictors, and the gradient does not exist. The exterior derivative of a function tells us how the function value changes as a predictor and its collinear terms are simultaneously changed, and it has applications in control engineering, physics and mathematics. In particular, most of our current work is in high-dimensional system identification for biological and control engineering systems. We interpret the regression coefficients in the presence of collinearities as the exterior derivative of the function. The exterior derivative interpretation is useful because it says that the regression coefficients only give derivative information in the directions parallel to the manifold, and the regression coefficients do not give any derivative information in the directions perpendicular to the manifold. If we restrict ourselves to computing regression coefficients for only the directions parallel to the manifold, then the regression coefficients are unique and they are uniquely given by the exterior derivative. This is not entirely a new interpretation. Similar geometric interpretations are found in the literature, but our interpretation is novel because of two main reasons. The first is that it is the first time the geometry is interpreted in the manifold context, and this is important for many application domains. The other reason is that this interpretation allows us to show that existing regularization techniques are really estimates of the exterior derivative, and this has important implications for the interpretation of estimates calculated by existing techniques. Our regularization scheme can improve estimation error, and it can be easily modified to include lasso-type regularization.

This work is presented in [1], which has been accepted to the Annals of Statistics.

Topology-based Control

Control theory has traditionally focused on a core group of goals: to stabilize a plant, to improve plant performance, to robustify a plant, to track a reference, or to perform motion planning. In engineering systems, these goals have been achieved through an analysis-design flow; this flow is rarely linear: there is often a need to go back to previous steps and incorporate things that were missed in earlier attempts. Beginning with design specifications, we write mathematical models for the engineering system, analyze these models, and devise a controller. We also implement the controller on actual hardware.

The traditional control scheme has been to input a signal into a plant, using either an open-loop or a closed-loop controller. Such a control strategy is possible if the plant is able to accept inputs or can be modified to do so. However, this situation is not always true in biological genetic networks; in these systems, there is often no input or obvious modification to allow inputs. Instead of inputs, genetic networks are more easily influenced through large-scale modifications. Genetic networks are different from traditional engineering systems and require a new paradigm for control. It is often easier to change the topology of a genetic network than it is to either change the states or elements of the network. For instance, a state could be the concentration of a protein within a cell, something which is difficult to affect to within any order of precision. Additionally, it is sometimes difficult or not feasible to modify or insert pathways by adding elements. Thus, for genetic networks it is important to develop a theory of control based on making largescale changes (e.g. genetic changes or pharmaceutical drugs) to the topology of the genetic network. Fundamentally, medical treatments seek to change how a cell operates and go beyond modifying the cellular environment.

Genetic networks can be modified in a variety of ways. The most basic is the use of pharmaceutical drugs, many of which prevent certain reactions from occurring or remove a state from a network. Biotechnology techniques allow for the insertion of genetic material into bacteria, and are commonly used for alternative energy and pharmaceutical applications.

In another technique, the genetic material of a virus is replaced with useful, genetic material. Next, the host is infected with the virus, and this inserts the useful, genetic material into the host. This control technique is being studied for use in pharmaceutical applications such as cystic fibrosis. Biologists continue to develop new techniques, amongst which include the use of microRNA and single interfering RNA. Though many of these techniques are established and used in practice, there is a lack of a systematic theory or methodology to determine which modifications to make or what to target with pharmaceutical drugs. Biological research often involves the use of intuition or trial-and-error to determine which changes are or are not beneficial for the purposes of controlling a biological system. This chapter proposes the idea of abstracting the effect of pharmaceutical drugs as modifying the topology of the biological network, and it also proposes how this abstraction might be used to do control by identifying drug targets.

Piecewise-affine (PWA) hybrid systems and ordinary differential equation (ODE) models of biological systems are considered in this work. Two different types of models are used for reasons of analysis: The simpler, hybrid systems models are easier to analyze for global behavior, and the more detailed, ODE models are easier to analyze for local behavior of small components of the network. Controllers using ODE theory are defined and analyzed, and these approaches are used to analyze and build a controller for the p53 pathway: a pathway that is related to cancer. This topological control changes the topology of the network by applying a pharmaceutical drug or other chemical, and the topology remains changed only in the presence of this pharmaceutical. As soon as it degrades away, the topology of the network goes back to an uncontrolled, unchanged state. Since the control is topological, it is crucial to have a correctly identified network. The approach described is unable to deal with latent variables that are unidentified, because the presence of latent variables can drastically change the behavior of the system.

This work is presented in [2] in the IEEE Transactions on Automatic Control, as well as in [3] and [4].

2.1.3 Robust Hybrid Systems

This year we worked on the use of hybrid systems tools to understand some issues in quantum computing [5].

2.1.4 Hybrid Systems and Systems Biology

The CHESS ITR enabled significant progress in Systems Biology in 2009-2010. The existing collaboration between Tomlin's lab and Mark Biggin's group in early Drosophila development at Lawrence Berkeley Labs was strengthened, with new results and several publications. A new collaboration between Tomlin's Lab and Joe Gray's group in breast cancer at Lawrence Berkeley Labs was begun, launching several new modeling efforts. CHESS provided critical funding for the start of this cancer modeling work, for which in mid-2010 new funding from the National Cancer Institute was awarded.

Statistical System Identification in Drosophila embryogenesis

Inferring regulatory networks in animals is challenging because of the large number of genes and the presence of redundant and indirect interactions. To build the highest quality models, it will be necessary to use multiple data sets, including: gene expression, genome wide binding, and network perturbation data. However, combining multiple data types to infer regulatory networks is still an open problem. An intermediate problem is to use only gene expression data to infer regulatory networks. The relationships between the expression levels of transcription factors and target genes are used to predict which genes are regulatory. While much work has been done in this area, it is critical to understand the maximum amount of information that can be obtained about the network using this strategy.

Typical approaches for inferring regulatory networks have been to assume a model formulation and have then fitted the data to this formulation (parametric system identification). Many models have been proposed, including coexpression networks, information-theoretic, regression onto dynamical systems and graphical models (including Bayesian networks). The primary differences between these models lie in the trade-off between statistical and interpretational issues. Techniques like Bayesian networks, graphical models, and information theoretic models have protections against over-fitting (i.e., fitting models with many parameters to a small amount of experimental data); however, these techniques do not provide dynamical models which can generate new biological insights. On the other hand, techniques such as nonlinear regression networks and regression onto dynamical systems provide more biologically interpretable models, but sometimes suffer from inaccurate assumptions or over-fitting of the model to the data. There is disagreement on the necessity of dynamical as opposed to static models. Dynamical models are more philosophically pleasing because regulatory networks contain temporal characteristics: For example, a protein binds to DNA and initiates transcription, which eventually leads to transport of the mature mRNA to the cytoplasm. Yet the argument is often made that static models provide a quasi-steady-state interpretation of the network that may provide a sufficient approximation. Rigorous comparison of the two approaches is lacking. Dynamical modeling of animal regulatory networks has a long history and is a powerful approach in which researchers hypothesize a set of nonlinear, differential equations to describe the network, but it requires significant prior knowledge about the network. If there is insufficient biological knowledge about the network, then the structure of the equations can be incorrectly chosen. And if the model is not carefully chosen, it will have a large number of parameters, possibly leading to weak biological effects being erroneously identified as strong effects. Furthermore, it is sometimes shown that a wide range of different parameter values can reproduce the biological behavior of the network, which could be taken as evidence for either network robustness or over-fitting.

We have applied the local linearization system identification described in the previous section to a network found in Drosophila embryos. The particular network which is studied here is the formation of eve mRNA stripes during stage 5 of embryogenesis, and this work was in collaboration with Mark Biggin and the Berkeley Drosophila Transcription Network Project (BDTNP). The dynamic model that we designed was shown to be significantly better than static models, and new biological insights about the developmental system were derived.

Compared to other dynamical methods, our approach requires minimal information about the mathematical structure of the ODE; it does not use qualitative descriptions of interactions within the network; and it employs our new statistical method to protect against over-fitting. It generates spatio-temporal maps of factor activity, highlighting the times and spatial locations at which different regulators might affect target gene expression levels. We identify an ODE model for eve mRNA pattern formation in the Drosophila melanogaster blastoderm and show that this reproduces the experimental patterns well. Compared to a non-dynamic, spatial-correlation model, our ODE gives 59% better agreement to the experimentally measured pattern. Our model suggests that protein factors frequently have the potential to behave as both an activator and inhibitor on the same cis-regulatory module depending on the factors' concentration, and implies different modes of activation and repression.

This work has been written up and appears in [6], in BMC Bioinformatics.

Optimization-based inference for temporally evolving networks with applications to biology

In this work, we design an optimization-based inference scheme to identify temporally evolving (Boolean) network representations of genetic networks from data. In the formulation of the optimization problem, we use an adjacency map as a priori information, and define a cost function which both drives the connectivity of the graph to match biological data as well as generates a sparse and robust network at corresponding time intervals. The main idea of our scheme lies in representing the captured relationship as a network path with a priori information (a given connectivity map) and using convex optimization techniques to find the time-varying sparsest graph consistent with experimental observations. Despite uncertainties about details for a given biological system, we often have reasonable qualitative knowledge about interactions of each gene, so we can use this information as a priori information. In this setting, the model behavior is solely based on this qualitative information which guarantees biologically reasonable behavior: robustness and sparsity in general. The ability of many biological networks to exhibit their function reliably despite noise or perturbation is often referred as functional robustness. We also note that biological regulatory networks are likely to be sparse especially at a fixed interval of time (for example, most transcription factors (TFs) do not regulate most genes). Also, there are expectations behind modeling efforts:

- Networks represent the structure of complex connections so viewing evolving networks as dynamical systems allows us to predict many of their properties analytically.

- If we can match signal propagation that drives the placement of links and nodes, then the topology or the structural elements will follow. It can help to move beyond architecture and uncover the laws that govern the underlying dynamic process.

In contrast to previous methodologies for dynamic graph analysis, we develop a convex optimization-based inference method, where we embed the dynamics of a linear time varying representation, and enforce sparsity at corresponding time intervals. We have applied the method to an example of biological network of HER2 over-expressed breast cancer.

This work has been written up [7] and is submitted to the 2011 American Control Conference.

An ODE Model for the HER2/3-AKT Signaling Pathway in Cancers that Overexpress HER2

During the 2009-2010 period, we designed new algorithms for recovering the structure of signaling pathways by using the power of perturbations. We created three versions of a dynamic model for the HER-AKT pathway in cancers that overexpress HER2. The different versions correspond to different PI3K mutants. The models showed some predictive ability and they inspired a treatment scheme that is being tested.

The Human Epidermal Growth Factor Receptor (HER) family activates signaling pathways which control crucial cellular processes such as cell division, motility and survival, and are among the best studied of all signaling networks]. HER family members are implicated in a wide variety of human cancers and are frequent targets of pharmaceutical intervention. Individual HER family members have diverse roles and may be specifically targeted in cancer therapies.

The HER2 receptor is overexpressed in 20-30% of breast cancers, in addition to some lung and other cancers. HER3 is also frequently overexpressed in cancers that overexpress HER1 or HER2. Increased expression of HER3 synergistically increases the potency of HER2 in the cancer, and the loss of HER3 removes the transforming ability of HER2. HER3 is in fact a necessary partner in HER2-mediated transformation. Additionally, cancers that overexpress HER2 have been found to be more aggressive than similar cancers that don't. Molecularly, activation of HER receptors leads to the activation of tyrosine kinases, enzymes which activate signaling molecules via phosphorylation, enabling propagation of signal along the signaling cascade. One of the targets of the HER family receptors is AKT, a well-known proliferation agent. The inhibition of AKT allows the protein BAD to be activated, and leads to cell apoptosis. This is why the selective inactivation of AKT in cancer cells is a desired therapeutic outcome.

A relatively new class of therapies for cancer is targeted drug treatments. Among these are selective tyrosine kinase inhibitors (TKIs), which have proven more effective than traditional chemotherapy in some cancers. However, their application to HER signaling has demonstrated limited anti-tumor activity, in spite of effective receptor activity inhibition. This inconsistency was resolved in an elegant study by Sergina et al (UCSF), in which it was demonstrated that HER3 and downstream pathway effectors PI(3)K/AKT evade inhibition by current HER-family TKIs due to a compensatory shift in the HER3 phosphorylation-dephosphorylation equilibrium. This shift is driven by increased membrane HER3 expression driving the phosphorylation reaction and by reduced HER3 phosphatase activity impeding the dephosphorylation reaction. These processes are driven by an AKT-mediated negative-feedback loop.

From a different point of view, the problem of creating a drug to treat a certain type of cancer can be thought of as a control engineering problem, in which the goal is to steer the state of a cancerous cells to healthy or dead. This point of view enforces the need for mathematical modeling for the behavior of cancer cells. Recently, a number of dynamical systems models of the HER2/3 pathway in various cell types has been introduced, enabling the formulation of nonintuitive hypotheses and setting the stage for potential novel therapeutics. Mathematical models for signaling pathways and cells are becoming more popular, because of the large amounts of data being produced by the molecular biology community and the great promises this point of view holds.

In this project, we use a compendium of known interactions of this pathway, as well as the data from the Sergina et al study, to formulate a dynamic model of HER2 and HER3 signaling in breast cancer. We aim to model the HER2/3 pathway in cancers that overexpress HER2, using data from the BT747 and SKBr3 cell lines. We create a model with ordinary differential equations (ODEs) that uses first and second order mass action kinetics and accounts for vesicular trafficking between cellular compartments. In addition to our tissue focus, our model is distinguished by its relatively large scope in terms of the number of protein species included, as well as the time span (two days, versus a time span of up to 2 hours in previous models). The extended time span is crucial to our model predictions and therapeutic application, as anti-tumor effects fail at longer time spans, even when indications appear favorable in the initial hours post drug application. With a model that accounts for extended time behavior of cancers upon treatment with TKIs, we hope to have an understanding that allows us to design better treatments for cancers that overexpress HER2.

Although the model employs first order approximations and represents a simplification of the underlying biological system, within the constraints of the measured experiment set currently available, it demonstrably captures the system behavior over time, successfully generalizes to predict unseen training data, and enables the formulation of testable hypotheses for therapeutic applications.

The data available to fit the parameters of the model is scarce at this point, and thus it might be expected that a large class of models would fit the data. Of course, this is only true if the network has a valid structure and enough degrees of freedom. Other properties of the system can be studied by making sure that they're robust over the irrelevant parameters. The idea is that if all of the models that fit the data show a certain behavior, the true model could be expected to exhibit that behavior. Additionally, a basic model can be used to determine the experiments needed to refine the model.

This work has been written up and appears in [8].

2.1.5 Embedded Software for National and Homeland Security

Autonomous Ground Vehicles

Ground vehicle research, although more stable than airborne vehicles, presents the subtle problem of providing intelligent behavior which operates in real-time, executes safely, and yet provides a "smooth" reaction to stimuli—i.e., the software behavior is somewhat humanized. Our application is the DARPA Urban Challenge, and we are using the ground vehicle testbed to show the performance of various algorithms and advancements in theory of switched systems, model-predictive control, parameter identification, time-triggered distributed components, and computer vision, as well as how these components work in real-time with one another. Additional concerns which this project addresses are distributed software testing, component-based design, and vehicle/sensor health monitoring. We are proving many of the theories and algorithms which have been developed in the last four years of the ITR.

Real-time Computer Vision

The goal of this task is to detect moving objects using a stereo camera system mounted on a car. This information will be used to detect, and estimate the trajectories of (possibly) mobile obstacles such as other vehicles. Note that, in this scenario, we wish to detect both objects in close proximity to our vehicle (for example, the vehicle in front of us), and also objects that are farther away (for example, oncoming vehicles), so we cannot make assumptions about whether height errors at adjacent pixels appear to be close to planar. Our solution must also be able to run in real time. In this application, we are interested in determining 3D motion in the scene, rather than the more-studied case of 2D motion in the image. In particular, we want to find the motion of objects in the scene relative to each other; if we have vehicle state data available, or if we can assume that the majority of the scene is not moving, we can determine which motion-segmented region corresponds to the background and remove its perceived motion from the perceived motion of the other objects, obtaining the motion of the other objects in a global frame.

Todd Templeton's focus in the past year has primarily been hardware and software infrastructure for the autonomous / assisted-driving vehicle testbed. This effort has culminated in external time synchronization for all car sensors, which is essential for perception algorithms that utilize multiple sensors, as well as a reusable and extensible library of communication and sensing software components that can run on the variety of hardware platforms carried in the car. In addition to a complement of four half-2U mini-ITX Linux PCs, the car carries Coldfire-based NetBurner embedded processor boards for sensors that do not support external triggering and time-synchronization; the cameras are the only sensors on the car that natively support an external trigger. A hardware interrupt on each microprocessor board is wired to the same (square wave) trigger signal as the cameras, through a solid-state relay for voltage conversion. The master external trigger (generated by a network-controllable signal generator in the car) is enabled after all of the microprocessor boards are powered on, and after the camera driver has initialized the cameras. The first falling edge of the trigger signal causes the first frame to be captured by each camera, and also provides a time-zero reference to the microprocessor boards. Subsequent falling edges of the trigger signal cause subsequent frames to be captured by the cameras, and also provide a time reference from which the microprocessor boards calibrate their internal clocks. Once the master trigger is enabled and a microprocessor board has performed the initial calibration of its clock (after receiving the second falling edge), it begins time-stamping data from the sensor (such as an INS or LADAR) to which it is attached and sending it over the car's Ethernet network using the Spread multicast messaging service. All perception algorithms, which run on the Linux PCs, use the time-stamp on each piece of sensor data, instead of the current system time, in their calculations.

The embedded Coldfire processors in the car present a challenge to software portability. To this end, Templeton has developed an architecture compatibility library that emulates the required subset of the POSIX C standard library on top of the embedded OS, embedded C library, and NetBurner system libraries. This compatibility library also provides standard I/O over a network socket, and functionality that runs before the program's main() to initialize the system clock against the external trigger.

The software collection for this platform, hopefully to be expanded to other robotic platforms within the group and open-sourced to the larger robotics community, is called the Intelligent Robotics Toolkit (IRT). It is primarily based on software used in a previous vision-based helicopter mapping and landing project, expanded to support distributed computing across a diverse hardware and software environment. We currently have a software engineering doctoral student from RWTH Aachen University in Germany helping us to strengthen the software engineering foundations of this toolkit, and to expand its simulation and visualization capabilities.

Recently, we have achieved time synchronization / triggering across all car sensors, and taken the car out on the roads near Richmond Field Station to capture sensor data from which to formulate and tune perception algorithms.

Templeton graduated in December 2009: in his final contributions, he focused on the problem of moving object segmentation while driving, using camera data collected from the above-mentioned data-gathering trip. Moving object detection and segmentation is important for both autonomous and assisted driving for two primary reasons: to estimate the trajectories and future positions of mobile obstacles, and to remove points on mobile obstacles from the map of the static portion of the scene (without which mobile obstacles become 3D walls as they move across the scene). Hence, the motion planning process becomes similar to that in a static environment (using the static scene map), with the addition of a list of mobile obstacles and their estimated trajectories and current positions.

The proposed motion segmentation algorithm is non-parametric in that it does not assume a particular shape or appearance of mobile obstacles, which allows it to detect moving obstacles such as bicyclists and pedestrians, and to not detect stationary objects that look like cars (such as parked cars)—these static objects will be handled like any other static part of the scene by the motion planner.

Verification of Driver Augmentation Systems

There has been an explosive growth in the use of embedded systems in cars. By some accounts, it is widely expected that by the end of the decade over 50% of the cost of a car will be vehicular electronics (or veitronics). Beyond the fundamental functionality of an automobile, such as driving, stopping, and turning, a modern car provides additional features for more passenger safety, better comfort, and lower environmental impact. Thus, the number of embedded processors used in a vehicle today is in the order of 80 for luxury cars, and is expected to grow further by some accounts. A contemporary car therefore is a networked embedded system, in which subsystems need to process data and communicate over special networks such as CAN, LIN or Flexray, often within hard real-time constraints. While there has been a great deal of attention paid to hybrid and electric cars and the veitronics for their drive trains, it is our contention that Advanced Driving Assistance Systems (denoted as ADAS) present a big opportunity to apply research that can potentially

have a big impact on the whole automotive industry. A key issue with the introduction of these safety critical technologies is the need to have them be verified, validated and in some cases certified. By verification, we mean formal results guaranteeing the performance (properties such as safety, liveness or non-blocking stability) of models of intelligent software systems embedded with the physical hardware on a car; and by validation we mean testing of the theoretical verified proofs of performance on hardware. Certification usually follows formal specifications laid down by regulatory/ insurance authorities and is accompanied by verification/ validation.
Automotive OEMs and tier 1 suppliers are currently suffering from the high cost of verification and validation of the kinds of features that are being demanded by customers. While the software itself to be used in automotive systems has to undergo rigorous processes for design, implementation, audit and test, using standards such as IEEE-610 or SEI CMMI, limited tools exist for verification of automotive ADAS. These systems provide an additional order of complexity, since they are not only semi-autonomous, but also interaction with a human operator for safety critical decisions is necessary.
The state-of-the art to the extend of our knowledge is a statistical approach: High-precision driving robots or automated mini-series cars such as the VW Golf GTi "53 plus 1" are used to rigorously repeat driving situations within sub-centimeter precision. If statistical measures such as MTTF or MTBF suggest safe operation, the system is deemed deployable in road traffic.
However, there are several problems with such an approach: Foremost, repeated execution with varying initialization parameters is only feasible for limited cases, such as a Parking Assistant. Further, the testing has to be rigorous and thus is costly in monetary terms and drastically increases time-to-market.
It is therefore our strong belief, that a formal verification method for ADAS is needed, to increase car functionality, safety and energy efficiency, as well as decrease development costs and time to market. We are thus developing a new theoretical framework and algorithms for the design of mathematically verified and validated cyber-human systems from beginning with ADAS as a specific case. We are exploring the next generation of cyber-human systems, which will employ much "better models of human cognition", and hence, better assist humans in autonomous or semi-autonomous ways. This research is enabled by recent major advances in the sub-areas of computing and communication, such as stochastic hybrid models, learning methods, signal-to-symbol transformation, distributed decision making and dynamic resource allocation in geographically distributed systems without communications infrastructure.
Over the last 40 years, the fields of knowledge representation, perception, robotics, control, and learning have evolved in their separate ways. The grand vision of cyber-human systems, in which these fields combine into complete agents, has all but disappeared. However, with new results and partial reunifications of these individual sub- disciplines, the time has come to propose a reconstruction of the science of cyber-human systems. The focus of the present research is developing integrated agent designs capable of performing simple tasks reliably in unstructured environments and generating purposive activity over an unbounded period.

Such a goal requires dealing with perceptual input, noisy, and partially known dynamics, real-time requirements, and complex environments containing many objects and agents.
The goal is to demonstrate our work on two examples: Automatic Parking and Highway Cruise Control with linear and lateral track, as well as negative and positive acceleration, authority. We chose the parking example, since several luxury cars already deploy such systems, and our approach can readily be compared with existing solutions. The highway example, however, is on the road map of all major OEMs, yet we argue that existing verification approaches will fail here. This is where our major contribution will be.
Automated Driver Automation Systems (ADAS) are required to perform safety-critical tasks with hard-real time guarantees. Simulation and validation based methods alone cannot guarantee that all specifications are met and that undesired behaviors are not executed. Formal methods offer a rigorous framework to prove that a mathematical abstraction of embedded software satisfies the required correctness properties. The mathematical abstraction of ADAS embedded software can be completely behavioral in that it closely describes all possible executions or it could be simply prescriptive in that it only specifies what the system should do. In both cases, the mathematical abstraction should specify evolution of the physical environment of the system such as the effect of forces, actions of the driver, presence of external entities like obstructions, other vehicles etc. In addition, the ADAS models should also include description of maneuvers such as cruise mode, lane changing, parking, and overtaking. The specification of required correctness properties can be generally translated into safety, i.e., the system never performs a bad execution; and liveness, i.e., the system eventually performs a good execution and does not deadlock.
The modeling of ADAS requires hierarchies of different models of computation, some discrete and others continuous. Such system models are referred to as hybrid system models. These models are especially suited to modeling compound behaviors arising from composition and interaction between heterogeneous sub-systems in automotive systems. A prototypical example is the hierarchical layering of finite state machines and nonlinear continuous time differential equations.
However, modeling uncertainty associated with human behavior in verification of ADAS raises new issues and challenges: First, the human-ADAS interface determines the mode and extent of information the user has about the behavior of the system. Secondly, the modes of interaction that are not proven to be correct can raise important safety issues of such systems. In particular, the human-machine interaction can lead to unpredictable behaviors if the assumption of human user does not match with the guarantee that the ADAS can provide. As an example, if the human user gives control of a safety- critical task to the ADAS but assumes a wrong model of the system's behavior, the ADAS may not be able to successfully execute the task. Lastly, such systems also raise an important modeling question: Can human cognitive limits be reasonably modeled and incorporated in ADAS verification framework? We plan to address these research challenges by formalizing probabilistic specifications of human-ADAS interaction. We are inspired by developments in modern cognitive science which holds the viewpoint that human mind is a computational system. Prior research has also demonstrated that significant headway can be made in

understanding how a human's cognitive resources can be integrated most effectively into the problem when mental models of the human, and information about the human's cognitive limits, are incorporated into the design of the control architecture in the very early stages of design specification. We will build a framework in which successive refinements of models of human cognitive functions will be checked for correctness against formal models of ADAS. Our research will benefit from the efforts to model interaction between human operators and air- traffic management systems by accounting for the physical constraints that come from applications. Horvitz at Microsoft Research has pioneered the use of Bayesian networks in the design of models of human interaction with automation: NASA operators with varying levels of expertise, interacting with the fuel control systems on board the Space Shuttle. Using these models to indicate the context and timing of the information to display to the operator, the expected utility of the operator's decisions can be greatly enhanced.

In our approach to verification, we use a multi-world semantics hierarchical system. We break up the verification task into individual components, and verify each in its own, and its interface to neighboring components in the hierarchy. An expression at each level is interpreted at the same level. Therefore, checking the truth-claim at that level is performed in its own declarative world. Relating these disconnected worlds together is a nontrivial task; however, we are collaborating with Jonathan Sprinkle from University of Arizona on using metamodels for the interface design between components. This both enables better analyzability as well as potential for using components developed by different research groups. Instead of semantic flattening, where an expression has to be both syntactically translated and semantic interpreted, we are following human reasoning and cognition, and propose the use of high-level expressions, which will be compiled into idealized lower-level expressions and then interpreted. Invariants will be used to show that higher-level truth-claims now become conditional lower level truth- claims. In this setting, higher-level truth-claims become necessary conditions. This is contrary to one-world semantics, in which higher- level claims are sufficient conditions.

Suppose one-world interpretation leads to falsification of higher- level claims that are true in multi-world semantics. This can happen, when lower-level faults are not accounted for in higher-level descriptions. Then, multi-world semantics must be split into multiple- frameworks, each dealing with the identified faults.

Our proposed framework is kept general enough to apply to various systems. From the sensory point-of-view, these levels have increasingly coarse representations of the world as one moves toward higher levels, while from the actuation point-of-view, the tasks become more strategic, supervisory and planning-oriented at the higher levels starting from "reactive" ones at lower levels.

The highest level, the Mission Planner, is the human operator in the automotive context. The output of this level could be in the form of "go to San Francisco". The Route Planner then takes this input and transforms it into an actual plan, in our context it would be an A* implementation detailing which roads to take. Next, the Maneuver Control layer takes these roads and transforms them into driving decisions, considering current sensor readings. From there, the Low- level Control transforms these driving decisions into actual control input for

the physical vehicle.
This model also has what we call in analogy to database design a "roll- back mechanism": If a component cannot perform a task given from a higher-level component, it has to communicate the reasons why it failed at execution. The higher level then has to roll-back: to revert to a more feasible strategy or pre-defined back-up plan.
There are several challenges associated with the verification of such hierarchical architecture. One has to establish such a hierarchy in a formal way with different models appropriate at the respective levels: define the semantics of interaction between different levels for operational and emergency goals, define assumptions about the model of interactions between the levels, the levels of service each level can provide, and define the set of services each level must accept in order to provide requisite services.
Within this hierarchy, we are also interested in knowing how uncertainty and fault propagates between different levels and their criticality to the goals.
Our mathematical framework for the hierarchy is currently under constant change. We defined a series of increasingly complex examples of (semi)-autonomous system that we try to capture in our framework, and apply our framework to these. If properties of the system cannot be described with the current definitions, these are altered or amended. Our examples range from the common water tank problem in hybrid systems theory via agents that have two simultaneous missions or objectives such as to traverse an area and to remain stealthy, to a collection of heterogeneous agents which each have unique sets of capabilities and have to work together towards a goal declared by the Mission Planner.
We are about to finalize this framework and intend to publish a detailed white paper in the upcoming Fall on our findings to invite feedback from the research community.
While the research effort described above is mainly aimed at verifying a (Semi-) autonomous system, an ADAS capable of making autonomous driving decisions also has to incorporate knowledge and assumptions about other traffic participants. We therefore propose the investigation into using an Algorithmic Game Theory (AGT) approach. While game theory itself is only descriptive, AGT however is prescriptive, and therefore could be used in this effort.
Clearly, interaction of cars (normally) cannot be described as an adversial game. However, it is also not really a collaborative effort, since one driver has no incentive to help another driver to achieve his goal. We are thus modeling traffic as an implicit collaboration:
Ideally, human drivers not only care for their own safety and goals, but also attempt to not inconvenience other traffic participants. Recent research conducted by our partners from the University of Aachen in co-operation with Volkswagen suggests that it suffices to use a 3x3 or 3x4 matrix to make informed LTT driving decisions. The car being the center in the matrix (X_{00} in its local neighborhood can be divided into segments to its sides, in front and behind, and on the diagonals. If there is any car detected in a segment, it's estimated speed is entered into the matrix. In the case of a sensor that can see further ahead than up to the next car (such as
Same advanced automotive radar), one column in driving direction extends the matrix. Each cell of the matrix contains the speed and the distance of a car, or Null if there is no car

detected.
To model the interaction between cars, we are using coordination games (CG). Coordination games model situations in which all agents can realize mutual gains as long as they make mutually consistent decisions. For such games, there typically exist several pure strategy Nash Equilibria, so that the agents have to chose the same or corresponding strategies. However, often the coordination has to be done without any ability to communicate. Finally, some equilibria may give higher payoffs, be more fair, or, most important to the application at hand, may be safer, thus occasionally leading to interest conflicts.
We argue that this fits the interaction of cars on a highway well. We do not want to force cars to be able to communicate with each other, for two major reasons: security and deployment. Communication between cars would have to be done wirelessly, and thus provides an easy point of attack for malicious parties. Also, for such a system, all traffic participants would have to be equipped with such communication facilities, rendering it unusable for the next decade. Further, just for the interaction of two cars and one possible disturbance, we can foresee several equilibria, and cars should settle for a common equilibrium that maximizes overall welfare.
In parallel to our verification efforts, we also created the Berkeley DRIVR Lab, in close collaboration with CHESS, with the intend to showcase successfully verified algorithms. After participating in the DARPA Urban Challenge, we published a paper in AAET 2008, outlining our ideas on how to apply lessons learned from autonomous driving towards ADAS. Discussions about our paper with other researchers in the field and OEMs led to our insight on the necessity of formal ADAS verification techniques. However, other lessons learned were on the need for a robust and scalable hardware platform, as well as test bed, and a lightweight software toolkit.
We have outfitted a fully automated 2008 Ford Escape Hybrid ByWire XGV as testbed for (semi-) autonomous driving Actuation is done transparently via a commercial toolkit (Torc Technologies) that directly communicates with the car's ECU. A cluster of customized quad-core mini-ITX computers is used to process data and make driving decisions, communicating among each other and with the car via Gigabit Ethernet. Various built-in screens and a KVM switch, as well as the possibility to connect notebook computers enables in-car debugging, while wireless network access enabled remote debugging, data streaming and offline analysis. All sensors, such as visual light cameras, thermal infrared cameras, laser scanners, radar, inertial measurement units and GPS are connected to time-synchronized embedded processors which trigger synchronized data-acquisition and timestamp data before relaying it to the computation cluster via Gigabit Ethernet for pre-processing, fusion and influencing driving decision making.
Since there is no robotics software framework available that satisfies all our demands for robustness, ease of use, reliability, platform independence, lightweight communication and interfaces to existing standards such as JAUS among others, we opted for implementation of our own.
Our software and computational approach is based on the following principles:

- Sensor drivers and communication methods must be abstracted as much as possible, to enable maximum flexibility in choosing different sensors and communication methods in the future.

- Transparent timestamping and synchronization facilities must be available at the sensor input level. To enable this, sensor drivers must be lightweight enough to run on an embedded processor; in particular, this means minimal or no threading.
- Transparent multicast, logging, and replaying facilities must be available at all levels (including the sensor-input level), in order to encourage the modularization of software, and to enable the independent development and verification of individual modules.
- Software modules must be able to be written in a variety of languages and run on a variety of platforms, with minimal invasiveness on the part of the infrastructure.
- Software modules should make optimal use of the computing hardware, from embedded processors via multi-core CPUs to manycore GPUS by utilizing frameworks such as Intel Thread Building Blocks, OpenMP and OpenCL or CUDA.
- Although the infrastructure is written in C/C++ due to its efficiency and ease of use, it is written with multi-platform compatibility in mind, and its interface is simple and can be easily wrapped for other programming languages, such as MATLAB.
- Where available, existing software components will be used, such as scientific and numerical libraries (BLAS, LAPACK), or vision libraries such OpenCV or the NASA Vision Workbench.

2.1.6 Control of Communication Networks

In a series of papers Abate and co-authors have continued to explore using stochastic hybrid systems congestion control schemes for both wired and wireless networks. These methods have tremendous applicability to other classes of network embedded systems as well (see).

2.2 ProjectFindings

Abstracts for key publications representing project findings during this reporting period, are provided here. A complete list of publications that appeared in print during this reporting period is given in Section 4 below, including publications representing findings that were reported in the previous annual report.

- [5]Milosh Drezgich, S. Shankar Sastry. 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, "Matrix Multiplicative Weights and Non-Zero Sum Games," 2010; Submitted.

  This article concentrates on the improvements in the vector and matrix version of a widely used multiplicative weights algorithm .The four results that we present are the following. We first present a small improvement in the existing upper bound for the matrix multiplicative weights algorithm for the particular set of two player zero sum games. Second, using the concavity property of unital positive linear maps between *-algebras, we establish the precise connection between the vector and matrix version of the multiplicative updates algorithm, that clearly reveals why the total loss for the both algorithms are the same. Third, as a corollary we present the matrix multiplicative weights algorithm with the iterative updates, that is both computationally less demanding and guarantees the same worst case performance as the conventional cumulative matrix multiplicative weight algorithm, resolving in part the open question stated at COLT 2010. Finally, unlike the vector version of the multiplicative weights algorithm, we present an evidence that the Nash equilibrium strategy for the non-zero sum Shapely game and the augmented Shapely game, can be found using matrix multiplicative weights updates algorithms.

- [3]Anil Aswani, Claire Tomlin. Unpublished article, "Computer-Aided Drug Discovery for Pathway and Genetic Diseases," 2010.

  Selecting drug targets in pathway and genetic diseases (e.g., cancer) is a difficult problem facing the medical field and pharmaceutical industry. Because of the complex interconnections and feedback found in biological pathways, it is difficult to understand the potential effects of targeting certain portions of the network. The pharmaceutical industry has avoided novel targets for drugs, largely because of the increased risk in developing such treatments. This necessitates the need for systems biology methods which can help mitigate some of the risks of identifying novel targets and also suggest further experiments to validate them. The primary goal of this paper is to introduce a mathematical framework for solving such problems, that is amenable to computational or mathematical study. The secondary goal is to suggest methods for solving problems posed in this framework. One of these methods is a heuristic which is designed to allow its computations to scale up to much bigger examples and pathways than presented here.

- [9]Milosh Drezgich, S. Shankar Sastry. Unpublished article, "On the NP in BQP," 2010.

One of the central questions of the theory of quantum computation is whether the quantum computers are able to solve, in polynomial time, the problems that are classically currently believed intractable, for example, unstructured search and 3SAT. While mostly believed unlikely, almost all current efforts toward the positive result have been utilizing the quantum adiabatic theorem in the design of the potential algorithm. We show that if the process of computation is governed by the system Hamiltonian, design of the prospective algorithm should not rely on the premise of adiabatic or ground state computation. In particular we prove that the prospective algorithm, represented by the dynamic system through the Schroedinger equation, does not fail simply because the system evolution did not satisfy adiabatic condition or drifted away from the ground state. As a corollary we both present the algorithm in $O(n^k)$ for the unique unstructured search and make conjecture on the algorithm for 3SAT.

- [10]J. Biermeyer, T. Templeton, C. Berger, H. Gonzalez, N. Naikal, B. Rumpe, S. S. Sastry. Proceedings des 11. Braunschweiger Symposiums "Automatisierungssysteme, Assistenzsysteme und eingebettete Systeme fr Transportmittel", ITS Niedersachsen, Braunschweig, "Rapid Integration and Calibration of New Sensors Using the Berkeley Aachen Robotics Toolkit (BART)," 2010.

  After the three DARPA Grand Challenge contests many groups around the world have continued to actively research and work toward an autonomous vehicle capable of accomplishing a mission in a given context (e.g. desert, city) while following a set of prescribed rules, but none has been completely successful in uncontrolled environments, a task that many people trivially fulfill every day. We believe that, together with improving the sensors used in cars and the artificial intelligence algorithms used to process the information, the community should focus on the systems engineering aspects of the problem, i.e. the limitations of the car (in terms of space, power, or heat dissipation) and the limitations of the software development cycle. This paper explores these issues and our experiences overcoming them.

- [1]Anil Aswani, Peter Bickel, Claire Tomlin. Annals of Statistics, "Regression on Manifolds: Estimation of the Exterior Derivative," 2010; To appear.

  Collinearity and near-collinearity of predictors cause difficulties when doing nonparametric regression on manifolds. In such scenar ios, variable selection becomes untenable because of mathematical difficulties concerning the existence and numerical stability of the regression coefficients. In addition, once computed, the regression coefficients are difficult to interpret, because a gradient does not exist for functions on manifolds. Fortunately, there is an extension of the gradient to functions on manifolds; this extension is known as the exterior derivative of a function. It is the natural quantity to estimate, because it is a mathematically well-defined quantity with a geometrical interpretation. We propose a set of novel estimators using a regularization scheme for the regression problem which considers the geometrical intuition of the exterior derivative. The advantage of this regularization scheme is that it allows us to add lasso-type regularization to the regression problem, which enables lasso-type regressions in the presence of collinearities. Finally, we consider the large p, small n problem in our context and show the consistency and variable selection abilities of our estimators.

- [11]Anil Aswani, Harendra Guturu, Claire Tomlin. "System identication of hunchback protein patterning in early Drosophila embryogenesis," Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference, 7723-7728, 18, December, 2009.

  Early patterning in the Drosophila melanogaster embryo occurs through a complicated network of interactions involving transcription factor proteins and the mRNA of their target genes. One such system is the pattern of hunchback mRNA and its regulation by Bicoid and Kruppel proteins. This system is well-studied, but there is disagreement amongst biologists on how exactly hunchback expression is regulated. We attempt here to provide evidence to distinguish between two models in contention, through system identification. Our general approach is to do nonlinear regression on a parametric, nonlinear partial differential equation model which incorporates transcription, diffusion, and degradation. We perform the regression, analyze the results and then interpret the results in the biological context. We also compare our results to previous work on this system.

- [12]Todd Templeton. PhD thesis, "Accurate Real-Time Reconstruction of Distant Scenes Using Computer Vision: The Recursive Multi-Frame Planar Parallax Algorithm," University of California, Berkeley, December, 2009.

  In this dissertation, we detail the Recursive Multi-Frame Planar Parallax (RMFPP) algorithm, a recursive extension of Irani et al.'s Multi-Frame Planar Parallax (MFPP) batch algorithm that allows real-time reconstruction of distant static scenes using computer vision, with expected error that increases only linearly with depth. We present an overview and comprehensive derivation of the theoretical foundation on which the RMFPP algorithm is built, including the seminal planar-parallax work by Sawhney. We derive a recursive cost function that preserves more of the problem's nonlinearity than does the cost function in the MFPP algorithm, which allows a more accurate recursive procedure. In order to obtain a recursive algorithm, we remove the geometry-refining optimization that is present in the MFPP algorithm; however, we empirically show that our algorithm degrades gracefully in the presence of geometric error. We present results using both synthetic and real imagery that show that the RMFPP algorithm is at least as accurate as the original MFPP batch algorithm in many circumstances, is preferred to both fixed- and dynamic baseline two-frame methods, and is suitable for real-time use.

- [13]YuLun Huang, Alvaro Cardenas, Saurabh Amin, Song-Zyun Lin , Hsin-Yi Tsai, S. Shankar Sastry. International Journal of Critical Infrastructure Protection, "Understanding the Physical and Economic Consequences of Attacks Against Control Systems.," 2(3):72-83, October 2009.

This paper describes an approach for developing threat models for attacks on control systems. These models are useful for analyzing the actions taken by an attacker who gains access to control system assets and for evaluating the effects of the attacker's actions on the physical process being controlled. The paper proposes models for integrity attacks and denial-of-service (DoS) attacks, and evaluates the physical and economic consequences of the attacks on a chemical reactor system. The analysis reveals two important points. First, a DoS attack does not have a significant effect when the reactor is in the steady state; however, combining the DoS attaack with a relatively innocuous integrity attack rapidly causes the reactor to move to an unsafe state. Second, an attack that seeks to increase the operational cost of the chemical reactor involves a radically different strategy than an attack on plant safety (i.e., one that seeks to shut down the reactor or cause an explosion).

- [14]Alvaro Cardenas, Saurabh Amin, Bruno Sinopoli, Annarita Giani, Adrian Perrig, S. Shankar Sastry. "Challenges for Securing Cyber Physical Systems," Workshop on Future Directions in Cyber-physical Systems Security, DHS, 23, July, 2009.

  We discuss three key challenges for securing cyberphysical systems: (1) understanding the threats, and possible consequences of attacks, (2) identifying the unique properties of cyber-physical systems and their differences from traditional IT security, and (3) discussing security mechanisms applicable to cyber-physical systems. In particular, we analyze security mechanisms for: prevention, detection and recovery, resilience and deterrence of attacks.

- [4]Anil Aswani, Nicholas Boyd, Claire Tomlin. "Graph-theoretic topological control of biological genetic networks," ACC 2009, 1700-1705, 10, June, 2009.

  The control of biological genetic networks is an important problem. If the system is abstracted into a graph, then the affect of drugs, pharmaceuticals, and gene therapy can be abstracted as changing the topology of the graph. We consider the control objective of removing the stable oscillations of the genetic network. This control is done using several theorems relating the topology of the network to the dynamics of the system. These theorems suggest that the controller should remove all the negative feedback in the networks.We prove that the problem of minimizing the edges and vertices to remove, in order to remove negative feedback, is NP-hard. In light of this result, a heuristic algorithm to solve this graph problem is presented. The algorithm is applied to several genetic networks, and it is shown that the heuristic gives reasonable results. Additionally, we consider the p53 network and show that the algorithm gives biologically relevant results.

- [15]Zong-Syun Lin, Alvaro Cardenas, Saurabh Amin, Yu-Lun Huang, Chi-Yen Huang, S. Shankar Sastry. "Model-Based Detection of Attacks for Process Control Systems," 16th ACM Computer and Communications Security Conference, ACM, submitted, 2009.

We present security analysis of process control systems (PCS) when an attacker can compromise sensor measurements that are critical for maintaining the operational goals. We present the general sensor attack model that can represent a wide variety of DoS and deception attacks. By taking example of a well studied process control system, we discuss the consequences of sensor attacks on the performance of the system and important implications for designing defense actions. We develop model-based detection methods that can be tuned to limit the false-alarm rates while detecting a large class of sensor attacks. From the attacker's viewpoint, we show that when the detection mechanisms and control system operations are understood by the attacker, it can carry stealth attacks that maximize the chance of missed detection. From the defender's viewpoint, we show that when an attack is detected, the use of model-based outputs maintains safety under compromised sensor measurements.

3 Outreach

3.1 Project Training and Development

3.2 Outreach Activities

Continuing in our mission to build a modern systems science (MSS) with profound implications on the nature and scope of computer science and engineering research, the structure of computer science and electrical engineering curricula, and future industrial practice. This new systems science must pervade engineering education throughout the undergraduate and graduate levels. Embedded software and systems represent a major departure from the current, separated structure of computer science (CS), computer engineering (CE), and electrical engineering (EE). In fact, the new, emerging systems science reintegrates information and physical sciences. The impact of this change on teaching is profound, and cannot be confined to graduate level.

This year we have continued our work to lay the foundation for a new philosophy of undergraduate teaching at the participating institutions.

3.2.1 Curriculum Development for Modern Systems Science (MSS)

Our agenda is to restructure computer science and electrical engineering curricula to adapt to a tighter integration of computational and physical systems. Embedded software and systems represent a major departure from the current, separated structure of computer science (CS), computer engineering (CE), and electrical engineering (EE). In fact, the new, emerging systems science reintegrates information and physical sciences. The impact of this change on teaching is profound, and cannot be confined to graduate level. Based on the ongoing, groundbreaking effort at UCB, we are engaged in retooling undergraduate teaching at the participating institutions, and making the results widely available to encourage critical discussion and facilitate adoption.

We are engaged in an effort at UCB to restructure the undergraduate systems curriculum (which includes courses in signals and systems, communications, signal processing, control systems, image processing, and random processes). The traditional curriculum in these areas is mature and established, so making changes is challenging. We are at the stage of attempting to build faculty consensus for an approach that shortens the pre-requisite chain and allows for introduction of new courses in hybrid systems and embedded software systems.

3.2.2 Undergrad Course Insertion and Transfer

At many institutions, introductory courses are quite large. This makes conducting such a course a substantial undertaking. In particular, the newness of the subject means that there are relatively few available homework and lab exercises and exam questions. To facilitate use of this approach by other instructors, we have engaged technical staff to build web infrastructure supporting such courses. We have built an instructor forum that enables submission and selection of problems from the text and from a library of submitted problems and exercises. A server-side infrastructure generates PDF files for problem sets and solution sets.

The tight integration of computational and physical topics offers opportunities for leveraging technology to illustrate fundamental concepts. We have developed a suite of web pages with applets that use sound, images, and graphs interactively. Our staff has extended and upgraded these applets and created a suite of PowerPoint slides for use by instructors.

We have begun to define an upper division course in embedded software (aimed at juniors and seniors). This new course will replace the control course at the upper division level at San Jose State. We also continued to teach at UC Berkeley the integrated course designed by Prof. Lee, which employs techniques discovered in the hybrid and embedded systems research to interpret traditional signals.

Course: Introduction to Embedded Systems (UCB EECS 149)
http://chess.eecs.berkeley.edu/eecs149

Instructors:
Prof. Edward A. Lee
Prof. Sanjit A. Seshia
Prof. Claire J. Tomlin

During the previous reporting period, Professor Tomlin assisted in development of the undergraduate Introduction to Embedded Systems course, EECS 124. For the Spring 2009 semester, the course was renumbered to EECS 149. The new material was taught in the Spring Semester of 2010 by Professor Sanjit A. Seshia.

The abstract for the class is below:

EECS 149 introduces students to the design and analysis of computational systems that interact with physical processes. Applications of such systems include medical devices and systems, consumer electronics, toys and games, assisted living, traffic control and safety, automotive systems, process control, energy management and conservation, environmental control, aircraft control systems, communications systems, instrumentation, critical infrastructure control (electric power, water resources, and communications systems for example), robotics and distributed robotics (telepresence, telemedicine), defense systems, manufacturing, and smart structures.

A major theme of this course will be on the interplay of practical design with models of systems, including both software components and physical dynamics. A major emphasis will be on building high confidence systems with real-time and concurrent behaviors.

This course is still under active development. This offering is therefore advised for advanced and adventurous undergraduates.

Course: Introduction to Control Design Techniques (UCB EECS 128)
http://inst.eecs.berkeley.edu/ ee128/fa08/

Instructor:
Prof. Claire J. Tomlin

During the previous reporting period Professor Tomlin has redesigned the undergrad-

uate control theory and engineering course, EECS 128, adding new labs and course material. The new material was taught in the Fall Semesters of 2008 and 2009.
The abstract for the class is below:
Root-locus and frequency response techniques for control system synthesis. State-space techniques for modeling, full-state feedback regulator design, pole placement, and observer design. Combined observer and regulator design. Lab experiments on computers connected to mechanical systems.

- Transfer function and state space models for control system analysis and synthesis. Pole locations and relationship to time response. Root locus methods. Stability.
- Feedback. Review of single-input single output (SISO) analysis and control methods in the frequency domain (Bode, Nyquist).
- SISO analysis and control using state space models. The matrix exponential and its relationship to time response. Controllability and observability. Combining state feedback with observers.
- Multi-input multi-output analysis and control using state space models.
- The linear quadratic regulator.

3.2.3 Graduate Courses

As part of the no-cost extension, a course in embedded systems was taught in the area of embedded and hybrid systems, as well as systems modeling. This course is a reflection of the teaching and curriculum goals of the ITR and its affiliated faculty.
Course: Linear System Theory(UCB EE221A)
http://inst.eecs.berkeley.edu/ ee221A/fa09/
Instructor: Claire J. Tomlin

Professor Tomlin is modernizing the graduate course in linear system theory, EECS 221A, adding units in linear programming and more general optimization. The new material was taught in the Fall Semesters of 2008, 2009 and now in 2010.
The abstract for the class is below:
This course provides a comprehensive introduction to the modeling, analysis, and control of linear dynamical systems. Topics include: A review of linear algebra and matrix theory. The solutions of linear equations. Least-squares approximation and linear programming. Linear ordinary differential equations: existence and uniqueness of solutions, the state-transition matrix and matrix exponential. Input-output and internal stability; the method of Lyapunov. Controllability and observability; basic realization theory. Control and observer design: pole placement, state estimation. Linear quadratic optimal control: Riccati equation and properties of the LQ regulator.

Advanced topics such as robust control and hybrid system theory will be presented based on allowable time and interest from the class.
This course provides a solid foundation for students doing research that requires the design and use of dynamic models. Students in control, circuits, signal processing, communications and networking are encouraged to take this course.

- Linear Algebra: Fields, vector spaces, subspaces, bases, dimension, range and Null spaces, linear operators, norms, inner products, adjoints.

- Matrix Theory: Eigenspaces, Jordan form, Hermitian forms, positive definiteness, singular value decomposition, functions of matrices, spectral mapping theorem, computational aspects.

- Optimization: Linear equations, least-squares approximation, linear programming.

- Differential Equations: existence and uniqueness of solutions, Lipschitz continuity, linear ordinary differential equations, the notion of state, the state-transition matrix.

- Stability: Internal stability, input-output stability, the method of Lyapunov.

- Linear Systems - open-loop aspects: controllability and observability, duality, canonical forms, the Kalman decomposition, realization theory, minimal realizations.

- Linear systems - feedback aspects: pole placement, stabilizability and detectability, observers, state estimation, the separation principle.

- Linear quadratic optimal control: least-squares control and estimation, Riccati equations, properties of the LQ regulator.

- Advanced topics: robust control, hybrid systems.

4 Publications and Products

In this section, we list published papers only. Submitted papers and in press papers are described in Section 2.2.

4.1 Technical reports

- None.

4.2 Software

- [16]Anil Aswani. "NODE model for eve mRNA modeling in Drosophila melanogaster," University of California, Berkeley, 2010.
- [17]Anil Aswani. "NEDE Estimator," University of California, Berkeley, 26, April, 2010.

4.3 PhD theses

- [12]Todd Templeton. PhD thesis, "Accurate Real-Time Reconstruction of Distant Scenes Using Computer Vision: The Recursive Multi-Frame Planar Parallax Algorithm," University of California, Berkeley, December, 2009.

4.4 Conference papers

- [14]Alvaro Cardenas, Saurabh Amin, Bruno Sinopoli, Annarita Giani, Adrian Perrig, S. Shankar Sastry. "Challenges for Securing Cyber Physical Systems," Workshop on Future Directions in Cyber-physical Systems Security, DHS, 23, July, 2009.
- [4]Anil Aswani, Nicholas Boyd, Claire Tomlin. "Graph-theoretic topological control of biological genetic networks," ACC 2009, 17001705, 10, June, 2009.
- [15]Zong-Syun Lin, Alvaro Cardenas, Saurabh Amin, Yu-Lun Huang, Chi-Yen Huang, S. Shankar Sastry. "Model-Based Detection of Attacks for Process Control Systems," 16th ACM Computer and Communications Security Conference, ACM, submitted, 2009.
- [11]Anil Aswani, Harendra Guturu, Claire Tomlin. "System identication of hunchback protein patterning in early Drosophila embryogenesis," Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference, 7723-7728, 18, December, 2009.
- [4]Anil Aswani, Nicholas Boyd, Claire Tomlin. "Graph-theoretic topological control of biological genetic networks," ACC 2009, 17001705, 10, June, 2009.

4.5 Books

- None.

4.6 Journal articles

- [13]YuLun Huang, Alvaro Cardenas, Saurabh Amin, Song-Zyun Lin , Hsin-Yi Tsai, S. Shankar Sastry. International Journal of Critical Infrastructure Protection, "Understanding the Physical and Economic Consequences of Attacks Against Control Systems.," 2(3):72-83, October 2009.

- [10]J. Biermeyer, T. Templeton, C. Berger, H. Gonzalez, N. Naikal, B. Rumpe, S. S. Sastry. Proceedings des 11. Braunschweiger Symposiums "Automatisierungssysteme, Assistenzsysteme und eingebettete Systeme fr Transportmittel", ITS Niedersachsen, Braunschweig, "Rapid Integration and Calibration of New Sensors Using the Berkeley Aachen Robotics Toolkit (BART)," 2010.

Although this is a long term project focused on foundations, we are actively working to set up effective technology transfer mechanisms for dissemination of the research results. A major part of this is expected to occur through the open dissemination of software tools.

4.6.1 The 2009-2010 Chess seminar series

The Chess seminar series provides a weekly forum for the problems and solutions found and solved by Chess members, as well as ongoing research updates. This forum works best when the audience is diverse in background, because the goal is to aid researchers in seeing how the other sub-disciplines are approaching similar problems, or to encourage them to work on problems they had not yet considered.
A full listing of this project-year's speakers is below. Most talks can be downloaded from the seminar website, at http://chess.eecs.berkeley.edu/seminar.htm

- "Bipedal Robotic Walking: Motivating the Study of Hybrid Phenomena" Professor Aaron D. Ames, Texas A&M University, May, 4, 2010.

- "Toward a Theory of Secure Networked Control Systems"
  Saurabh Amin, UC Berkeley, April 27, 2010.

- "Saving Energy by Reducing Traffic Flow Instabilities"
  Berthold K. Horn, MIT, April 26, 2010.

- "Planning and Learning in Information Space"
  Nicholas Roy, MIT, April 20, 2010.

- "SILVER: Synthesis Using Integrated Learning and Verification"
  Susmit Jha, University of California Berkeley, April 08, 2010.

- “Advances in Embedded Software Synthesis from Dataflow Models”
  Soheil Ghiasi, University of California Davis, April 06, 2010.
- “A Model-Based End-to-End Tool Chain for the Probabilistic Analysis of Complex Systems”
  Alessandro Pinto, United Technologies Research Center , March 30, 2010.
- “An Algebra of Pareto Points”
  Marc Geilen, Eindhoven University of Technology, March 16, 2010.
- “Robust Uncertainty Management Methods Applied to UAV Search and Tracking Problems”
  Andrzej Banaszuk, United Technologies Research Center, March 09, 2010.
- “Embedded Convex Optimization”
  Jacob Mattingley, Stanford University, March 02, 2010.
- “Software Failure Avoidance Using Discrete Control Theory”
  Terence Kelly & Yin Wang, HP Labs, February 23, 2010.
- “Simulation and Hybrid Control of Robotic Marionettes”
  Professor Todd Murphey, Northwestern University, February 16, 2010.
- “A Natural-Language-Based Approach to System ModelingFrom the Use of Numbers to the Use of Word”
  Professor Lotfi A. Zadeh, UC Berkeley, February 09, 2010.
- “On Relational Interfaces”
  Stavros Tripakis, UC Berkeley, February 02, 2010.
- “Modeling, Planning, and Control for Robot-Assisted Medical Interventions”
  Allison Okamura, Johns Hopkins, January 28, 2010.
- “Residential Load Management using Autonomous Auctions”
  William Burke, UC Berkeley, January 19, 2010.
- “From High-Level Component-Based Models to Distributed Implementations”
  Borzoo Bonakdarpour, Verimag Laboratory, January 12, 2010.
- “Extending Ptolemy to Support Software Model Variant and Configuration Management”
  Charles Shelton, Robert Bosch Research and Technology Center, December 15, 2009.
- “Active Traffic Management using Aurora Road Network Modeler”
  Alex A. Kurzhanskiy, University of California at Berkeley, December 08, 2009.

- “Distributed Coverage Control with On-Line Learning”
  Mac Schwager, Computer Science and Artificial Intelligence Lab, MIT, December 1, 2009.
- “Mobile Floating Sensor Network Placement using the Saint-Venant 1D Equation”
  Andrew Tinka, University of California Berkeley, November 24, 2009.
- “High-Level Tasks to Correct Low-Level Robot Control”
  Hadas Kress-Gazit, Cornell University, November 17, 2009.
- “Mobile Device Insights - Virtualization and Device Interoperability”
  Jorg Brakensiek, Nokia Research Center, November 10, 2009.
- “Understanding the Genome of Data Centers”
  Jie Liu, Microsoft Research, November 03, 2009.
- “Robust Control via Sign Definite Decomposition”
  Shankar P. Bhattacharyya, Department of Electrical Engineering, Texas A&M University, October 29, 2009.
- “Mathematical Equations as Executable Models of Mechanical Systems”
  Walid Taha, Rice University, October 22, 2009.
- “Data-Driven Modeling of Dynamical Systems: Optimal Excitation Signals and Structured Systems”
  Bo Wahlberg, Automatic Control Lab and ACCESS, KTH, Stockholm, Sweden , October 13, 2009.
- “Simulating Print Service Provider Using Ptolemy II”
  Jun Zeng, Hewlett-Packard Laboratories, October 6, 2009.
- “Robust Distributed Task Planning for Networked Agents”
  Han-Lim Choi, Massachusetts Institute of Technology, September 29, 2009.
- “User-generated 3-D Content: Personalizing Immersive Connected Experiences”
  Yimin Zhang, Intel Labs China, September 25, 2009.
- “The Relation of Spike Timing to Large-Scale LFP Patterns”
  Ryan Canolty, University of California, Berkeley, September 22, 2009.
- “Concurrency and Scalability versus Fragmentation and Compaction with Compact-fit”
  Hannes Payer, University of Salzburg, September 15, 2009.
- “Avoiding Unbounded Priority Inversion in Barrier Protocols Using Gang Priority Management”
  Harald Roeck, University of Salzburg, September 15, 2009.

- “The Design, Modeling, and Control of Mechatronic Systems with Emphasis on Betterment of Quality of Human Life”
  Kyoungchul Kong, University of California, Berkeley, September 8, 2009.

- “Large Monitoring Systems: Data Analysis, Design and Deployment”
  Ram Rajagopal, University of California, Berkeley, September 3, 2009.

- “Reasoning about Online Algorithms with Weighted Automata”
  Orna Kupferman, Hebrew University, August 25, 2009.

4.6.2 Workshops and Invited Talks

In addition to the below invited and workshop organizational activities, Chess faculty have delivered numerous plenary talks, invited talks, as well as informal dissemination of Chess goals and research.

- C. Tomlin, Short courses on hybrid systems, Royal Institute of Technology, Stockholm, Spring 2010

- S. Sastry, Short course on Generalized Principal Component Analysis, Royal Institute of Technology, Stockholm, Spring 2010.

4.6.3 General Dissemination

The Chess website, http://chess.eecs.berkeley.edu, includes publications and software distributions. In addition, as part of the outreach effort, the UC Berkeley introductory signals systems course, which introduces hybrid systems, is available.

4.7 Other Specific Products

The following software packages have been made available during this review period on the Chess website, http://chess.eecs.berkeley.edu:

- http://bdtnp.lbl.gov/Fly-Net/bioimaging.jsp?w=node. Implements NODE model for eve mRNA modeling in Drosophila melanogaster. Author: Anil Aswani

- http://hybrid.eecs.berkeley.edu/ NEDE/EDE_Code.zip Implements NEDE estimator for regression on manifolds or with significant data collinearity. Author: Anil Aswani

5 Contributions

5.1 Human Resource Development

Several panels in important conferences and workshops pertinent to embedded systems (e.g., DAC, ICCAD, HSCC, EMSOFT, CASES, and RTSS) have pointed out the necessity of upgrading the talents of the engineering community to cope with the challenges posed by the next generation embedded system technology. Our research program has touched many graduate students in our institutions and several visiting researchers from industry and other Universities so that they now have a deep understanding of embedded system software issues and techniques to address them.

Specifically, our directors played a major role in the development of workshops and briefings to executives and researchers in the avionics industry to motivate increased research spending due to an anticipated drop in research funds available to train graduates in embedded software and embedded systems. One particular intersection with our efforts is the Software Producibility Initiative out of the Office of the Secretary of Defense.

The industrial affiliates to our research program are increasing and we hope to be able to export in their environments a modern view of system design. Preliminary feedback from our partners has underlined the importance of this process to develop the professional talent pool.

5.2 Integration of Research and Education

In this report, we have touched multiple times on research and education especially in the outreach section. In addition, there has been a strong activity in the continued update of the undergraduate course taught at Berkeley on the foundations of embedded system design. The graduate program at Berkeley and at Vanderbilt has greatly benefited from the research work in the ITR. EE249 at Berkeley has incorporated the most important results thus far obtained in the research program. EE 290 A and C, advanced courses for PhD students, have featured hybrid system and the interface theories developed under this project. EE219C, a course on formal verification, has used results from the hybrid theory verification work in the program. Finally, many final projects in these graduate courses have resulted in papers and reports listed in this document. The course EE291E on Hybrid Systems: Computation and Control is jointly taught at Berkeley and Vanderbilt and is benefiting a great deal from comments of students as far as the development of new text book material.

In addition to the influence on graduate students, we have endeavored to show hybrid and embedded systems as emerging research opportunities to undergraduates. We have also demonstrated that for advanced undergraduates these topics are not out of place as senior design courses, or advanced topics courses, which may in the future lead to the integration of these as disciplines in engineering across a broader reach of universities.

5.3 Beyond Science and Engineering

Embedded systems are part of our everyday life and will be much more so in the future. In particular, wireless sensor networks will provide a framework for much better environmental monitoring, energy conservation programs, defense and health care. Already in the application chapter, we can see the impact of our work on these themes. In the domain of transportation systems, our research is improving safety in cars, and foundationally improving control of energy conserving aspects such as hydrocarbon emissions. Future applications of hybrid system technology will involve biological systems to a much larger extent showing that our approach can be exported to other field of knowledge ranging from economics to biology and medicine. At Berkeley, the Center for Information Technology Research in the Interest of Society is demonstrating the potential of our research in fields that touch all aspects of our life. Some key societal grand challenge problems where our ITR research is making a difference includes health care delivery, high confidence medical devices and systems, avionics, cybersecurity, and transportation.

References

[1] Anil Aswani, Peter Bickel, and Claire Tomlin. Regression on manifolds: Estimation of the exterior derivative. Annals of Statistics, 2010. To appear.

[2] A. Aswani and C. J. Tomlin. Monotone piecewise affine systems. IEEE Transactions on Automatic Control, 54(8), 2009.

[3] Anil Aswani and Claire Tomlin. Computer-aided drug discovery for pathway and genetic diseases. Unpublished, 2010.

[4] Anil Aswani, Nicholas Boyd, and Claire Tomlin. Graph-theoretic topological control of biological genetic networks. In ACC 2009, page 17001705, June 2009.

[5] Milosh Drezgich and S. Shankar Sastry. Matrix multiplicative weights and non-zero sum games. 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, 2010. Submitted.

[6] A. Aswani, S. Keranan, J. Brown, C. Fowlkes, D. Knowles, M. Biggin, P. Bickel, and C. J. Tomlin. Nonparametric identification of regulatory interactions from spatial and temporal gene expression data. BMC Bioinformatics, 11(413), 2010.

[7] Y. H. chang, J. Gray, and C. J. Tomlin. Optimization-based inference for temporally evolving networks with applications to biology. In Proceedings of the American Control Conference, July 2011. submitted.

[8] S. Itani, J. Gray, and C. J. Tomlin. An ode model for the her2/3-akt signaling pathways in cancers that overexpress her2. In Proceedings of the American Control Conference, July 2010.

[9] Milosh Drezgich and S. Shankar Sastry. On the np in bqp. unpublished, 2010.

[10] J. Biermeyer, T. Templeton, C. Berger, H. Gonzalez, N. Naikal, B. Rumpe, and S. S. Sastry. Rapid integration and calibration of new sensors using the berkeley aachen robotics toolkit (bart). Proceedings des 11. Braunschweiger Symposiums "Automatisierungssysteme, Assistenzsysteme und eingebettete Systeme fur Transportmittel", ITS Niedersachsen, Braunschweig, 2010.

[11] Anil Aswani, Harendra Guturu, and Claire Tomlin. System identication of hunchback protein patterning in early drosophila embryogenesis. In Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference, pages 7723–7728, December 2009.

[12] Todd Templeton. Accurate Real-Time Reconstruction of Distant Scenes Using Computer Vision: The Recursive Multi-Frame Planar Parallax Algorithm. PhD thesis, University of California, Berkeley, December 2009.

[13] YuLun Huang, Alvaro Cardenas, Saurabh Amin, Song-Zyun Lin, Hsin-Yi Tsai, and S. Shankar Sastry. Understanding the physical and economic consequences of attacks against control systems. International Journal of Critical Infrastructure Protection, 2(3):72–83, October 2009.

[14] Alvaro Cardenas, Saurabh Amin, Bruno Sinopoli, Annarita Giani, Adrian Perrig, and S. Shankar Sastry. Challenges for securing cyber physical systems. In Workshop on Future Directions in Cyber-physical Systems Security. DHS, July 2009.

[15] Zong-Syun Lin, Alvaro Cardenas, Saurabh Amin, Yu-Lun Huang, Chi-Yen Huang, and S. Shankar Sastry. Model-based detection of attacks for process control systems. In 16th ACM Computer and Communications Security Conference, page submitted. ACM, 2009.

[16] Anil Aswani. Node model for eve mrna modeling in drosophila melanogaster, 2010. Software.

[17] Anil Aswani. Nede estimator, April 2010. Software.