BOSTON
UNIVERSITY

Cyber Security Research Projects at BU

Azer Bestavros

Professor and former Chair of Computer Science
Director of Hariri Institute for Computing
Founding Member of RISCS

NSF US/Mideast Workshop on Trustworthiness in
Emerging Distributed Systems and Networks
Istanbul, Turkey -- June 4, 2012

Active Research Projects (~$3M/yr)

- Trustworthy Cloud Computing [Bestavros et al]
- Formal Verification of Software and Networks [Kfoury & Bestavros]
- Memory-Safe Programming Languages and Software [Xi & West]
- Distributed Spatial Anomaly Detection [Crovella]
- Attack Resistant Cryptographic Hardware [Karpovsky & Tobin]
- Data Authentication for Outsourced Databases [Kollios & Reyzin]
- Anonymous Peer-to-Peer Overlays [Bestavros, Goldberg, & Matta]
- Market-Based SPAM Management [van Alstyne]
- Privacy-preserving Mining of Social Networks [Terzi]
- Mining Network Data for Influence & Authority [Terzi et al]
- Towards Composable Security Analysis [Canetti]
- Secure BGP Routing [Goldberg & Reyzin]
- Clean-Slate Internet Architectures using RINA [Matta]
- eXpressive Internet Architecture [Byers et al]
- Securing the Open Softphone [Crovella et al]

Two Example Projects

- ***CloudCommons:***
  Team: Bestavros, Appavoo + Ishakian, Skowyra, Sweha, Bahargam
  - Mechanism design and system support for expressing and exploiting cloud supply & demand elasticity
  - Techniques for cloud-assisted security-performance tradeoffs with application to anonymity
- ***iBench:***
  Team: Kfoury, Bestavros + Lapets, Reynolds, Skowyra, Bassem
  - 0-day malware detection using a software certification portal for Java and ActionScript binaries
  - Unified environment for verification of safety and security properties of distributed protocols and systems

CC: Mechanisms for Cloud Markets

- ***Colocation Services:***
  Allow workloads (demand) to aggregate rationally
  - Colocation Games, applied to Xen VM colocation (XCS)
  - Cap and Trade, applied to bandwidth management
- ***Morphosys:***
  Exploit (and reward) supply/demand flexibility
  - Expressive languages to specify SLAs, including dials for performance and security
  - Provably-safe SLA transformations
  - Shapley pricing mechanisms to reward (incentivize) the expression of flexibility

CC: Security-Performance Tradeoffs

- ***Crypsis:***
  How to protect cloud and P2P content distribution systems from Zenith attacks aiming to identify popular content
- ***Cyclops:***
  How to use cloud resources to guarantee the availability of P2P-hosted content on a budget
- ***AngelCast:***
  How to use cloud resources as insurance against QoS degradation – e.g., for peer-assisted streaming
- ***TorAssist:***
  How to use cloud resources to improve performance of Tor without compromising anonymity

CC: TorAssist

- Identify cause of Tor's poor performance
  - High variability in available node capacities
  - Noisy feedback of its load-balancing controller
- Investigate whether strategically adding cloud Tor "angels" could ameliorate performance
  - Node capacity
  - Node placement
- Suggest incentives for users to buy in
  - New functionality that benefit users with the right client

iBench: Zero-Day Malware Detection

- Signature-based, anti-malware solutions work for known malware, but not for novel malware
- There is an arms race between malware authors and anti-malware software vendors
- There are numerous techniques to disguise existing exploits and evade detection
- Malware authors have a substantial financial incentive to also invent novel exploits
- The number of "0-day" exploits continues to rise

iBench: Alloy-based software verification

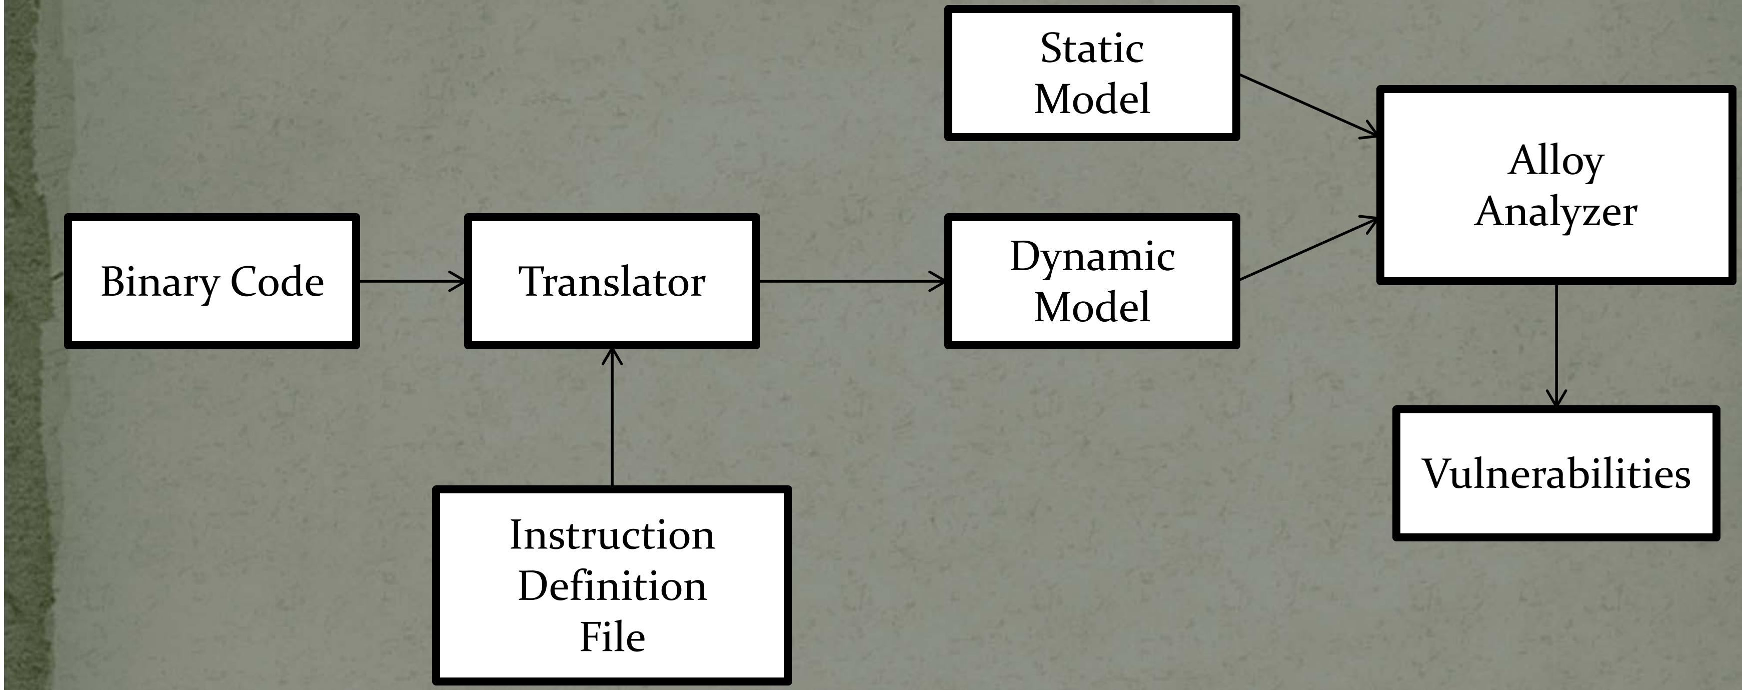

Can analyze any binary code that can be described by an IDF
Alloy system analysis time is measured in hours
Standard analysis time using conventional tools on live malware is months
This system works now on Java class files and ActionScript Flash files
Extensions planned for Android, Microsoft .NET, and ARM

iBench: Unified Verification Assistance

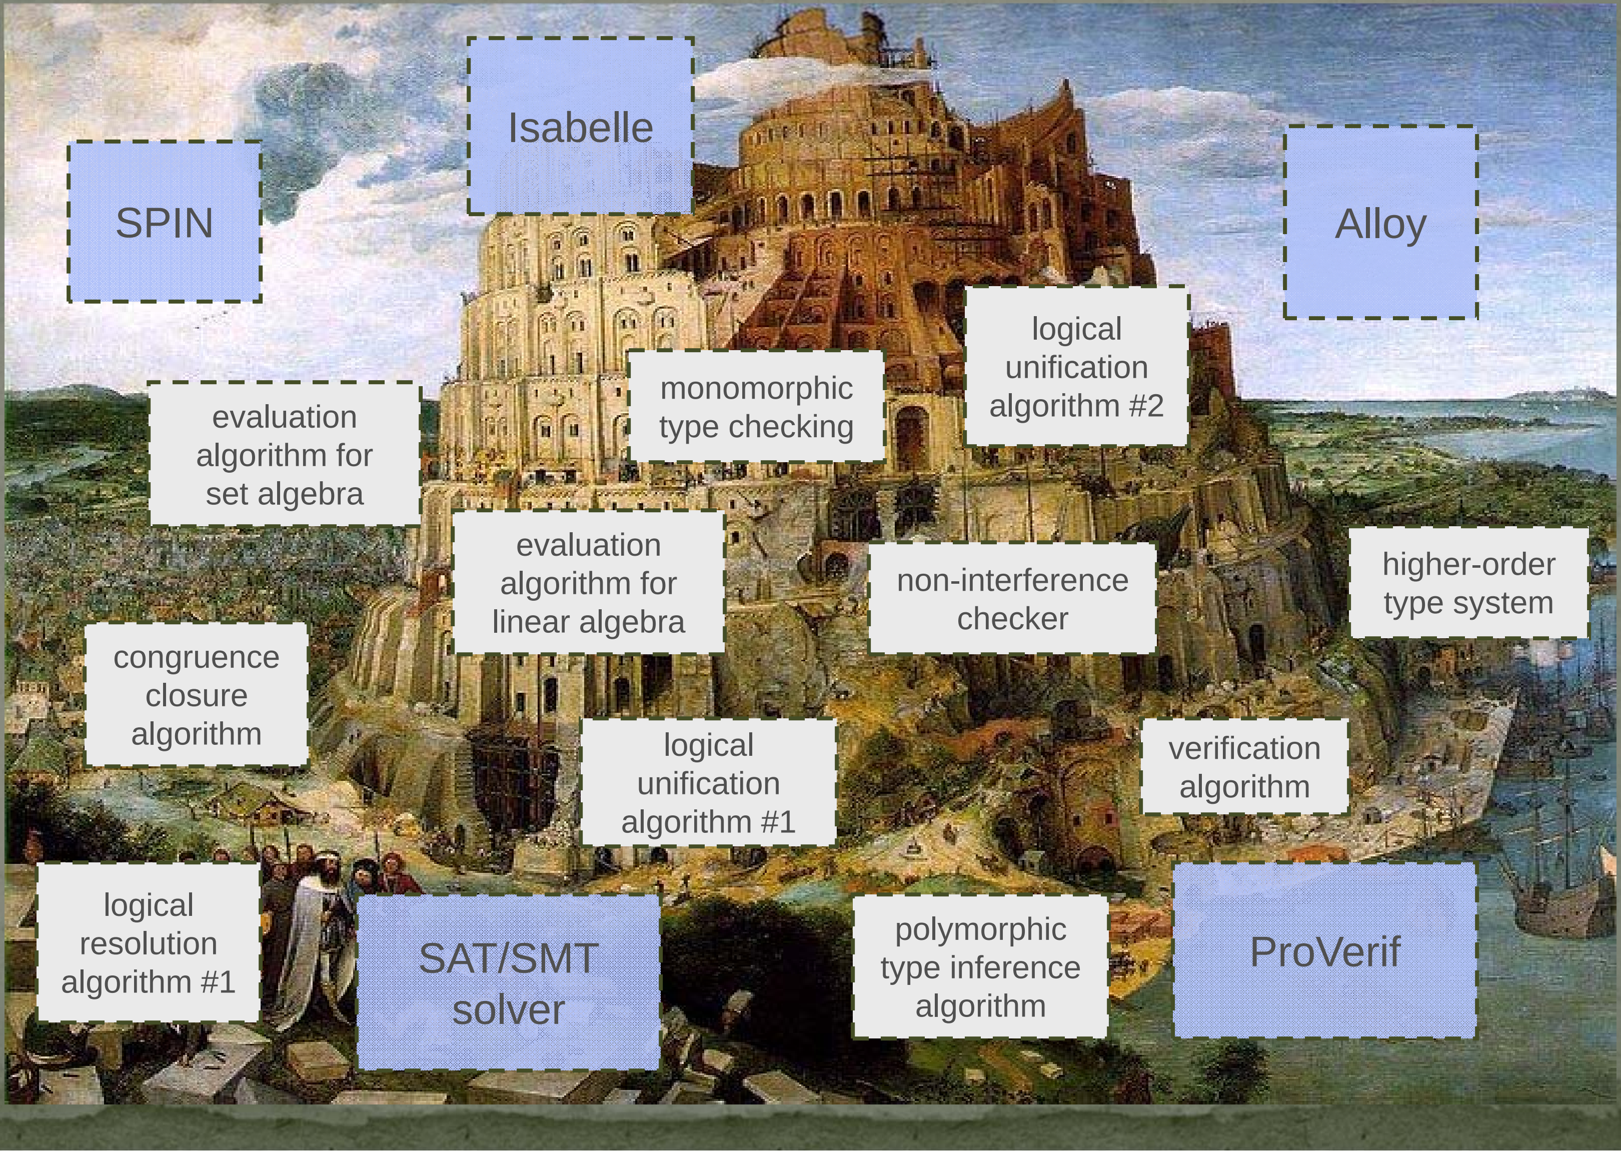

Isabelle

SPIN

Alloy

logical unification algorithm #2

monomorphic type checking

evaluation algorithm for set algebra

evaluation algorithm for linear algebra

non-interference checker

higher-order type system

congruence closure algorithm

logical unification algorithm #1

verification algorithm

logical resolution algorithm #1

SAT/SMT solver

polymorphic type inference algorithm

ProVerif

iBench: Integrate and Hide Complexity

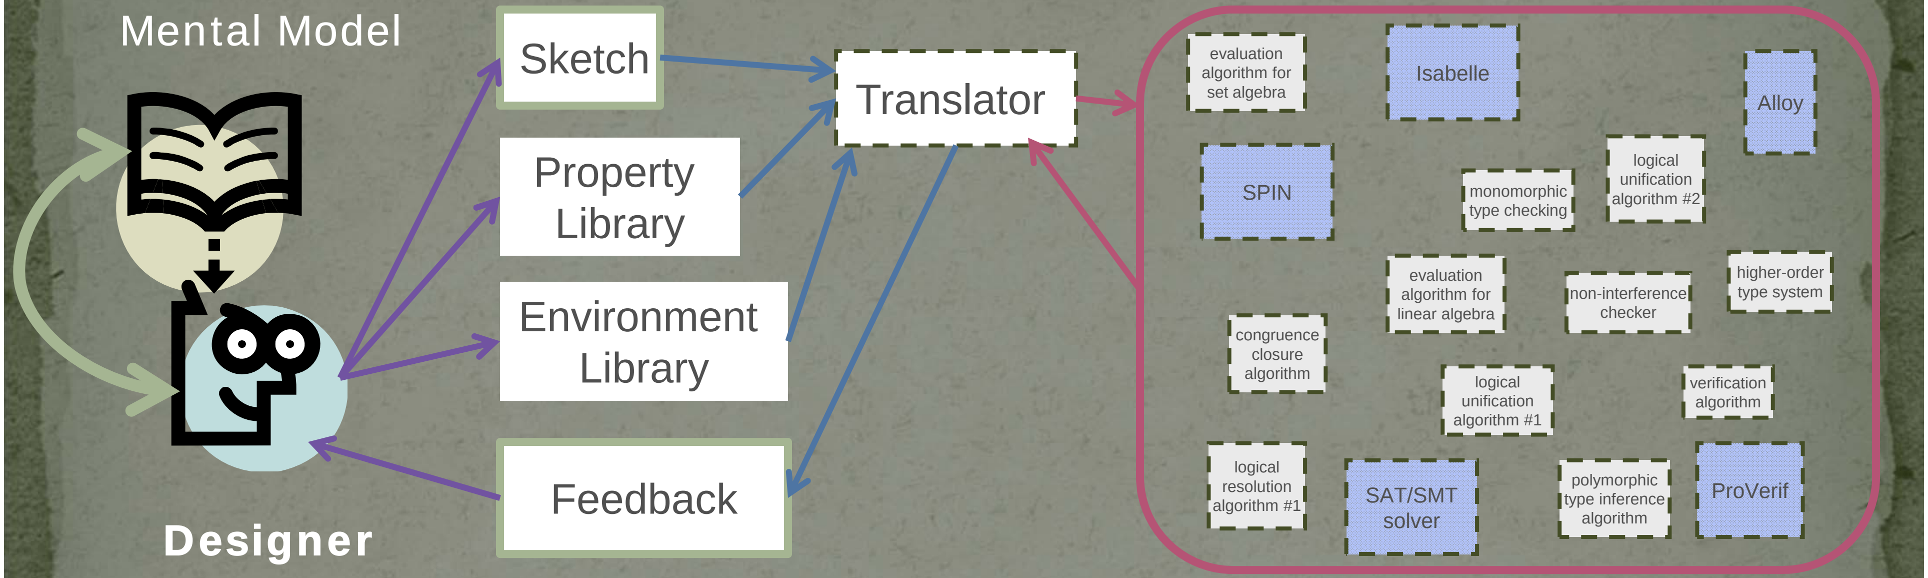

Mental Model

Designer

Sketch

Property Library

Environment Library

Feedback

Translator

evaluation algorithm for set algebra

Isabelle

Alloy

SPIN

monomorphic type checking

logical unification algorithm #2

evaluation algorithm for linear algebra

non-interference checker

higher-order type system

congruence closure algorithm

logical unification algorithm #1

verification algorithm

logical resolution algorithm #1

SAT/SMT solver

polymorphic type inference algorithm

ProVerif

1. Designer makes a sketch, selects interesting properties and environments from provided library (or makes his own)
2. Translator converts to formal specifications
3. Formal tools provide analysis
4. Translator converts formal output to statements about sketch

iBench: What Kind of Assistance?

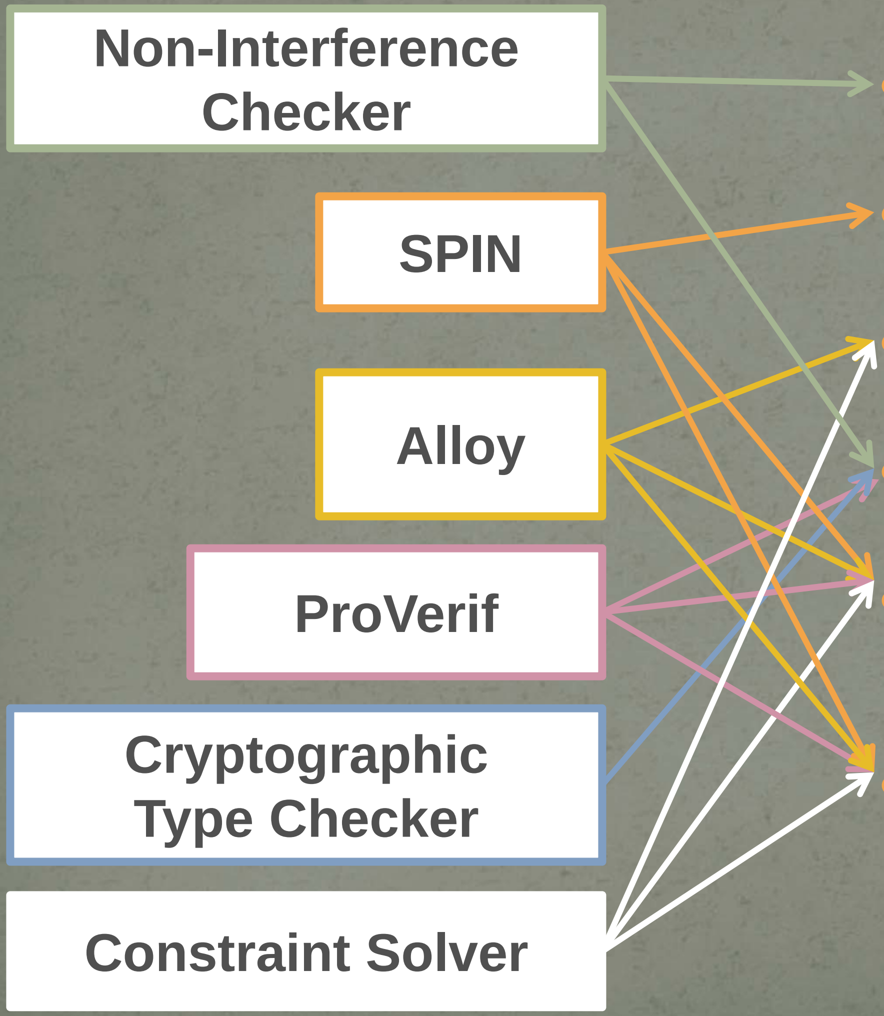

Non-Interference Checker

SPIN

Alloy

ProVerif

Cryptographic Type Checker

Constraint Solver

- ‘This (public) variable affects this (secret) variable’
- ‘Deadlocks are possible in the following scenarios…’
- ‘These properties are inconsistent…’
- ‘This code is insecure. One way to make it secure is…’
- ‘The protocol only works of the following assumptions are made…’
- ‘The protocol still works if the following assumptions are removed…’

Active Research Projects (~$3M/yr)

- Trustworthy Cloud Computing [Bestavros et al]
- Formal Verification of Software and Networks [Kfoury & Bestavros]
- Memory-Safe Programming Languages and Software [Xi & West]
- Distributed Spatial Anomaly Detection [Crovella]
- Attack Resistant Cryptographic Hardware [Karpovsky & Tobin]
- Data Authentication for Outsourced Databases [Kollios & Reyzin]
- Anonymous Peer-to-Peer Overlays [Bestavros, Goldberg, & Matta]
- Market-Based SPAM Management [van Alstyne]
- Privacy-preserving Mining of Social Networks [Terzi]
- Mining Network Data for Influence & Authority [Terzi et al]
- Towards Composable Security Analysis [Canetti]
- Secure BGP Routing [Goldberg & Reyzin]
- Clean-Slate Internet Architectures using RINA [Matta]
- eXpressive Internet Architecture [Byers et al]
- Securing the Open Softphone [Crovella et al]

Trustworthy Cloud Computing [Bestavros ++ @ Brown U & UCI]

- ***Motivation & Goals:*** Cloud computing introduces opportunities and challenges for security – need to make security an integral part of cloud SLAs
- ***Approach & Results:*** Develop expressive SLAs and associate delivery and validation mechanisms to enhance trust in cloud interactions, including
  - Ability to check data integrity and consistency
  - Develop SLA mechanisms for fair market valuation
  - Develop protocols for safe SLA transformations for automated service colocation, negotiation, and optimization

Formal Verification of Software and Networks [Kfoury & Bestavros]

- ***Motivation & Goals:*** Verify/infer overall safety and security properties from component specification

- ***Approach & Results:*** Design domain-specific formal languages to
  - Encapsulate safety properties
  - Support compositional/scalable verification
  - Applied to real-time & QoS properties of cyber-physical systems and flow networks

Safe Programming Languages and Software [Xi & West]

- ***Motivation & Goals:*** Enhance security by making software artifacts less vulnerable to program exploits

- ***Approach & Results:*** Develop "safe" programming languages and execution environments that are not vulnerable to attacks through software exploits
  - Provably ensure memory safety
  - Enable programmers to assert security properties
  - Enable verification of asserted security at compile time
  - Applied to development of device drivers (using ATS) as well as to virtualization environments (using sandboxing)

Distributed Spatial Anomaly Detection [Crovella]

- ***Motivation & Goals:*** Detect Internet Traffic Volume Anomalies

- ***Approach & Results:*** Leverage observations at multiple locations based on following principles:
  - Avoid global communication and centralized control
  - Augment current parametric anomaly detection methods with non-parametric methods
  - Annotate anomalies with probabilistic quantifier of its importance, (not just identify possible anomalies)
  - Used effectively for Internet – basis for "Guavus" startup

Low-Rate Exploits of Network Dynamics [Bestavros & Matta]

- ***Motivation & Goals:*** Harden systems and networks against stealthier DoS and RoQ attacks that exploit protocol dynamics

- ***Approach & Results:*** Develop signatures for low-rate attacks and study vulnerability of multiple protocols
  - Used control theory to define and evaluate exploits of network and system adaptation dynamics
  - Applied to attacks mounted against congestion control, admission control, load balancers, virtual machines, among others

Attack Resistant Cryptographic Hardware [Karpovsky & Tobin]

- ***Motivation & Goals:*** Transactions are moving into open and mobile environment, resulting in new threats and attacks

- ***Approach & Results***: Design secure, low-cost, low-power special-purpose hardware devices based on asynchronous fine grain pipelining and robust encoding of data, resulting in
  - Unique tools for secure hardware design
  - Best performance per Watt
  - Multiple fault injection attack tolerance

Data Authentication for Outsourced Databases [Kollios & Reyzin]

- ***Motivation & Goals:*** Enable clients at the edge of an untrusted cloud to access and query the data efficiently, while getting assurance of integrity

- ***Approach & Results***: Several new approaches are proposed, and analytically and experimentally studied
  - Solutions extend existing indexing structures (e.g., using Merkle Trees)
  - Applied to a range of DB query processing forms, including range queries
  - Shown to work very well even for very large datasets

Anonymous Peer-to-Peer Overlays [Bestavros, Goldberg & Matta]

- ***Motivation & Goals:*** P2P structured overlays could be potentially used to enhance secure communication and circumvent censorship technologies
- ***Approach & Results**:* Identified potential (Zenith) attacks against P2P overlays targeting popular content and developed appropriate, efficient defenses
  - Techniques tested on multiple DHT structured overlays
  - Novel DHT lookup protocols that are immune to Zenith attacks have been developed and tested
  - Trustworthy resource discovery in P2P overlays without reliance on a centralized trust authorities

Market-Based SPAM Management [van Alstyne]

- ***Motivation & Goals:*** Apply economic rather than technological or regulatory screening to manage SPAM
- ***Approach & Results:*** Instead of just blocking SPAM, recognize and promote valuable communication and provide feedback to spammers and users
  - Shift focus away from the information in the message to the information known to the sender
  - Use principles of information asymmetry to cause the spammer to incur higher costs than senders of legitimate information
  - Often outperforms "perfect" filter

Privacy-preserving Mining in Social Networks [Terzi]

- ***Motivation & Goals:*** Information leakage through social networks threatens privacy even in the presence of privacy controls

- ***Approach & Results:*** Develop models and analysis techniques to evaluate and counter the threats to privacy from “second hand” information leakage
  - Developed and tested techniques to recover information from randomized social network graphs
  - Developed and tested a framework for computing the privacy score of users in online social networks
  - Developed identity anonymization techniques for social nets

Mining Network Data for Influence & Authority [Terzi and Bestavros & Byers++ @ Harvard]

- ***Motivation & Goals:*** Identify authoritative and influential groups of nodes by mining network data (social, e-commerce, communication, media, CPS, etc.)

- ***Approach & Results:*** Develop new graph models, metrics and algorithms to measure group centrality
  - Characterized impact of targeted Groupon offers on popularity and reputation
  - Developed and tested techniques to identify expertise and influence in collaboration networks.
  - Developed and tested techniques to identify optimal placement of information aggregators and filters.
  - Developed and tested techniques for the management of centrality with applications to targeted information gathering and advertisement.

Composable Security Analysis [Canetti]

- ***Motivation & Goals:*** Combining individually-secure protocols may result in new vulnerabilities; need systematic approach to decide on composable securit
- ***Approach & Results:*** Study conditions and limitations of composability of cryptographic constructs. Research includes
  - Universal composability with global set-up
  - Composability of cryptographic protocols
  - Trading off soundness, simplicity and efficiency
  - Application to software obfusctation

Secure BGP Routing on the Internet [Goldberg & Reyzin]

- ***Motivation & Goals:*** Routing remains the "weakest link" on the Internet due to the lack of authentication of route advertisement in BGP

- ***Approach & Results:*** Develop new secure BGP protocols that are provably correct and study approaches to their deployment
  - Showed security vulnerabilities in many proposed S*BGP protocols and developed alternatives
  - Studied market-driven approaches to the deployment of S*BGP on the Internet

Clean-Slate Internet Architectures using RINA [Matta]

- ***Motivation & Goals:*** Security is an after tought in current Internet architecture – plugging holes is hopelessly inadequate; need clean-slate design

- ***Approach & Results:*** Adopt RPC as the main and only building block for Internet protocols and services, which can be recursively constructed
  - No standard protocols or naming convention
  - Security is tailored for each application
  - Approach demonstrated for applications in mobile and wireless settings

Situational Awareness via Anomaly Detection [Barford, Cassandras, Crovella , Paschalidis]

- ***Motivation & Goals:*** Coordinated (network-wide) anomaly detection could be used for situational awareness.
- ***Approach & Results:*** Leverage the potential from better anomaly detection from multiple vantage points
    - Distributed anomaly detection algorithms and tools
    - Coordination of local and network-wide views
    - Use clustering and pattern recognition approaches to identify and classify specific cyber attack scenarios

Securing the Softphone [Crovella ++]

- ***Motivation & Goals:*** New smart phones are increasingly open and easily susceptible to exploits due to ubiquity of "apps" and of multi-channel communication

- ***Approach & Results:*** Develop a multi-pronged approach using clean-slate designs
  - Hardens the physical layer (hardware)
  - Develop incentive-compatible protocols
  - Develop centralized and distributed defenses