

Scale Changes Everything

Linda Northrop
Director, Product Line Systems Program
Software Engineering Institute

ICGSE 2006
October 17, 2006



Software Engineering Institute

Carnegie Mellon

© 2006 Carnegie Mellon University

Software Engineering Institute (SEI)

- Federally Funded Research and Development Center
- Created in 1984
- Sponsored by the U.S. Department of Defense
- Locations in Pittsburgh, PA; Washington, DC; Huntsville, AL; Los Angeles, CA; Frankfurt, Germany; Doha, Qatar
- Operated by Carnegie Mellon University
- Works directly with global commercial and government organizations




Ultra-Large-Scale Systems (ULS)

**Scale
Changes
Everything**



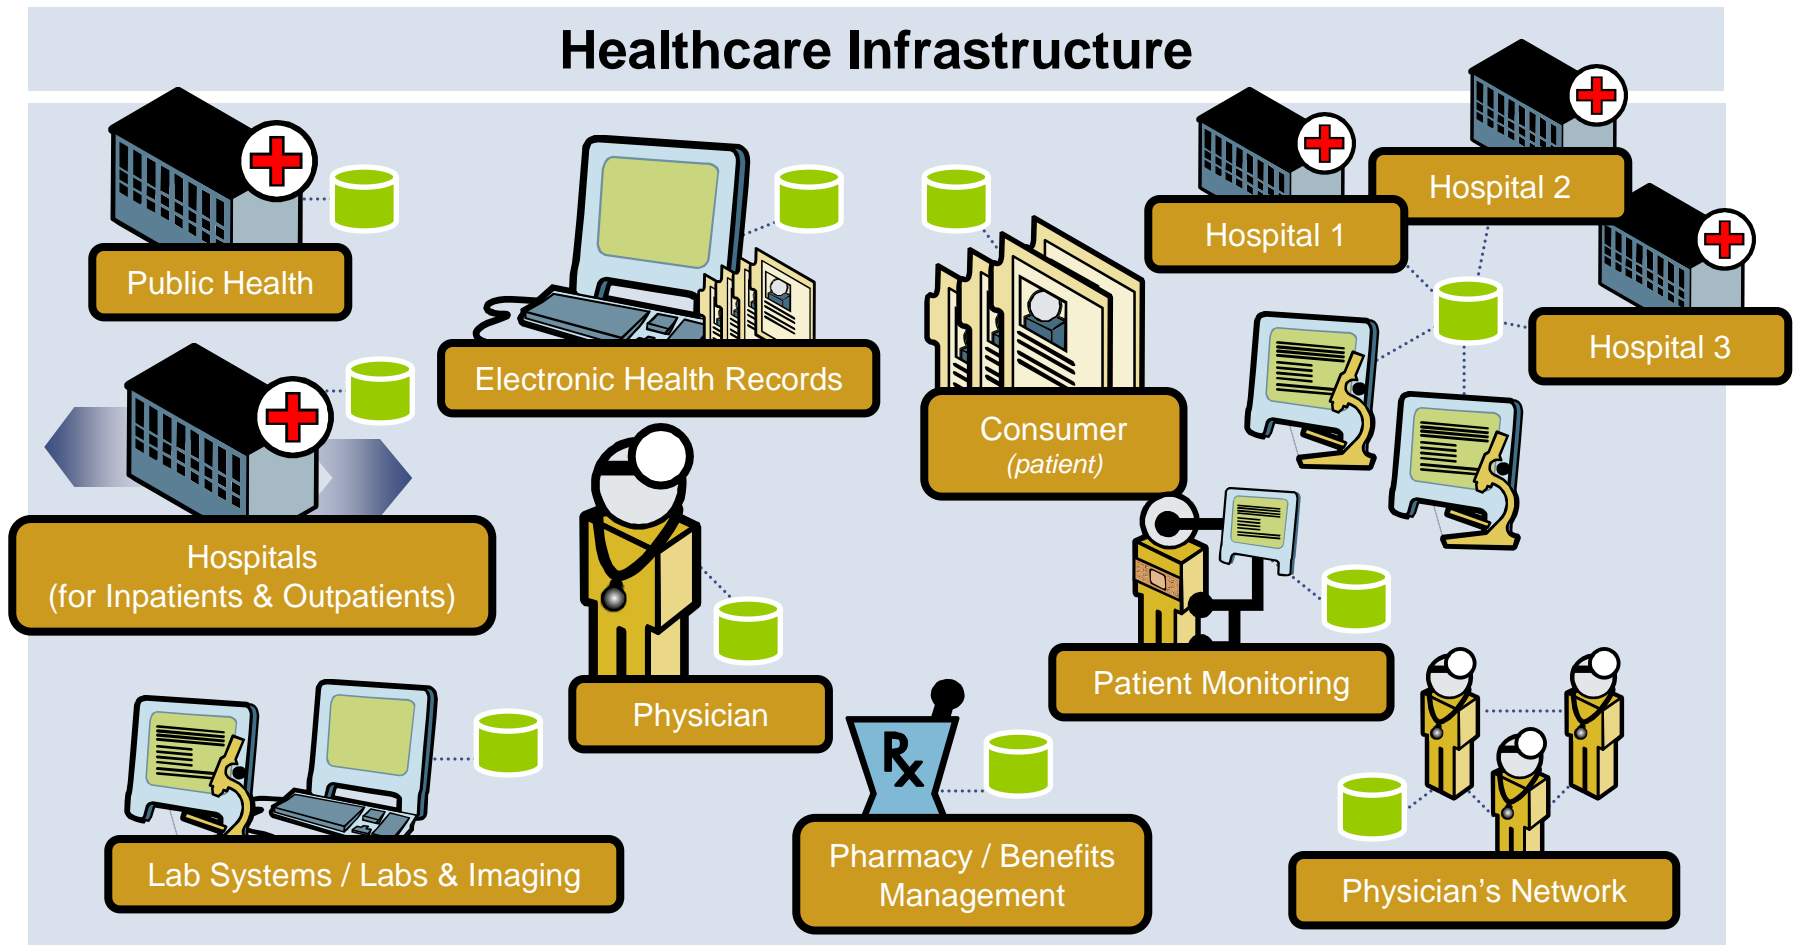
Trend Toward Increasing Scale-1



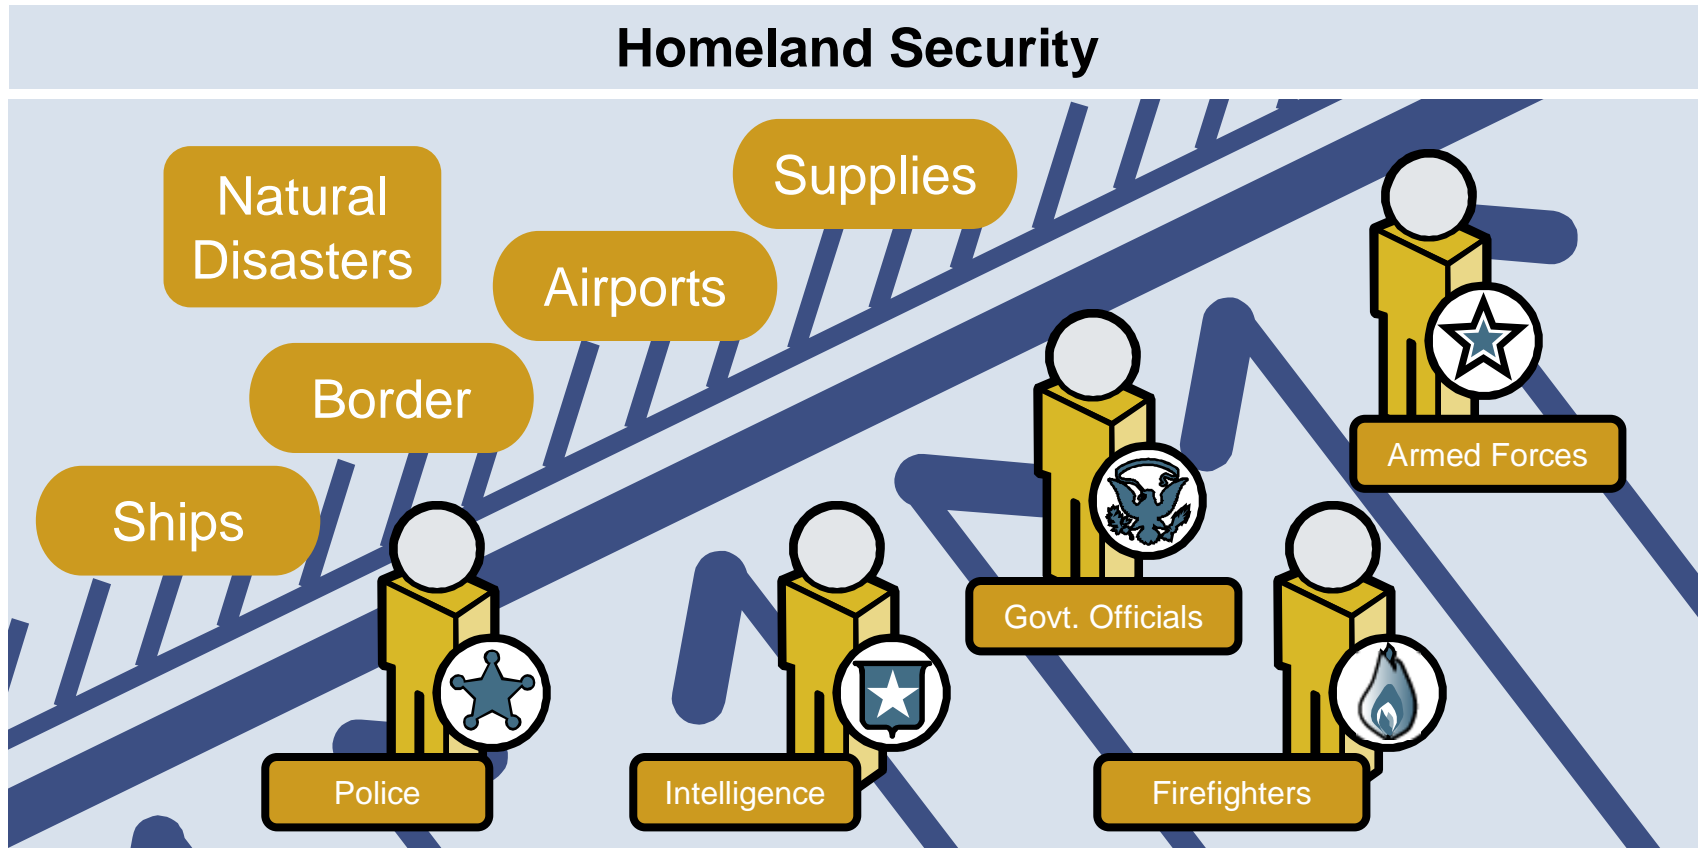
- Enormous web service and computing infrastructure
- Supply chain systems
- Software-based engineering systems



Trend Toward Increasing Scale-2



Trend Toward Increasing Scale-3



Increasing Scale In Military Systems

Increasingly Complex Systems

- ultra-large, network-centric, real-time, cyber-physical-social systems
 - thousands of platforms, sensors, decision nodes, weapons, and warfighters
 - connected through heterogeneous wired and wireless networks

Goal: Information Dominance

- *Transient and enduring resource constraints and failures*
- *Continuous adaptation*
 - changes in mission requirements
 - changes in operating environments
 - changes in force structure
 - perpetual systems' evolution
 - addition of new systems
- *Sustainable - legally, technically, politically*



A Reason for Concern

Such systems are going to be larger and more complex than any previously seen

- very serious technical challenges, obvious and undoubtedly to-be-discovered
- many vendors, many technologies, many systems
- evolving doctrine + evolving technology + (or \Rightarrow ?) ill-defined requirements

The US Army is concerned that the scale of future systems is beyond our reach.



The Challenge



*The Honorable
Claude M. Bolton, Jr.*

“Our soldiers depend on software and will depend more on software in the future.

The Army’s success depends on software and the software industry.

We need better tools to meet future challenges, and neither industry nor government is working on how to do things light-years faster and cheaper.

How can future systems, which are likely to be a **billion lines of code**, be built reliably if we can’t even get today’s systems right?”

— Asst Sec Army Claude Bolton
August 16, 2005



Ultra-Large-Scale (ULS) Systems Study

Gather leading experts to study:

- characteristics of ULS systems
- challenges and breakthroughs required
- promising research and approaches

Intended outcomes:

- ULS System Research Agenda
- program proposal
- collaborative research network

About the Effort

Funded by the Army (ASA ALT)

*Staffing: 9 member SEI team
13 member expert panel*

Duration: one year (04/05 -- 05/06)



SEI Team

Linda Northrop

Software Design and Architecture,
Software Product Lines

Rick Kazman

Adaptive Architectures, Design Methods

Peter Feiler

Methodologies,
Configuration Management

Mark Klein

Real-time Performance Analysis,
Software Architecture Design and Analysis

John Goodenough

Software Reliability,
Safety Assurance

Mark Pleszkoch

Rigorous Software Engineering Methods

Rick Linger

Rigorous Software and
Systems Engineering

Kurt Wallnau

Software Components,
Program Generation, Language Semantics

Tom Longstaff

Security and Survivability Engineering
in Complex Systems

Bill Pollack

Chief Editor

Daniel Pipitone

Chief Graphical Designer



SEI Team (and Expert Panel)

Gregory Abowd

Georgia Institute of Technology

Peter Neumann

SRI International Computer Science Laboratory

Carliss Baldwin

Harvard Business School

Douglas Schmidt

Vanderbilt University

Mary Shaw

Carnegie Mellon University

Bob Balzer

Teknowledge Corporation

Richard P. Gabriel

Sun Microsystems

Dan Siewiorek

Carnegie Mellon University

Gregor Kiczales

University of British Columbia

Kevin Sullivan

University of Virginia

Ali Mili

New Jersey Institute of Technology

John Lehoczky

Carnegie Mellon University

Jack Whalen

PARC



How This Study Is Different?

It presents an overall research agenda -- not just for new tools or a new software method.

It is based on the challenges associated with ultra-large scale.

It focuses on the future.


It involves an multi-disciplinary base.

It takes a fresh perspective on the development, deployment, operation, and evolution of software-intensive systems.

Germes of these ideas are present today in small research pockets; these efforts are currently too small to have much impact on next-generation DoD ULS systems.



ULS Systems Research Agenda




Describes

- the characteristics of ULS systems
- the associated challenges
- promising research areas and topics

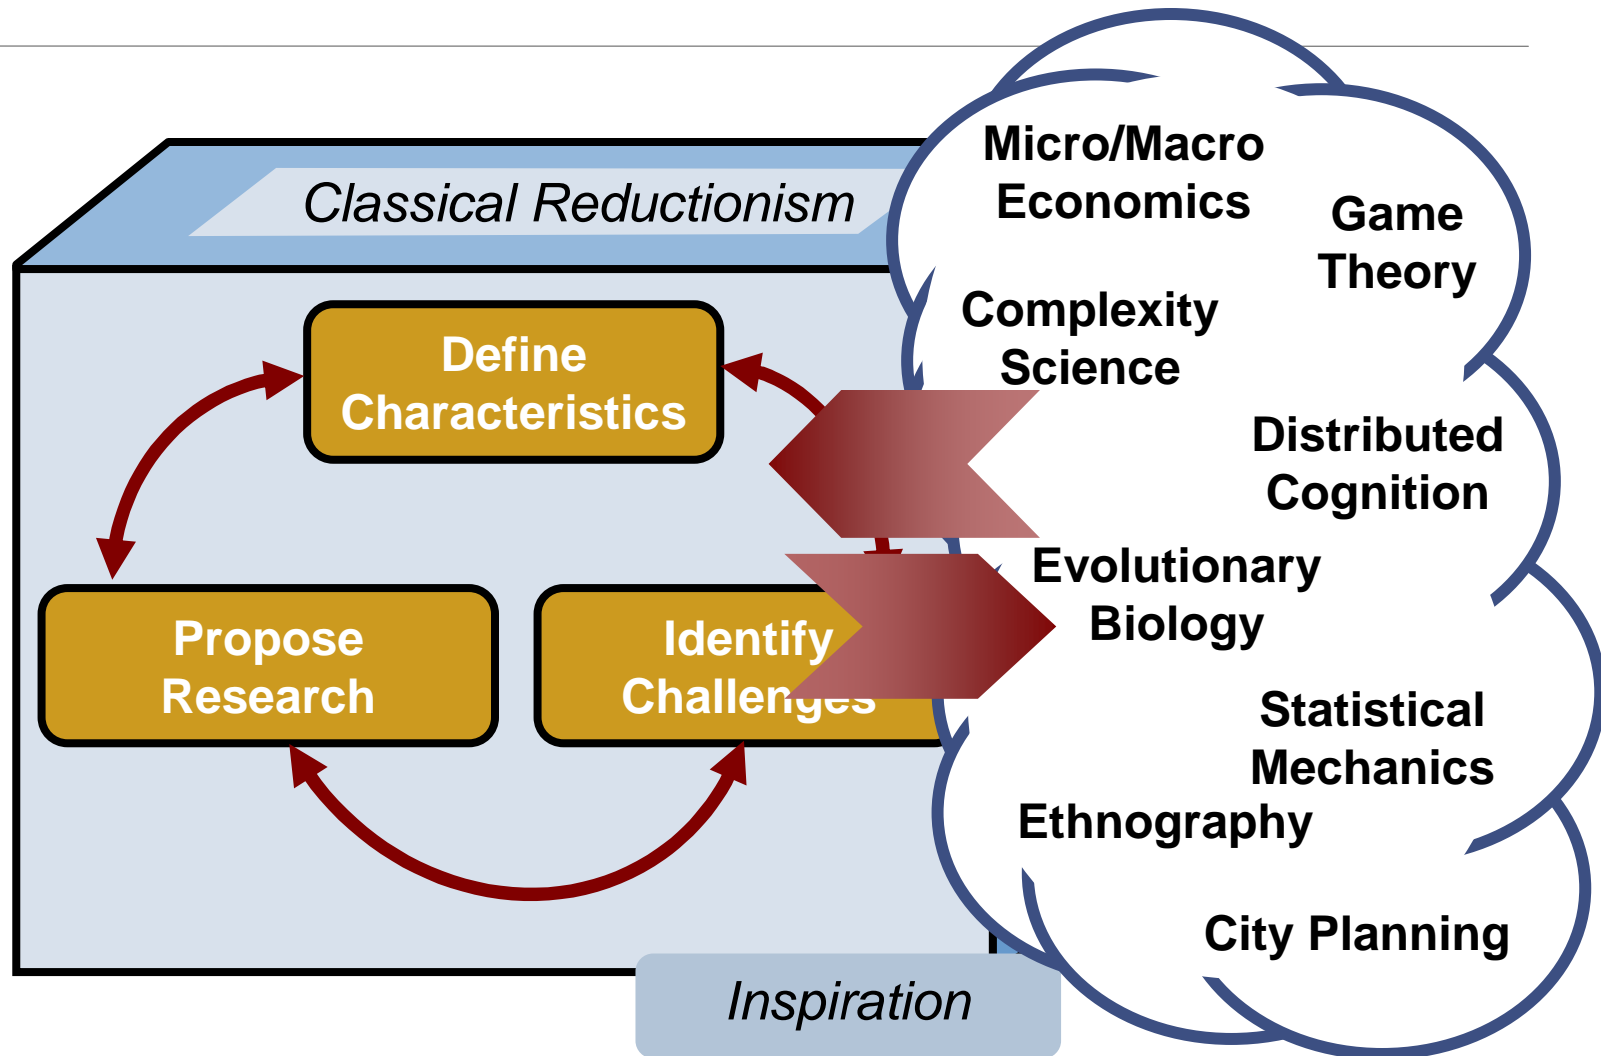
Is based on a new perspective needed to address the problems associated with ultra-large-scale systems.



Working Inside and Outside the Box




Working Inside and Outside the Box




Ultra-Large Scale

Ultra-large size in terms of

- Lines of code
- Amount of data stored, accessed, manipulated, and refined
- Number of connections and interdependencies
- Number of hardware elements
- Number of computational elements
- Number of system purposes and user perception of these purposes
- Number of routine processes, interactions, and “emergent behaviors”
- Number of (overlapping) policy domains and enforceable mechanisms
- Number of people involved in some way
-



What Is an Ultra-Large-Scale (ULS) System?




A ULS System has unprecedented scale in some of these dimensions:

- Lines of code
- Amount of data stored, accessed, manipulated, and refined
- Number of connections and interdependencies
- Number of hardware elements
- Number of computational elements
- Number of system purposes and user perception of these purposes
- Number of routine processes, interactions, and “emergent behaviors”
- Number of (overlapping) policy domains and enforceable mechanisms
- Number of people involved in some way

ULS systems will be interdependent webs of software-intensive systems, people, policies, cultures, and economics.

ULS systems are systems of systems at internet scale.





Scale Changes Everything

Characteristics of ULS systems arise because of their scale.

- Decentralization
- Inherently conflicting, unknowable, and diverse requirements
- Continuous evolution and deployment
- Heterogeneous, inconsistent, and changing elements
- Erosion of the people/system boundary
- Normal failures
- New paradigms for acquisition and policy

These characteristics may appear in today's systems and systems of systems, but in ULS systems they dominate.

These characteristics undermine the assumptions that underlie today's software engineering approaches.



Today's Approaches

The Engineering Perspective - for large scale software-intensive systems

- largely top-down and plan-driven
- requirements/design/build cycle with standard well-defined processes
- centrally controlled implementation and deployment
- inherent validation and verification

The Agile Perspective - proven for smaller software projects

- fast cycle/frequent delivery/test driven
- simple designs embracing future change and refactoring
- small teams and retrospective to enable team learning
- tacit knowledge

Today's approaches are based on perspectives that fundamentally do not cope with the new characteristics arising from ultra-large scale.



A New Perspective is Required


“The older is not always a reliable model for the newer, the smaller for the larger, or the simpler for the more complex...Making something greater than any existing thing necessarily involves going beyond experience.”



Henry Petroski


Pushing the Limits: New Adventures in Engineering





Today We Build “Buildings”






We Need To Think Cities





We Need to Think Ecosystem




Diverse users with complex networked dependencies and intrinsic adaptive behavior

Has:

- Robustness mechanisms: achieving stability in the presence of disruption
- Measures of health: diversity, population trends, other key indicators

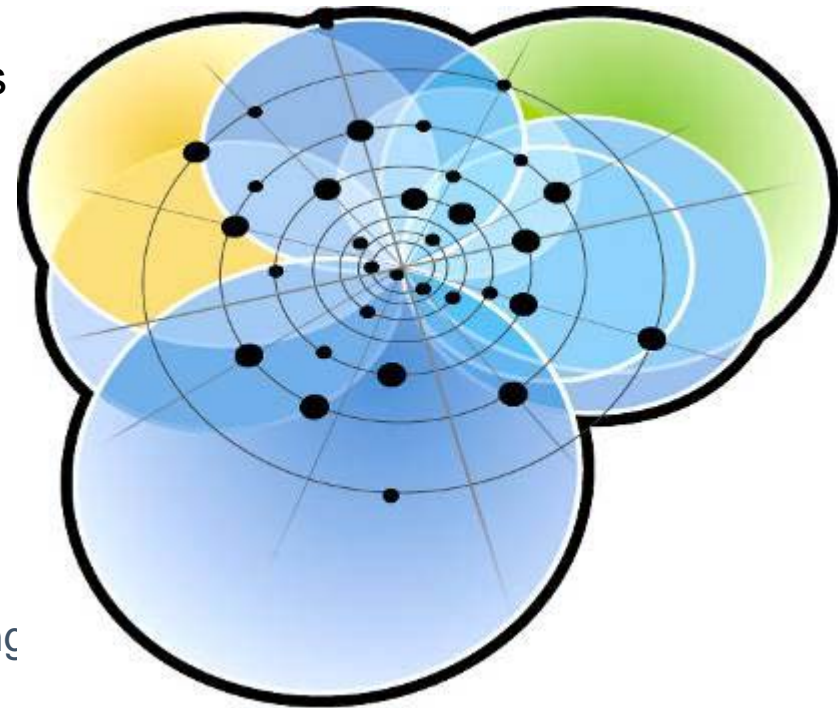





We Need to Think Socio-Technical Ecosystems

Socio-technical ecosystems include people, organizations, and technologies at all levels with significant and often competing interdependencies.

- There will be competition for resources.
- There will be organizations and participants responsible for setting policies.
- There will be organizations and participants responsible for producing ULS systems.
- There will need to be local and global indicators of health that will trigger necessary changes in policies and in element and system behavior.






Why a New Perspective?

There are fundamental assumptions that underlie today's software engineering and software development approaches that are *undermined* by the characteristics of ULS systems.

There are challenges associated with ULS systems that today's perspectives are very unlikely to be able to address.






ULS Systems vs Today's Approaches - 1

Characteristics	Today's assumptions undermined
Decentralized control	All conflicts must be resolved and resolved centrally and uniformly.
Inherently conflicting, unknowable, and diverse requirements	Requirements can be known in advance and change slowly. Tradeoff decisions will be stable.
Continuous evolution and deployment	System improvements are introduced at discrete intervals.
Heterogeneous, inconsistent, and changing elements	Effect of a change can be predicted sufficiently well. Configuration information is accurate and can be tightly controlled. Components and users are fairly homogeneous.





ULS Systems vs Today's Approaches - 2

Characteristics	Today's assumptions undermined
Erosion of the people/system boundary	People are just users of the system. Collective behavior of people is not of interest. Social interactions are not relevant.
Normal failures	Failures will occur infrequently. Defects can be removed.
New paradigms for acquisition and policy	A prime contractor is responsible for system development, operation, and evolution.



Challenges

ULS systems will present challenges in three broad areas:

- Design and evolution
- Orchestration and control
- Monitoring and assessment




Design and Evolution



Specific challenges in ULS system design and evolution stemming directly from the characteristics of ULS systems:

- Economics and industry structure
- Social activity for constructing computational environments
- Legal issues
- Enforcement mechanisms and processes
- Definition of common services supporting the ULS system
- Rules and regulations
- Agility
- Handling of change
- Integration
- User-controlled evolution
- Computer-supported evolution
- Adaptable structure
- Emergent quality






Orchestration and Control

Orchestration and control refers to the set of activities needed to make the elements of a ULS system work together in reasonable harmony to ensure continuous satisfaction of mission objectives.

Orchestration is needed at all levels of ULS systems and challenges us to create new ways for

- Online modification
- Maintenance of quality of service
- Creation and execution of policies and rules
- Adaptation to users and contexts
- Enabling of user-controlled orchestration





Monitoring and Assessment

The effectiveness of ULS system design, evolution, orchestration, and control has to be evaluated.

There must be an ability to monitor and assess ULS system state, behavior, and overall health and well being.


Challenges include

- Defining indicators
- Understanding why indicators change
- Prioritizing the indicators
- Handling change and imperfect information
- Gauging the human elements





Where Do We Focus Our Research

- Address the predominant characteristics of ULS systems and the three challenge categories.
- Look for breakthroughs not incremental improvement in current approaches.
- Take a more expansive view of software research and include its interactions with associated research in the physical and social sciences.



Potential Tools: Game Theory



Algorithmic Mechanism Design



- games + microeconomics + computation
- computational markets for any scarce ULS resource?

Institution Design

- learning games + self-reinforced expectations + cultural norms
- better formal models of acquisition in non-prime-dominated landscape?



Networks, Statistical Mechanics, Complexity



HIV Partners



Stability

Networks Are Everywhere

Recurring “scale free” structure

- internet & yeast protein structures

Analogous dynamics?

- epidemiology, robustness and vulnerability

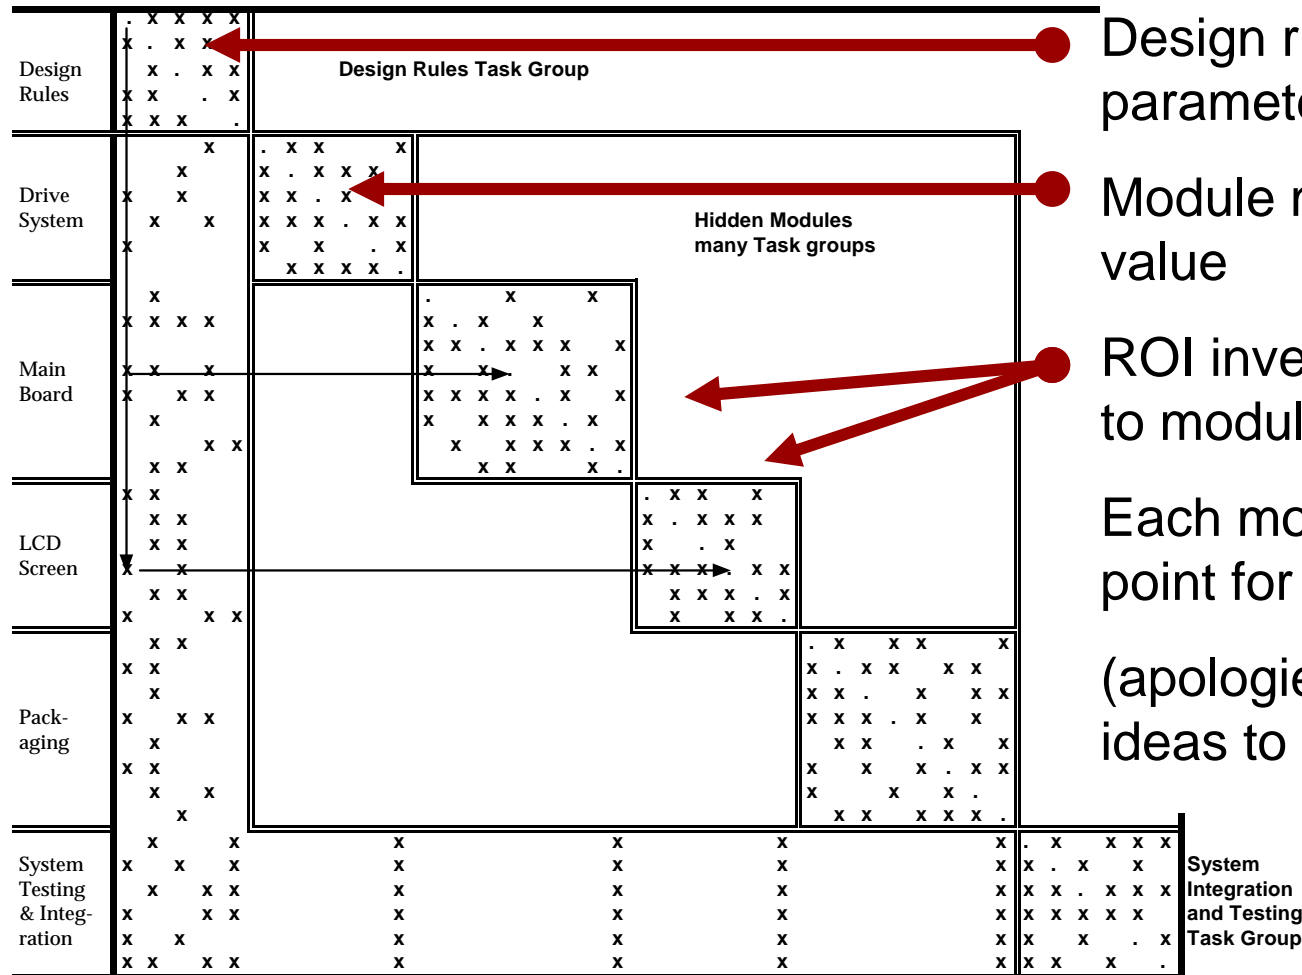
Unstable Equilibrium

How many changes before a system becomes unstable?

What scale and frequency of disruptions can be expected?



Economics (Finance) As Design Criteria



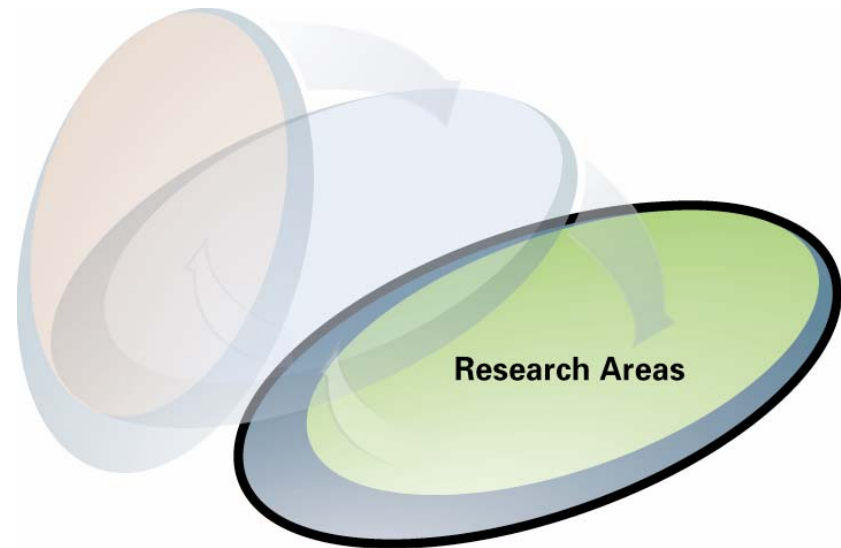
- Design rules (feature parameterization)
- Module maximizes option value
- ROI inversely proportional to module footprint
- Each module a potential point for competition
- (apologies for abuse of ideas to Carliss Baldwin)




Research Portfolio

We recommend an interdisciplinary portfolio of seven research areas and suggested topics for breakthrough research needed to meet the challenges associated with ULS systems.

- Is not expressed in terms of today's "hot" technologies.
- Does not supplant current software research.
- Expands today's horizons.




Research Areas - 1



6.1 Human Interaction: Involves anthropologists, sociologists, and social scientists conducting detailed socio-technical analyses of user interactions in the field, with the goal of understanding how to construct and evolve such socio-technical systems effectively.

- Context-Aware Assistive computing
- Understanding Users and Their Contexts
- Modeling Users and User Communities
- Fostering Non-Competitive Social Collaboration
- Longevity










Research Areas - 2

6.2 Computational Emergence:

Explores the use of methods and tools based on economics and *game theory* (e.g., *mechanism design*) to ensure globally optimal ULS system behavior by exploiting the strategic self interests of the system's constituencies; explores *metaheuristics* and *digital evolution* to augment the cognitive limits of human designers.


- Algorithmic Mechanism Design
- Metaheuristics in Software Engineering
- Digital Evolution



A GAME THEORY ①	A GAME THEORY ②	A GAME THEORY ③	A GAME THEORY ④
			
Oskar Morgenstern 1902 – 1987	John von Neumann 1903 – 1957	John F. Nash, Jr. 1928	Reinhard Selten 1913
Monograph Efficiency: 5.0	Monograph Efficiency: 3.7	Monograph Efficiency: 6.8	Monograph Efficiency: 6.2
Standard Rate of Return: 6.9	Standard Rate of Return: 12.0	Standard Rate of Return: 6.8	Standard Rate of Return: 15.0
Field: Peer-review Index: 1.3	Field: Peer-review Index: 22.7	Field: Peer-review Index: 17.1	Field: Peer-review Index: 2.7
Prevalence/Prevalent Index: -0.24	Prevalence/Prevalent Index: -0.1	Prevalence/Prevalent Index: +0.3	Prevalence/Prevalent Index: +11.6
Dependent Utility: 1.0	Dependent Utility: 1.0	Dependent Utility: 0.7	Dependent Utility: 0.9




Research Areas - 3




6.3 Design: Broadens the traditional technology-centric definition of design to include people and organizations; social, cognitive, and economic considerations; and design structures such as *design rules* and government policies.

- Design of All Levels
- Design Spaces and Design rules
- Harnessing Economics to Promote Good Design
- Design Representation and Analysis
- Assimilation
- Determining and Managing Requirements



Research Areas - 4




6.4 Computational Engineering: Focuses on evolving the expressiveness of representations to accommodate the semantic diversity of many languages and focuses on providing automated support for computing the evolving behavior of components and their compositions.

- Expressive Representation Languages
- Scaled-Up Specification, Verification, and Certification
- Computational Engineering for Analysis and Design




Research Areas - 5



6.5 Adaptive System Infrastructure: Investigates integrated development environments and runtime platforms that will support the decentralized, “always-on,” nature of ULS systems as well as technologies, methods, and theories that will enable ULS systems to be developed in their deployment environments.

- Decentralized Production Management
- View-Based Evolution
- Evolutionary Configuration and Deployment
- In Situ Control and Adaptation






Research Areas - 6

6.6 Adaptable and Predictable System Quality: Focuses on how to maintain quality in a ULS system in the face of continuous change, ongoing failures, and attacks and how to identify, predict, and control new indicators of *system health* (akin to the U.S. gross domestic product) that are needed because of the scale of ULS systems.

- Robustness, Adaptation, and Quality Attributes
- Scale and Composition of Quality Attributes
- Understanding People-Centric Quality Attributes
- Enforcing Quality Requirements
- Security, Trust, and Resiliency
- Engineering Management at Ultra-Large Scales



Research Areas - 7



6.7 Policy, Acquisition, and Management: Focuses on transforming acquisition policies and processes to accommodate the rapid and continuous evolution of ULS systems by treating suppliers and supply chains as intrinsic and essential components of a ULS system.


- Policy Definition for ULS Systems
- Fast Acquisition for ULS Systems
- Management of ULS Systems



(The Selling Of) Mechanism Design: Game Theory + Microeconomics + Computation

Relevance to DoD Missions

- ULS systems must support war fighters at all echelons who are engaged in information-intensive activities and who must share critical but finite information technology resources.
- Future combat missions will require robust and decentralized resource allocation mechanisms that are strategy proof, support a diversity of interests, and provide fully predictable and near-optimal global outcome.



Key Concepts

- *Game theory* provides mathematical tools to study the outcomes of interactions among self-interested, and possibly deceptive, players, where the interactions are governed by a set of rules.
- *Mechanism design* is the inverse of game theory: it seeks to discover the rules of games that will result in a desired outcome despite self-interested and deceptive behavior.
- Mechanism design is concerned primarily with *microeconomics* the economic behavior of agents in the face of scarcity.



Toward a Roadmap for a ULS Systems Research Program

There are many possible approaches to structuring a research program from the ULS Systems Research Agenda

We provide three possible support structures based on

1. Specific DoD missions and capabilities
2. DoD research funding types required
3. Estimates of the relative starting points of the research

We expect that sponsors with different needs will likely choose to support different combinations of research and perhaps different paths through (or projects within) the research program.

The envisioned outcome of the proposed research is a spectrum of technologies and methods for developing ULS systems, with national-security, economic, and societal benefits that far extend beyond ULS systems themselves.



Study Conclusions

There are fundamental gaps in our current understanding of software development at the scale of ULS systems.

These gaps

- present profound impediments to the technically and economically effective achievement of the DoD goal* of deterrence and dominance based on information superiority
- require a broad, fresh perspective and multi-disciplinary, breakthrough research

We recommend

- a ULS Systems Research Agenda that includes research areas based on a fresh perspective aimed at challenges arising from increasing scale
- short-term startup program and a long-term, substantive research program for ULS systems

** As stated in the Quadrennial Defense Review (QDR) Report, Feb 2006*



ULS Systems Research Study Report

Acknowledgements

Executive Summary


Part I

1. Introduction
2. Characteristics of ULS Systems
3. Challenges
4. Overview of Research Areas
5. Summary and Recommendations

Part 2

- 6 Detailed Description of Research Areas
- Glossary

<http://www.sei.cmu.edu/uls/>



Software Engineering Institute

Carnegie Mellon

ICGSE 2006: Scale Changes Everything
Linda Northrop, Oct 2006

© 2006 Carnegie Mellon University

The Start of a Collaborative Research Network

As a result of this study, a community of interest in ULS systems and the needed research is already beginning to grow.

- Contributors to this study have begun to describe and advocate ULS system challenges and potential research.
- Over the course of the past year, keynote presentations and invited talks related to this study have been given at a diverse set of forums.
- Workshops and panel discussions are being organized and the term “ultra-large-scale systems” is gaining traction with others outside the original group charged to conduct the study.
- Several organizations and many individuals have expressed interest in getting involved in a ULS Systems Research Program.



What We Learned

- New, multi-disciplinary perspective and new research in building ultra-large-scale systems is long overdue.
- Manifestations of scale and its attendant complexity arise in many disciplines, and can be understood as a phenomenon in its own right.
- The ULS Systems research proposal, if funded, will provide a clearing in which new ideas can be explored.



What's Next

- ULS System Senior Leader Forum
- Initiation of pockets of ULS System Research
- Promulgation of ULS System Ideas
 - Lectures
 - Talks
 - Workshops
 - Publications
- Lobbying for a ULS System Research Program



Thanks To Those Who Made This Report Possible

Report Author Team:

Peter Feiler, Richard P. Gabriel, John Goodenough, Rick Linger, tom Longstaff, Rick Kazman, Mark Klein, Douglas Schmidt, Kevin Sullivan, Kurt Wallnau, Bill Pollak (Chief Editor), Daniel Pipitone (Information Designer)

Support System:

Hon Claude Bolton, Paul Nielsen (SEI CEO), Clyde Chittister (SEI COO), Hal Stevens (SEI/Army Liaison), Jim Linnehan (Army/SEI Liaison)



thank you
again



Contact Information

Linda Northrop

Director

Product Line Systems Program

Telephone: 412-268-7638

Email: lmn@sei.cmu.edu

U.S. Mail:

Software Engineering Institute

Carnegie Mellon University

Pittsburgh, PA 15213-3890

SEI Fax: 412-268-5758

