CRICIS:

CRITICAL REAL-TIME COMPUTING AND INFORMATION SYSTEMS

A Report from a Community Workshop

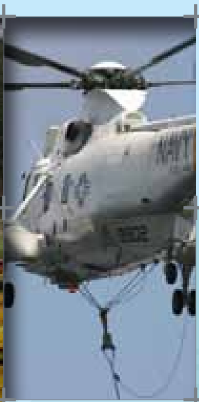

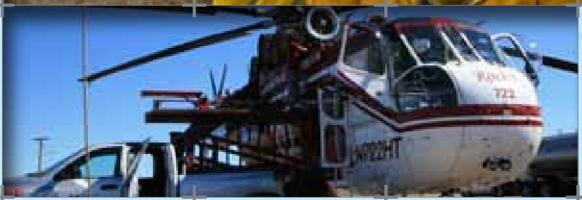

www.cra.org/ccc/

CRICIS EXECUTIVE SUMMARY

We propose a new research area—Critical Real-time Computing and Information Systems (CRICIS)—that has the potential to revolutionize our nation's preparedness and resilience in the face of future disasters by adopting a computational perspective to fundamental scientific, engineering, and social barriers in disaster management. Computing is essential for the collection and transformation of data into usable and secure forms of information tailored for disaster professionals, citizens, and their information devices. A revolutionary socio-technical investigation is needed to harness the opportunities offered by the spread of new classes of devices, sensors, networks, and social media in the past decade. The needed research investigation will encompass multiple diverse disciplines and will have to address significant challenges of scale, complexity, and uncertainty.

CRICIS requires **fundamental new research** in *socio-technical systems that enable decision-making for extreme scales under extreme conditions*. This research cuts across physical and engineered artifacts, information technology, and human-computer collaboration. It is an example of the general shift in science and industry from physical devices and computational packages to socio-technical information systems used by a diverse population. CRICIS consists of five themes of cross-disciplinary research: *integrating computing, physical science, and social science; working and comprehending at scale; real-time awareness and modeling; methods and metrics; and training and education.*

The expected societal benefits are profound. CRICIS will save lives, improve the quality of life for the injured, accelerate economic recovery, preserve the continuity of government, and create new jobs and a new sector of the economy. Computing, together with advances in other disciplines such as social science, can revolutionize the training of workers, including returning veterans, and exploit the attractiveness of "helping others" to recruit larger and diverse populations to science, technology, engineering and mathematics (STEM) disciplines.

Disasters pose four broad challenges for fundamental research in computing, which are distinct from existing programs and initiatives at the Federal level:

- **Simultaneous Extreme-Scale Challenges.** Disaster computation must deal with challenge problems that occur at several exceptionally stressful scales simultaneously, including:
  - ***Time:*** Disaster management encompasses preparedness and prevention (past), real-time response (present), and recovery (future). Events can be discrete and short lived, as in the case of a tornado, or long-term, as with climate change.
  - ***Space:*** Damage can be highly localized, e.g., a bridge collapse, or spread over states, e.g., a hurricane or earthquake, while the economic impacts can be national or global.
  - ***Stakeholders:*** Citizens, governments (municipal, county, state and federal), industry, and non-governmental organizations all have roles to play.
  - ***Data:*** Data is heterogeneous, takes many forms and content, comes from different sources, arrives in different volumes at different times, and exhibits different priorities for different phases of the disaster.

- **Computational Complexity Exacerbated by Interdependencies.** The disaster event, as well as the associated response, is one in which the behavior is non-linear and there exist large interdependencies between variables, multiple temporal and spatial scales, and no single optimal solution. This complexity propagates to algorithmic and data complexity and challenges in modeling chaotic systems. Other sources of computational complexity include the need for maintaining privacy and security of the data; the politics, sociology, psychology, and native language issues; and the impact of the resilience of the underlying electrical, communications, transportation, and financial infrastructure.

- **Human-System Integration in Adverse Environments.** Critical Real-time Computing and Information Systems (CRICIS) is about enhancing human capabilities in making decisions for extreme events that challenge comprehension under extreme situations during which decision makers are tired, dirty, and hungry; they may be working remotely and under dangerous, uncomfortable, demanding conditions that are physiologically and cognitively disrupting. Research is needed in sensemaking, comprehension, visualization, trustworthy data, decision support tools, and determining physiological and cognitive impacts.

- **A Different Style of Research.** The manner in which research on Critical Real-time Computing and Information Systems (CRICIS) is conducted is distinctive in at least three ways. Similar to the cross-cutting Smart Health and Wellbeing program at the National Science Foundation, which aims to tackle the myriad challenges of health care delivery within a large, expensive, *and ever-evolving system, Computing for Disasters research is holistic and conducted almost solely in the context of the domain—and thus it is essential to have peer-partnerships with stakeholders. In addition, it relies primarily on empirical methodologies.*

Dramatic advances in disaster computation are possible with a long-term research investment to facilitate and sustain an organic research community. A living roadmap would create a common ground for stakeholders and researchers and serve as a basis to define and evangelize desirable computing capabilities, lower the bar of entry for new researchers, and encourage adoption of technologies by interested agencies and NGOs. A portfolio is envisioned comprised of *traditional small, medium,* and *large grants* with *regional centers* that would connect stakeholders with researchers and provide testbeds of sufficiently high fidelity and scale to support formative experimentation and evaluation of progress.

1. IMPORTANCE OF CRITICAL REAL-TIME COMPUTING AND INFORMATION SYSTEMS (CRICIS)

The purpose of this report[1] is to capture the role and benefits of computing in disaster management, articulate the fundamental research challenges posed by the extreme scales and complexity of disasters, and provide a roadmap to inform investments. Disasters are often viewed as an application domain for computing, in which data must be gathered, transmitted, transformed, and presented to stakeholders so that they can make decisions. However, the unique attributes of disasters (extreme temporal, geographical, stakeholder, and data scales, unpredictability, compromised wireless and physical information infrastructures, lack of access to the disaster sites, distributed sources of information and expertise, etc.) pose hidden computing challenges in real-time data collection, integration, and preservation, privacy and security, visualization, analysis, and predictive capabilities for on-the-fly decision-making. Addressing these challenges can not only advance emergency response and recovery, but also drive forward areas of computer and information science and engineering.

1.1 Disasters

Figure 1 The four phases of a disaster.
Disasters are often associated by the Public with immediate life-saving or mitigation efforts, leading to "disaster response" or "disaster management" being used as the overall term for what are actually four phases of a disaster: *prevention, preparation, response, and recovery.*

[1]. Each phase conjures up expectations of different agencies. For example, the U.S. Forest Service created Smokey the Bear to encourage prevention of forest fires; state Departments of Transportation direct evacuations in preparation for an event; response is synonymous with urban search and rescue and National Guard teams; and recovery is associated with insurance inspectors, business recovery loans, and medical or mental health workers. The four phases are not independent and sequential; indeed response and recovery operations start instantaneously, and sub-populations have different long-term recovery needs. Life-saving response operations rarely last beyond 10 days, though local officials may be reluctant to declare the response phase over. Recovery operations can go on for months and are generally marked by when roads, schools, and hospitals are re-opened. Economic and public health recovery often takes years beyond that.

1.2 What is CRICIS?

CRICIS addresses all four phases of disasters by fostering a systems approach to enabling timely and effective decision making by all stakeholders. As highlighted by the President's Council of Advisors on Science and Technology (PCAST) [2, 3] and the Computing Research Association [3], disaster response constitutes a societal Grand Challenge. Over a dozen studies since 2003 by PCAST [2, 4], the National Science and Technology Council

[1] This report reflects findings from a Critical Real-time Computing and Information Systems (CRICIS) Workshop held April 24-25, 2012, in Arlington, Virginia, and subsequent feedback from the computing community. The workshop was co-sponsored by the National Science Foundation (NSF) and the Computing Community Consortium (CCC; http://cra.org/ccc). A complete list of workshop participants appears in the Acknowledgments.

(NSTC) [5], and the National Academies [6-13] have concluded success will not be found in component technologies, manpower, or physical resources by themselves. While advances in wireless networks, unmanned systems, embedded sensors, pattern recognition, surface reconstruction, data fusion, and scheduling algorithms can be expected to continue and even accelerate, these advances individually will not necessarily result in usable information or better decisions. Rather each study has found that success will come from systems approaches to developing the information for timely decision making and synchronizing the flow of this information to the shifting demands of disasters and resources.

Disasters, and by extension CRICIS, are already a national priority, and the increasingly negative impact of disasters will amplify the visibility and urgency of handling disasters. The most recent World Disasters Report [14] found that, between 2000 and 2009, a total of over 7,000 disasters resulted in 1.1 million people casualties worldwide, affected another 2.5 million directly, and yielded a loss of $986.7 billion. The number of disasters each year during this 10-year period remained relatively constant, around 700. However, the impact of the disaster was a function of whether the event occurred in an urban area, not the sophistication of the country. The report argues that deaths and impacts of disasters will continue to increase due to the increasing numbers of people living in concentrated urban environments, and because of the consequences of damage to urban structures where people live and work (e.g., office and apartment buildings) and that serve to mitigate disasters (e.g., hospitals, transportation and other critical infrastructure).

1.3 The Time is Now

The nation, and the world, stands at an inflection point today. Disasters have long impacted society, resulting in mass casualties and huge financial tolls. However, key advances in computing over the past several decades—smartphones, worldwide mobile Internet access, social media, ultra-cheap small computers, sensors, visualization, human-computer interaction, artificial intelligence, robotics, Cloud computing, high-performance computing, ultra low cost storage, data mining, information retrieval, machine learning, geospatial databases, computing for economics, and game theory—position the U.S. at the cusp of a watershed of advances in this revolutionary socio-technical discipline.

2. WHY CRICIS IS DIFFERENT

CRICIS is different from existing research areas and paradigms. It transcends current research in topics such as artificial intelligence, evolutionary network dynamics, GIS, human-centered computing, natural language understanding, data mining, digital libraries, remote sensing, visualization, universal usability, unmanned systems, and wireless networks by exploring how these disciplines address extreme (and dynamically changing) temporal, geographical, stakeholder, and data scales. CRICIS has at least five distinct, cross-cutting themes that shift research from reductionism and static models to a holistic focus *on how dynamic socio-technical systems work, how stakeholders can comprehend data at scale, how models can be adapted in real-time, and how to effectively train and educate the population to exploit technical improvisation* in order to respond to disasters.

2.1 Expands Fundamental Research Applied to Disasters

Many of the areas of exploration described below have a long history of fundamental research, sometimes motivated by disaster situations (e.g., unmanned systems and social networks). However, as illustrated in Figure 2, disasters *create extreme scales and complexity for those areas of computing and require pervasive integrative human-computer systems* approaches. Thus, Critical Real-time Computing and Information Systems (CRICIS) will extend and expand on research in these areas, benefiting both the disaster management community and providing feedback and new directions for the computing research community.

2.1.1 Computing at the Extremes

CRICIS must operate over ***extreme temporal, geographical, stakeholder, and data scales.*** Worse yet, the scales of operation cannot be predicted because they depend on the specifics of the disaster (e.g., type, duration, geographic area impacted, population size affected, etc.) and the phase of consideration (prevention, preparedness, response, and recovery). Temporally, disasters start before the incident, have intense data-driven decision demands in the first weeks following the incident, and then taper off over the years, with gradual economic recovery. Geographically, disasters may directly impact large regions, with secondary impacts at the national or international scale (such as the loss of Taiwanese chip manufacturing facilities in 1999) or engage external resources and expertise. However, disasters may be geographically limited, such as the September 11, 2001, World Trade Center collapse. The number of stakeholders also varies from a small population affected by the 2011 Joplin Missouri tornado or the diaspora following Hurricane Katrina in 2005. The data elements may also be extreme. The input to algorithms or sensors may be huge; consider the number of cell towers to restore, bridges or nodes in critical infrastructure to manage, houses to be inspected, or people (and cars) to evacuate. The data itself may be extreme, especially given the staggering amount of information (and misinformation) available from social media (i.e., Twitter, Facebook, Flickr, YouTube) and other tools and services with user-generated content.

Of these scales, data and the ensuing challenges of *extreme, diverse data integration* are especially important to CRICIS. Data is critical for predicting infrastructure and social vulnerabilities, preventing disasters, saving lives and mitigating damage, and minimizing recovery time. However, the needed information must be extracted from extremely diverse data sets; usually one type of data is not sufficient and thus stakeholders must contend with heterogeneous data sets. These data sets may be diverse in terms of time, ranging from historical records (e.g., last census, analyses of previous responses, etc.) to real-time social media feeds (e.g., tweets, Facebook posts, volunteered geographic information from NGOs such as ushahidi, Crisismappers, etc.). The timing of the data is often unrelated to the priority of the data. The type of GIS data varies as well, as Public Safety databases may use place names, street addresses, and property owners, expecting responders to rely on street signs – but these street signs may be swept away. The type of GIS data also varies, with GPS and map coordinates being expressed in diverse reference systems such as latitude-longitude and state-plain coordinates. Types of data content vary as well, ranging from numbers and text to photographs and videos. The sources for data types vary in institutional heritage ranging from government, private section, non-profit organization and recently formed volunteer groups,

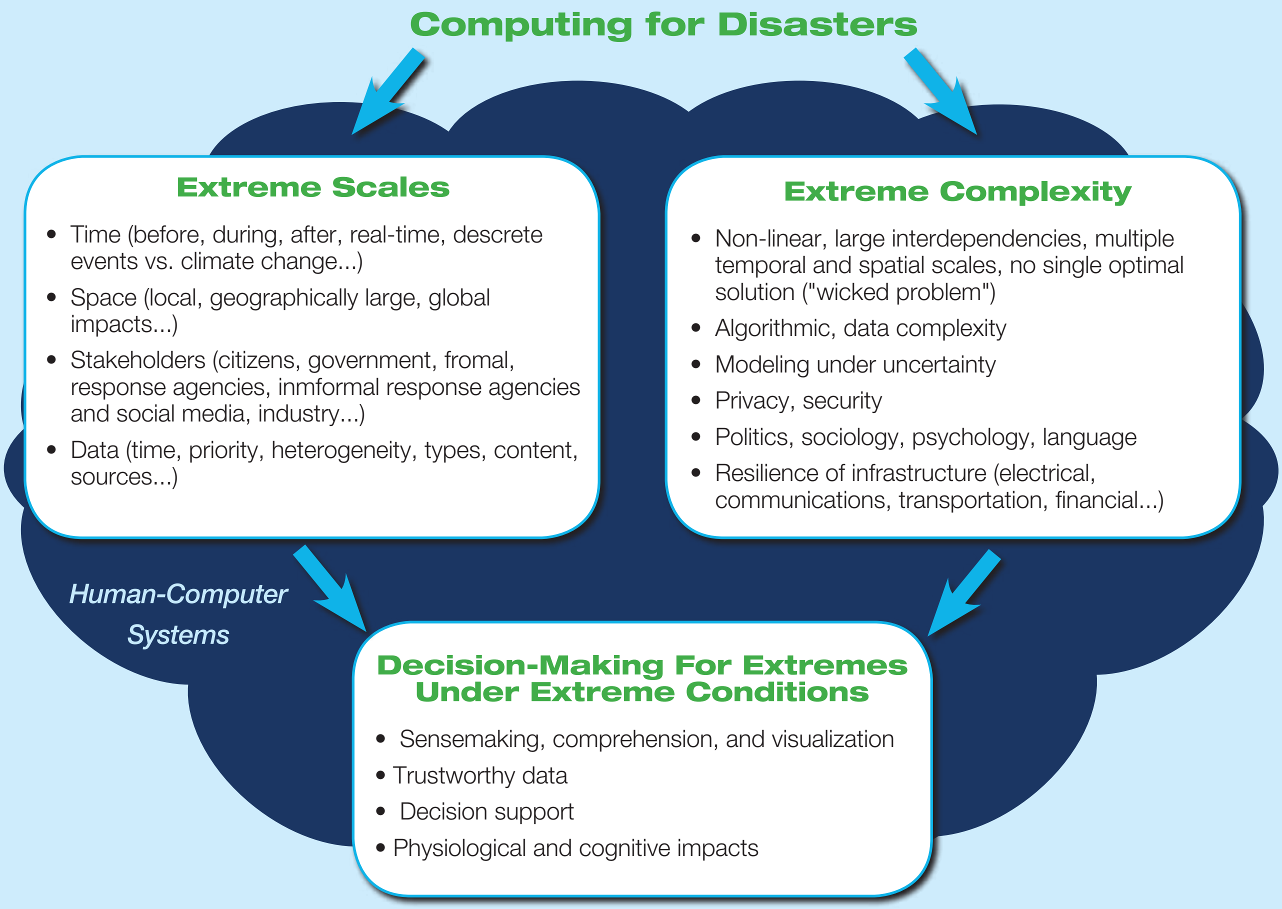

Figure 2.
How Critical Real-time Computing and Information Systems (CRICIS) Expands Fundamental Research through Extremes and Human-Computer Systems.

which may include a diverse set of languages, ethnicities, cultures, motivations, etc. These sources vary widely in accuracy and credibility requiring a diversity of validation methods including automated methods (e.g., voting), manual methods (e.g., site visited by a trusted personnel) and a combination of these methods.

CRICIS must also address the behavioral ***complexity*** within a system. Disasters do not have a simple solution or a single technological fix; they are complex events with large interdependencies, multiple temporal and spatial scales that exhibit nonlinear behavior, and no single optimal solution or analytic framework. This complexity poses a major challenge for modeling and simulation. The complexity challenges are exacerbated by the extreme scales diversity of data and time that computations have to work over. Complexity can be increased by the lack of resiliency in the critical infrastructure, which disrupts interdependencies. For example, large disasters such as Hurricane Katrina and the Haiti Earthquake introduced a large number of faults ranging from electrical power failure to building and transportation infrastructure decimation; these cascaded into computer failures, data-center failures, communication network failures, financial networks outage (automatic teller machines), telephone network outage, more transportation failures (gas pumps), etc. Even a resilient infrastructure may not be able to adequately cope with a disaster; consider that celluar communication networks may survive or cell towers on wheels (COWS) be rapidly deployed, but these are unable to handle the volume of high-resolution maps, images, and video. Complexity also is evident in recovery, as recovering from an event involves diverse resource and scheduling, as captured in [15]. Disaster managers need local data collection to assess situation on the ground and local supply chain to distribute relief supplies (e.g., food, water, fuel, shelter, medicines, etc.) to affected population.

Extreme complexity also arises from the *human element* in disasters. Data about people, their location, income, medical condition, etc. entails issues of privacy and security. States, counties, and municipalities have different policies and attitudes towards disasters. For example, the politics of privacy has led to situations such as those at Hurricanes Charley and Ike when data from functioning traffic cameras could not be used by state transportation officials to monitor evacuation on surface roads or by law enforcement to monitor for looting due to restrictions in some municipalities that such data be limited to specific traffic infractions. Human complexity arises because of diversity in sociology, psychology, and language. Both Florida and Texas have large Spanish-speaking populations; the sociology of these populations leads to a different style of communication of public risk and different patterns of behavior during and after a disaster.

CRICIS must not only handle these extreme scales and complexities but also ensure technology supports *decision-making under extreme conditions.* Decisions during a disaster have significant life-saving and economic

consequences, introducing stress. Decision-making is harder because of the extreme conditions. Formal stakeholders such as responders will have cognitive deficits due to sleep deprivation, disruption of routine, having to work in unfamiliar settings and/or with technology or data sources, and in many cases working remotely through computer mediation. Informal stakeholders such as citizens in the disaster area may be experiencing cognitive deficits due to emotional upheaval.

CRICIS is distinct from existing research areas because of the challenges in managing scales and in comprehension. The safety-critical, accelerated decision-making and concentrated problem solving for disasters complicates the problem of individual and collective comprehension. CRICIS is needed to support rapid sensemaking for the current phase of the disaster and for a particular stakeholder, with inclusion of trustworthy data from any number of sources, both machine sourced, and human sourced (and often technologically-mediated). CRICIS is not only about improving the acquisition or quality of human-sourced or machine-sourced data but rather about getting, transforming, and tailoring the entire set of relevant data.

2.1.2 Human-Computer Systems at the Extremes

CRICIS is also distinct from existing research because of the view of computing for disasters within a human-computer system framework. Because CRICIS is about enhancing human capabilities, the technology must be tailored to the stakeholders' needs. Stakeholder needs vary significantly; the erroneous assumption that disaster response is isomorphic with military command and control functions have resulted in brittle systems that do not address the real needs and constraints of formal responders. Likewise, the oversimplification of Citizens as stakeholders within a rational social system ignores how people really behave. Thus CRICIS research is characterized by the pervasive presence of human-computer system principles and methods.

A fundamental principle of CRICIS is that *the technology must be tailored to the human stakeholders and needed capabilities.* For example, a robust low-bandwidth network is of little value to responders attempting to share high-resolution imagery early in the initial phase of the response, but can be of great value for long-term monitoring of evacuation and return patterns, environmental sensing, etc. A sophisticated algorithm that produces a highly optimal, but counter-intuitive, allocation of resources may be rejected by an emergency support function manager if the algorithm does not provide a means to explain or justify its recommendation. Technology is not sufficient; technology must be designed and evaluated in terms of the larger socio-technical context.

Another fundamental principle is that *formal stakeholders are not the same as military stakeholders.* Investments by the U.S. Department of Defense have produced advances in unmanned systems, sensors, and situation awareness displays. However, this technology was designed for a very different ecology; formal stakeholders such as fire rescue departments have significantly less time set aside for training on new technology than the military and have reduced training budgets that in turn prevent responders from training at scale. Therefore human-computer system research must be involved in extending military solutions to disasters. Another misconception is that disaster response is command-and-control based following a hierarchical decision-making organization. Disasters involved multiple agencies and stakeholders who maintain their own individual hierarchies and are trained separately.

A third principle is that *informal stakeholders, particularly Citizens, have an important role in providing data and consuming information.* Research in disasters has largely been either/or: focused on tools for formal stakeholders or on tools for informal stakeholders. The assumption has been that formal stakeholders make decisions that help the Public at large while informal stakeholders make decisions for individuals or small groups. However, computing is enabling Citizens to move from the role of passive victim to active collaborator in disaster management. CRICIS can foster crossing these boundaries, for example, enabling valuable information from informal sources such as ad hoc NGOs or social networking reaches formal stakeholders in a trustworthy usable form, and information generated by formal stakeholders is shared with informal stakeholders such as insurance agencies.

The final principle is that *stakeholder behavior, even for formal responders, cannot be modeled as a rationalized social system.* A tenet of persuasive computing is that people do not blindly follow recommendations, even from authority figures. One example was the grid-locked evacuation of the Houston-Galveston region in anticipation of Hurricane Rita. CRICIS favors solutions which both incorporate realistic models of human behavior (versus consistent, compliant behavior) and real-time feedback (what people are actually doing). It is important to keep in mind that formal stakeholders also behave unexpectedly. Formal stakeholders often take on initiative and unimagined responsibility, such as when the FAA Air Traffic Controllers shut down all air traffic during the 9/11 attacks. CRICIS should enable the needed flexibility and adaptability to permit extraordinary responses to unimaginable situations.

2.2 Poses Unique Computer Science Research Directions

CRICIS contributes to fundamental computer science research by expanding and creating new, distinct research directions that cut across physical and engineered artifacts, information technology, and human-computer collaboration. Five themes of cross-disciplinary research in systems with extreme scales and complexity have been identified: *integrating computing, physical science, and social science; working and comprehending at scale; real-time modeling; methods and metrics; and training and education.*

1. **Integrating computing, physical science, and social science.** CRICIS research is characterized by its explicit focus on the socio-technical system and the creation of novel integrative capabilities. A new embedded sensor for recording damage to a structure from a meteorological event is a great advance in physical science. An algorithm that can in real-time fuse and analyze multiple sensor readings and represent uncertainty in the analysis is a powerful computation advance. A system that applies the fusion algorithm to any set of embedded shock sensors and then is able to display what that means to a city manager (e.g., Red-Yellow-Green tagging of unsafe, work needed, and safe for re-entry) and create trust is a valuable capability.

2. **Working and comprehending at scale.** CRICIS research explicitly addresses the issues of extreme temporal, geographical, and data scales and diversity but goes further to consider how stakeholders comprehend the meaning of the data in the context of their goals and responsibilities. For example, a worthy CRICIS project would create systems for supporting information flow and sense-making among distributed communities with differing levels of technology, education, skills, etc. to build consensus that supports scalable real-time decisions on a particular aspect of the response, such as searchers, city managers and inspectors, insurance agencies, construction firms sharing information to speed inspection, bidding, permitting, repair, and re-opening of critical facilities such as hospitals and schools. Another project might increase the speed and efficacy of actionable situation awareness by integrating sensors, text, voice, and video into a device-aware display.

3. **Real-time Modeling.** Accurate models of the physical, social, and economic impacts of disasters are essential to prevention and preparedness for disasters; these models also initialize the response and recovery activities. However, for a large-scale problem with complex interdependencies, a priori models are unlikely to be completely accurate or inclusive of all relevant situations. Therefore, CRICIS research focuses on how to incorporate observations into existing models in real-time or to create new models which can then be used to project damage, vulnerabilities, social consequences, etc. Behavioral models which move away from the rationalized social response and capture the effects and propagation of trust, impact of noisy data, etc., are also needed.

4. **Methods and Metrics.** CRICIS research is distinct from other research areas in part because the methods and metrics do not wholly exist. Socio-technical systems are difficult to reproduce or simulate at scale. Can evaluations be conducted without scale exercises? Existing metrics for task performance, especially with teams of human and agents, cannot assign the credit for performance: was it a helpful technology or was it a mediocre technology operated by a masterful human? CRICIS projects might include creating data-flow models (or new cognitive analysis techniques for modeling) of key data throughout a particular decision-making process.

5. **Training and Education.** Training of stakeholders and the education (or re-education) of professionals and students is a major theme of CRICIS research. Normal stakeholders such as Public Safety agencies have extremely limited training budgets and thus computing for cost-effective training is highly desirable. However, CRICIS is an example of the general shift in science and industry from “widgets” and “computational devices” to socio-technical systems used by a diverse population. Thus, CRICIS should include research in how to educate computing, engineering, and social professionals in socio-technical system design as well as create new programs for graduates and undergraduates. Examples of CRICIS projects would be systems that use familiar and widely available tools such as smart phone apps for facilitating coordination and for simplifying “just in time” training and the use of “serious games” for training. However, it is expected that all CRICIS projects would explicitly consider the limited training time and access to physical facilities as well as how to compensate for the cognitive deficits encountered in the field that can lead to stakeholders discarding tools or procedures.

The types of innovation involved in integrating computing, physical science, and social science for CRICIS include but are not limited to:

- Unmanned systems and new sensors for safe, remote exploration and sensing in structural collapse (e.g., World Trade Center) and chemical, biological, or radiological hazardous materials releases (e.g., Bhopal, Sarin release on the Tokyo subway, Fukushima). In these cases, responders work remotely through technology to perceive and act in the access-denied area.

- Artificial intelligence and computer vision tools expressed through human-computer interaction for rapidly generating, modifying and tracking a "bird's-eye" situational overview of a complex and diffuse event. Without appropriate visualizations of the data, the decision maker may miss critical cues that will either create secondary crises or prolong economy recovery; the benefit of the availability of the data is lost if the system doesn't include how to display it effectively.
- Optimization and search algorithms leveraging collective behavior and action in diffuse events where solutions might be highly localized, or where the search space is large. These algorithms must account for actual human behavior or allow real-time adaptation to observed behavior.
- Re-establishing communication networks under adverse conditions based on dynamically changing priorities and opportunities. Restoring networks must support decision-making priorities; this was seen in Hurricane Andrew where large portions of Florida were stranded without response help because there was no way to communicate the depth of the damage. Algorithms must balance where the needs are rather than restrict themselves to an orderly, frontier-like resumption of services.
- Use of "serious games", computer and mixed-reality simulations to allow responders to train and experiment with new technologies and scenarios on demand.

2.3 Conducts Research Differently

The manner in which research Critical Real-time Computing and Information Systems (CRICIS) is conducted is noteworthy in at least three ways: it is holistic and conducted in the context of the domain; it relies primarily on empirical methodologies; and it is based on partnerships with stakeholders. These are discussed in more detail below.

Is holistic. A reductionist model does not apply to decision-making for disasters, or indeed complex non-linear problems in general. Decomposing the complex physical and social interactions over time and space neglects that it is the interactions that enable rapid, effective decision-making. Rather than extrapolating how an atomic phenomenon tested in isolation will scale, CRICIS research is holistic. It places the research in the context of the socio-technical system, exhibiting a "because gap G in computing for disasters exists and is preventing stakeholder S from capability C, we propose to investigate P to address G…" flavor.

Relies primarily on empirical methodologies. CRICIS research is characterized by its observational nature. Empiricism is appropriate because there is no significant, comprehensive corpus of data (yet) on how computing or information and communication technology has been used in disasters. This corpus would be acquired through detailed ethnographic observation and rigorous case studies, which then leads to the formation of principles and theories that can be examined through field tests with stakeholders in simulated and realistic situations.

Controlled experimentation and simulation is less relevant because, each disaster is unique and thus provides new innovations, bottlenecks, and requirements for computing. If researchers aren't in place to ethically deploy a solution, they should be carefully committed to identifying solutions for innovation and to the study of those identified problems in the event to inform innovation for the next disaster.

Is based on meaningful partnerships with a stakeholder(s). Because CRICIS research is holistic, it must derive requirements from direct contact with practitioners throughout the disaster cycle phases. Involving practitioner organizations will ensure that the NSF's research community begins with realistic requirements, while addressing significant problems. Then as the research develops, feedback from practitioners will guide implementations, sympathetic partners will be available for pilot testing, and mature projects can be appropriately integrated for evaluation in genuine crises. CRICIS projects should require field study partnerships to ensure requirements are accurate and guidance for refinement is readily available, much as bioinformatics programs generally require medical practitioners on the research team. However, university Human Review Committees (HRC) must be trained and prepared to readily guide and approve such research in a rapid time frame to make such field commitments possible.

2.4 Comparison to Other Programs

The previous sections have elaborated on the many ways in which fundamental research at the intersection of computing and disaster management stands to yield advances that will vastly improve the way in which society prepares for, respond to, and recovers from natural and man-made disasters. However, it is important to emphasize that these advances are not simply dependent upon computing; rather, an initiative that brings together researchers from multiple disciplines—to include computer scientists, as well as electrical engineers, civil and transportation engineers, systems engineers, geoscientists, social, behavioral, and economic scientists, and so on—is crucial for achieving the kind of revolutionary breakthroughs that are necessary in this space. Equal partnerships between researchers across these disciplines—combined with disaster stakeholders—must be enabled.

Furthermore, CRICIS integrates areas addressed through existing programs within NSF: physical science and engineering (e.g., sensors, unmanned systems, models, etc.), social sources of data (e.g., social networking, citizen science, etc.), transfer and storage of data (e.g., networks, cloud computing, context-aware networks, data warehousing, security, privacy, etc.), transformation of data into useful information (e.g., computer vision, human computation, etc.), real-time formal and informal decision making support (e.g., artificial intelligence, optimization, visualization, interfaces, cognitive science, CSCW, policy, etc.). Similarly, it operates over large sets of heterogeneous data, a key focus of the recent multi-agency Core Technologies and Technologies for Advancing Big Data Science & Engineering (BIGDATA) solicitation. And it targets human-computer partnerships where computing becomes a key tool in facilitating decision-making.

However, as previously have described, disaster management requires a concerted effort beyond these programs given the unique attributes of scale and complexity. It includes non-digital sources of information, such as the decision-making behavior of the stakeholders, and considers the physical domain-computer relationship, incorporating the physicality of disasters and the technological challenges of obtaining data remotely (in other words, CRICIS constitutes a more focused, but more comprehensively socio-technical, exploration of a single domain).

3 BENEFITS OF CRITICAL REAL-TIME COMPUTING AND INFORMATION SYSTEMS (CRICIS)

Intuitively, CRICIS offers multiple benefits to society by empowering formal and informal stakeholders: computing will enable them to increase lives saved, improve quality of life for the injured, accelerate economic recovery, and preserve the continuity of government. However, CRICIS will also lead to the creation of new jobs and a new sector of the economy, leveraging investments in Defense technologies and expertise of returning veterans, but will also requiring new methods of training and education. It will broaden the perspectives of researchers and encourage fundamental advances in computing for socio-technical systems in general. This section attempts to capture the less intuitive benefits of CRICIS through a set of examples, then a discussion of the benefits to stakeholders, the economy, and research in general.

3.1 Examples of How Computing Can Revolutionize Disasters

Computing can revolutionize multiple facets of the disaster management challenge. The core element of a computational perspective is the construction and analysis of algorithms and systems for the collection, transmission, and transformation of data to provide semantically rich forms of information tailored for stakeholders via multiple interface modalities. The general principle in each of the examples below is that computing cuts across the diverse range of actionable information needs of disaster response.

Computing is a key element in gathering data. Data for disasters is acquired from diverse sources both a priori and in real-time. Maps, building and land-use plans, and predictive models of the location and severity of damage serve to ground prevention and preparedness and to bootstrap response and recovery planning. For example, computing could provide an accurate, timely event detection by fusing data using diverse components including sensors, the acquisition network, the Cloud, and human volunteers and trained staff. After the incident has occurred, the sources of information increase: direct observations of responders; unmanned systems and sensors deployed by government agencies, corporations, and private citizens; photographs and other information uploaded from phones; information in social networks; and news feeds are all in play. Data sources are diverse in type (numerical, audio, video, natural language), format, precision, trustworthiness, responsiveness, and location. Moreover, all these parameters may change with time. Computing can exploit this huge number of data sources, while managing their immense heterogeneity, uncertainty, and security restrictions. Developments in data mining, machine learning, large multimedia archiving and retrieval systems, efficient search, human-computer interfaces and visualization, and statistics enable these challenges to be overcome today whereas they were insurmountable as recently as a decade ago.

Computing can transmit and transform the data into actionable information, while hiding complexities of superfluous information, and help provide accurate situation awareness across all echelons of decision makers. Data from multiple sources can be fused across distributed data-acquisition networks and servers, possibly in Cloud computing systems. Disaster management systems have to operate in environments in which power may be intermittent and communication bandwidth is limited and unpredictable. However, not all data are best represented geospatially; fortunately computing, coupled with perceptual psychology, can match the specific decision-making needs for a particular role with the appropriate visualization. Artificial intelligence can generate commands to machine components (“close the gas valve”) and recommendations to Citizens (“go to the high school for immediate mustering and evacuation”). Computing can provide resilient networking, allowing the right information to get to the right person at the right time, rather than the current approach of sending all the data at the highest resolution to everyone.

Computing can explicitly represent and mitigate uncertainty. A priori knowledge such as maps may be compromised by a meteorological event or a structural collapse. The relative infrequence of disaster events suggests that simulations of damage severity, plume tracking, or evacuation patterns will have areas of inaccuracies or regions where the projections simply do not apply. Computing using probabilistic and other evidential reasoning techniques can explicitly determine the amount of certainty in projections; this data then can be used to temper decisions or to direct gathering of new data. Computing advances in modeling and simulation can create methods to update and refine models in real-time.

Computing can assist with the social, behavioral, and economic consequences, and opportunities, of disasters. Formal preventions, preparations, responses, and recoveries from disasters are orchestrated by multiple agencies–each of which has mechanisms for providing situation awareness to its network of agents and making decisions (what Nobel laureate Elinor Ostrom describes as polycentric control architectures [16]). Informally, Citizens respond as individuals based on their own goals and their knowledge of unfolding situations. Responders (the “tired, dirty, and hungry”) may be cognitively impaired during disasters and reluctant to adopt new technology or trust new sources of data. Citizens may not behave “rationally;” they may not follow recommendations, opting instead to focus on finding their families. The economic consequences of poor situation awareness can be immense. However, few applications allow those on-site to contribute data or for the worldwide participation of people who desperately want to help. Computing, integrated with the social sciences, can develop systems that orchestrate responses and harness the human dimension to changing critical situations in the most efficacious ways.

Computing can optimize resources and logistics, as well as help predict and maintain the critical infrastructure. Computing can provide timely situation awareness and management of resources across large regions of energy, water, terrorism, and logistics. Early warning of disasters or of secondary crises such as epidemics or dwindling fuel reserves are often beyond the ability or scope of agencies to notice until it becomes too late. Efficient transport of responders and relief supplies, repairs to the transportation and energy grids, as well as evacuation and re-entry models can be optimally computed.

Computing can be a tool to reunite families, identify and triage victims, and optimally allocate resources as part of a coordinated real-time response. Advances in computer vision and digital libraries have already resulted in computer programs that allow parents to more rapidly find children who have been dislocated during a disaster in the immediate aftermath of a crisis. During an event response, real-time Disaster ISR (Intelligence, Surveillance, and Reconnaissance) can provide decision-makers and first responders with information vital to resource allocation and triage. Ground-based robot searchers could assess the medical state of victims found inside buildings, while aerial/wide area search could employ coordinated fleets of low-cost UAVs to survey structure damage, navigable transportation routes, and locations of displaced populations. A general principle of Disaster ISR would be to explore how to combine on-board autonomous intelligence and crowd-sourced human-computation from remote participants to effect faster, better medical triage or structure stability estimation.

Computing can revolutionize training of workers. Because disasters are rare and hard to physically duplicate and because government Public Safety budgets are diminishing, training for these events is often inadequate in practice. "Serious games" or training via computer games, online learning, mixed reality simulation, and computer supported workgroup tools can provide more cost-effective and readily available training.

Computing can exploit the attractiveness of helping others to attract larger and more diverse populations to STEM disciplines. Few application domains can capture the interests of young students as much as systems that integrate STEM disciplines to help them deal with disasters that impact their families and friends. From kindergarten onwards, students in earthquake-prone regions, are naturally interested in systems that sense and respond to earthquakes, just as students along the East Coast and Gulf are interested in hurricanes. An excellent way of exciting students about STEM disciplines and information and communication technology is to enable them to participate in direct ways – deploying inexpensive sensors, posting photographs and videos, helping connect families with lost pets in faraway locations.

3.2 Benefits Both Formal and Informal Stakeholders

CRICIS will benefit all stakeholders in a disaster. As shown in Figure 3, disasters involve both formal and informal stakeholders[2]. Formal stakeholders consist of groups with a formal role under the National Response Framework: federal, state, and local government, non-governmental organizations (NGOs), and industry associations. Informal stakeholders include response-related industries (e.g., CrisisMappers, insurance companies, etc.) and industries with key technologies (e.g., telecommunications, networks, robotics, construction, etc.) as well as citizens

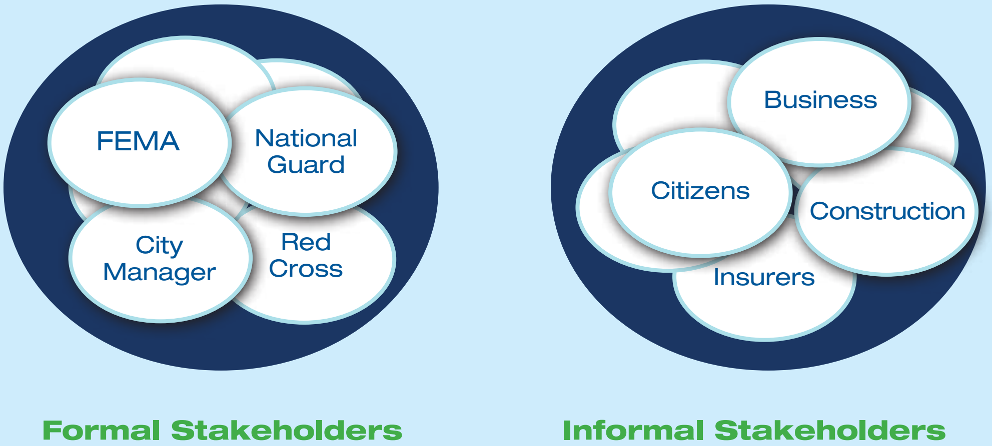

Figure 3.
Two types of stakeholders based on role in disaster response.

[2] It is worth noting that formal and informal does not correspond to government and Public. For example, NGOs such as the American Red Cross have specific responsibilities under the National Response Framework.

Disasters directly impact many informal stakeholders as victims or industries with economic interests, but indirectly impact the formal stakeholders who must make decisions. Figure 4 shows the current relationship of these two classes of stakeholders. In making decisions, formal stakeholders rely upon data provided by projections from damage models, or by observations from manned and increasingly unmanned systems transmitted via wireless networks and displayed in various types of GIS representations. The data and decision making process used by formal stakeholders is opaque to the informal stakeholders, who often only receive information such as evacuation/re-entry postings. The informal stakeholders may exploit social networking and wireless networks to self-organize. Self-organization can also lead to data gathering, often expressed as GIS mash-ups that can in turn be used by all stakeholders. Industry and insurance agencies may be forced to duplicate data collection in order to support their needs.

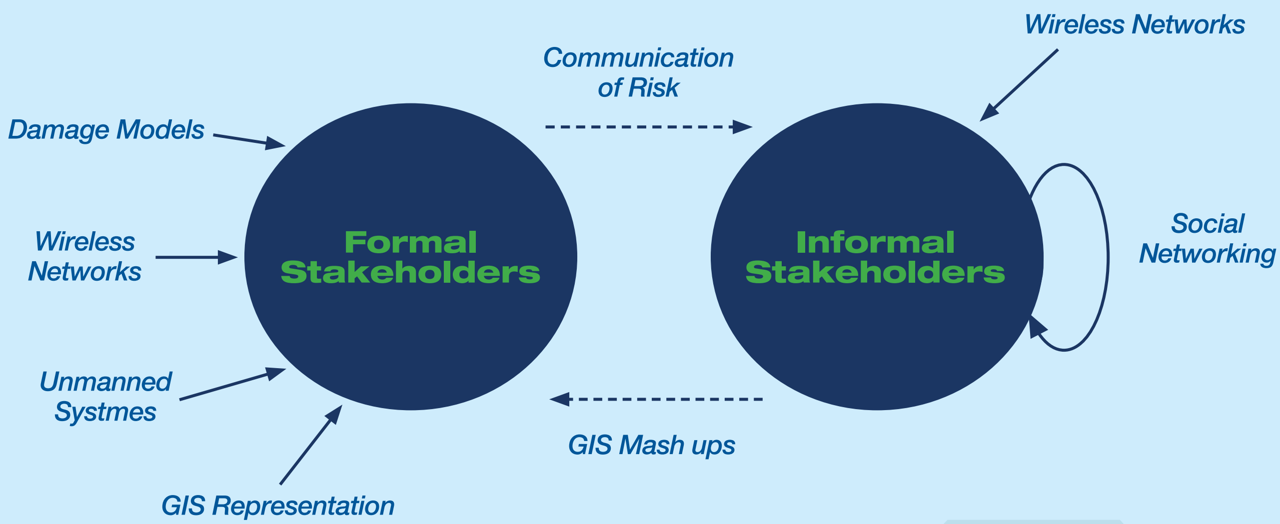

Figure 4.
Data sources for stakeholders.

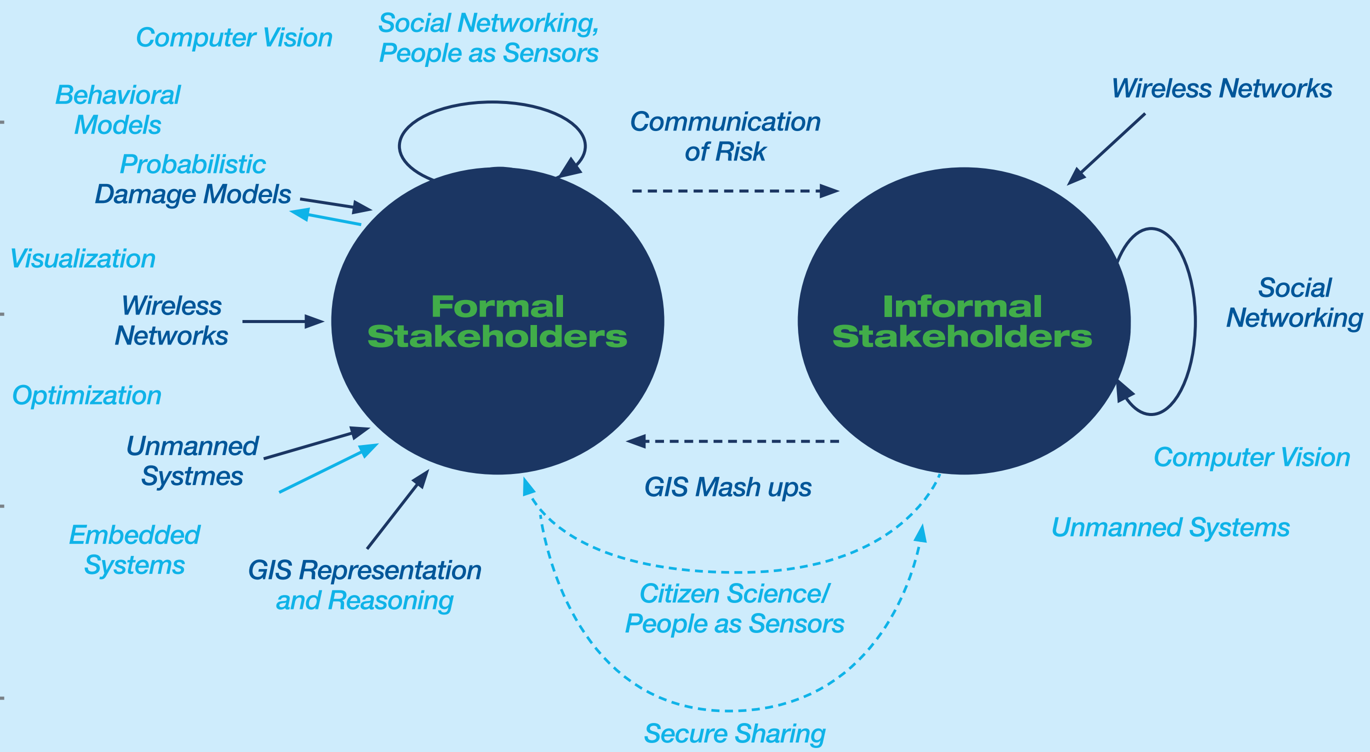

Figure 5.
Advances in data available to stakeholders.

3.3 Benefits Economy and Research

CRICIS research has four major benefits to society. First, at the national level, it can provide the technology needed to save lives and preserve continuity of government, while accelerating economic recovery. Second, CRICIS goes beyond what government, either DHS or DoD, and industry can do alone and it addresses the inherent national security considerations that historically prevent inter-agency dissemination of data, expertise, and advances. Third, the focus on researcher-stakeholder partnerships creates the potential for new collaborations and new ways of thinking. Fourth, CRICIS offers the prospect of extending new discoveries in socio-technical systems for decision-making to other domains.

CRICIS has the right timing and scope to leverage the defense investment in information technologies in order to revitalizing public safety and emergency preparedness while creating new jobs for returning veterans, fostering new companies, and sustaining competitiveness in existing industries. The Department of Defense has created technologies such as unmanned systems, sensor networks, visualization tools, and situation awareness displays that have a high potential benefit for disaster response. However, most of the technologies are still in their infancy in terms of innovation and still need improvement in order to be adopted by the Public Sector. Additionally, public safety exhibits entirely different characteristics with respect to training, frequency of use, budget constraints, and socio-organizational inertia, and these have historically hampered the transfer of technology from defense to homeland security applications. CRICIS research will bridge this gap between core technological capabilities and the very different social and organizational ecologies.

CRICIS will contribute to understanding decision-making in socio-technical systems, which is expected to reverberate to other applications besides disasters. The new territory of Computing for Disaster brings national priority challenges that require researchers to broaden their perspectives beyond familiar topics to unchartered territories where new skills are required. They will have to shift their focus from working with small amounts of data or narrow capabilities to building models, testing hypotheses, identifying patterns, developing simulations, and conducting longitudinal and cross-disciplinary studies at enormous scales. The opportunity is large, but these new tasks require fresh thinking, novel research methods, and interdisciplinary collaborations. Thus the broad perspective endowed by CRICIS is expected to lead to new solutions for complex socio-technical systems in general. Advances in decision making for disasters might create new directions in healthcare during pandemics or in coordinating and optimizing critical infrastructure systems, such as power grids and transportation.

4 A ROADMAP AND PORTFOLIO FOR FOSTERING CRICIS

As seen in the previous sections, CRICIS requires a different type of research community, one in which researchers are engaged with stakeholders as much as they are with one another as they address extreme scales, complexity, and human-system integration. This focus on meaningful connections with stakeholders is shared with the bioinformatics and health IT communities which have coalesced and progressed due to programmatic shepherding. However, CRICIS is not just adding computing to a different discipline: working with computing at the unique interdependencies of extreme scales and complexity functioning within human-computing systems will challenge and grow basic research in computing that will have benefits beyond disaster management.

Fostering such research requires a sustained, multi-phase effort to establish a roadmap and a balanced portfolio of research that will attract and introduce researchers, build partnerships between researchers and stakeholders, create testbeds and evaluation opportunities, and facilitate technology transfer. While history shows innovations rarely follow roadmaps [17], roadmaps can provide a common ground for stakeholders and researchers to envision desirable computing capabilities. Such roadmaps are valuable because they help lower the bar of entry in computing for disasters research. They also facilitate adoption of capabilities by agencies because the agencies can anticipate likely innovations and begin to plan for insertion, evaluation, and training. The National Academies study “Improving Disaster Management: The of IT in Mitigating, Preparedness, Response and Recovery” [18] recommended a living roadmap one that evolved continuously.

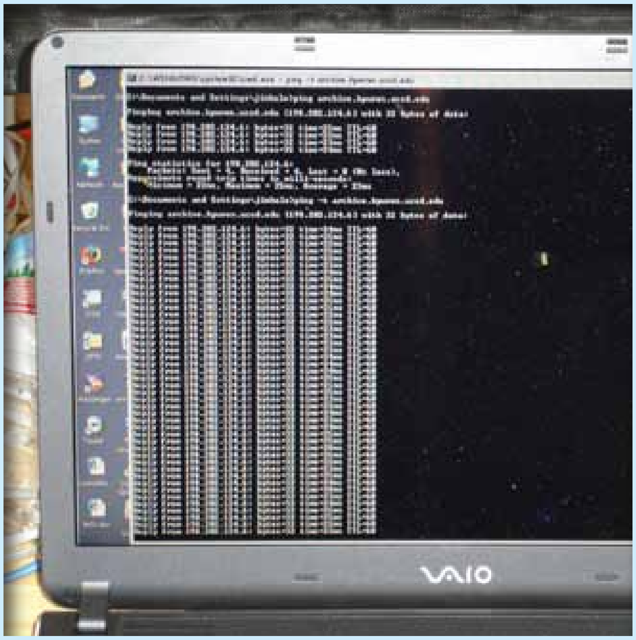

4.1 A Living "Computing for Disasters" Roadmap

The living research roadmap must reflect a consensus of stakeholders and researchers. One mechanism would be to i) hold a domain knowledge elucidation event to create a baseline, ii) advertise and refine the roadmap through technology-centric workshops or sessions at relevant conferences, and iii) host an annual workshop at NSF with stakeholders in conjunction with a demonstration similar to NSF Hazards Research Showcase held September 6-7, 2011, at the Hart Senate Office Building (aka "Disasters on the Hill"). This roadmapping workshop could serve to encourage agency stakeholders to contribute to funding the research portfolio.

A domain knowledge elucidation workshop would convene formal and informal stakeholders with researchers to define the state of the practice, identify and codify information-intensive disaster scenarios, capture training and education goals and constraints, propose benchmarks, and to the degree possible, specify metrics, methods, and focus areas. The inclusion of multiple government agencies, such as NSF, NIST, NASA, and NOAA, would be desirable. Formal stakeholders would consist of federal and state government agencies (e.g., FEMA, USGS, National Geospatial Agency, Department of Transportation, Army Corps of Engineers, Centers for Disease Control, National Guard, etc.), and non-governmental organizations and industry associations with a formal role under the National Response Framework (e.g., American Red Cross, CHEMTREC, etc.). Informal stakeholders would include response-related industries (e.g., CrisisMappers, insurance companies, etc.) and industries with key technologies (e.g., telecommunications, networks, robotics, simulation, etc.).

Workshops or sessions at relevant disaster-oriented computing conferences would help introduce researchers to the systems issues and opportunities and reinforce the distinct style of collaborative research. Any new NSF program in CRICIS should make a strong effort to work with these existing communities and benefit from their skills but help to expand their horizons and opportunities. The sessions would be expected to generate new ideas and possibilities as well. An effort should be made to cover the major conferences within two years, and it is expected that these sessions would become integral to many of the conferences. In addition, sessions at more mainstream conferences (e.g., ACM, CHI, IEEE, FedCSIS, etc.) may be appropriate; however, targeting researchers already pre-disposed to work in Public Safety may be more immediately productive. A partial list of conferences to be targeted includes:

- Conference on Earthquake Engineering
- Crisis, Tragedy, and Recovery Network associated with the ACM/IEEE Joint Conference on Digital Libraries
- IEEE Conference on Technologies for Homeland Security (IEEE HST)
- IEEE International Symposium on Safety, Security and Rescue Robots (IEEE SSRR)
- Information Systems for Crisis Response and Management (ISCRAM)

A two-day annual workshop/"Disasters on the Hill" event would reinforce the commitment of NSF to the formation of this vital research community, increase the visibility of the community to the larger research community and Congress, allow stakeholders to meet and interact with the technology and researchers, and create new stakeholder-research partnerships. The first day could consist of a "Disasters on the Hill" like event to introduce stakeholders and researchers, with the second day devoted to feedback from the stakeholders and collective envisioning.

4.2 Funded Research Portfolio

In order to effectively advance the state of the art and practice in disaster management, the envisioned funded research program should be structured in parallel tracks that reinforce each other. The envisioned time frame for the program overall is 15 years, with rapid response grants, traditional PI grants of 2 to 5 years, and centers having a minimum of 5 year duration, with the option to renew for another 5 years. Projects should be cognizant that stakeholders will not be able to make initial sizable infrastructure investments or change legally mandated procedures in order to accept new technology, instead the results of computing research must transform the system from within by giving stakeholders the tools that allow these experts to do a better job.

1. **Rapid response grants.** It is essential that researchers be able to document and immediately engage in disasters to capture ephemeral data. While the NSF RAPID program has a strong track record of supporting disaster research as noted in a recent workshop synthesizing lessons learned from RAPIDs related to the New Zealand earthquake and the Tohoku Japan earthquake and tsunami [19], there are often delays of weeks and even months. This prevents comprehensive data collection and effective interventions in the critical first two weeks after a major disaster. Therefore a pool of funds and protocol for immediate distribution should be available to researchers who maintain readiness to act and collect data on behalf of the larger research community.

2. **Traditional PI grants.** Three sizes of grants provide a suitable investment and entry point for individual or groups of researchers as well as more mature teams of investigators or teams conducting extensive fieldwork:
   - ***Seedling Projects***, in which a small group of researchers partners with a stakeholder to address a single topic within computing for

disasters. Seedling projects would be of short duration, i.e., no more than two years, and would be expected to lead to a Medium or Large Project. These projects allow researchers to make connections with response professionals and identify their possible computational contributions.

- ***Medium Projects***, in which a multi-disciplinary group of researchers partner with one or more stakeholders to address one or more topics and to evaluate the results within a systems context in appropriate physical exercises, testing of social science theories and hypotheses against huge archives of data and information, or computer simulations. The results should include some measure of human-computer interaction. These projects would last up to 5 years in duration, as field work with human subjects is often time-consuming and difficult to schedule.
- ***Large Projects*** are more ambitious in scope or in implementation and testing than Medium projects. Large projects can range from the establishment of a testbed or simulation to projects that capture computing at scale. These projects are expected to have multiple stakeholders as partners; for example, a project on computing for nuclear disasters would be expected to have representatives from the nuclear power industry, the response community, and nuclear regulatory agencies. The Large projects would last up to 5 years, as field work with human subjects is often time-consuming and difficult to schedule.

5. **Centers.** Following [18], multiple, large multidisciplinary centers of universities, response organizations (local or federal), and other relevant participants should be established to focus on promising technologies that would advance disaster resilience or response. The centers would couple research with practice by developing testbeds and exercises in order to rapidly evaluate concepts and elicit stakeholder requirements for emerging technologies. It is expected that at least one center will be capable of large-scale physical simulation in order to understand issues that only emerge at scale. Centers would be expected to address simulation, both for evaluation and for training. These grants would support the development of new curriculum and build new degree programs, professional certification, and minors in disaster management. One approach is to consider centers organized around regional resources and disaster professionals so as to facilitate the close partnerships and rapid prototyping of ideas needed for successful systems research. Center grants would be expected to persist for 5 years, with renewal for another 5 years. The establishment of centers would be spread out over the first five years in order to reap the benefits of growing the community.

To maximize the impact of the research conducted in this program, both through individual grants and through the multi-disciplinary institutes, the participants will be encouraged to utilize the testbeds and metrics to assess progress and convey to the stakeholder (e.g., eventual adopter communities) the value of their results. Regardless of size, projects should successfully integrate basic and applied, theoretical and practical, mission-driven and curiosity-driven research to advance both CRICIS and computing in general.

ACKNOWLEDGMENTS

The Critical Real-time Computing and Information Systems (CRICIS) Workshop was jointly sponsored by the National Science Foundation and the Computing Community Consortium (CCC; **http://cra.org/ccc**). We especially thank ***Dr. Erwin Gianchandani*** for his guidance and valuable feedback.

This report was produced by the steering committee for the workshop, comprising:

- ***Robin Murphy***, *co-chair, Texas A&M University*
- ***Trevor Darrell***, *co-chair, University of California, Berkeley*
- ***Hans-Werner Braun***, *University of California, San Diego*
- ***K. Mani Chandy***, *California Institute of Technology*
- ***Ed Fox***, *Virginia Polytechnic Institute*
- ***Susan Metros***, *University of Southern California*
- ***Elena Messina***, *National Institute for Standards and Technology*
- ***Paul Oh***, *Drexel University*
- ***Leysia Palen***, *University of Colorado, Boulder*
- ***Ramesh Rao***, *University of California, San Diego*
- ***Shashi Shekhar***, *University of Minnesota*
- ***Ben Shneiderman***, *University of Maryland*

It was informed and reviewed by the workshop participants:

- ***Anish Arora***, *Ohio State University*
- ***David Boyd***, *Department of Homeland Security*
- ***Tim Brown***, *University of Colorado, Boulder*
- ***Amy Bruckerman***, *Georgia Institute of Technology*
- ***James Caverlee***, *Texas A&M University*
- ***Mani Chandy***, *California Institute of Technology*
- ***ChienHung Chen***, *Georgia Institute of Technology*
- ***Sarita Chung***, *Harvard University*
- ***Ram Dantu***, *University of North Texas*
- ***Trevor Darrell***, *University of California, Berkeley*
- ***Daniel Fagbuyi***, *George Washington University*
- ***Scott Fahlman***, *Carnegie Mellon University*

- ***Ed Fox***, *Virginia Polytechnic Institute*
- ***Peter Freeman***, *Georgia Institute of Technology*
- ***Arif Ghafoor***, *Purdue University*
- ***Marc Hafner***, *California Fire*
- ***Milton Halem***, *University of Maryland*
- ***Jose Holguin-Veras***, *Rensselaer Polytechnic University*
- ***Andrea Kavanaugh***, *Virginia Polytechnic Institute*
- ***Eva Lee***, *Georgia Institute of Technology*
- ***Alan MacEachren***, *Pennsylvania State University*
- ***Frank McKenna***, *University of California, Berkeley*
- ***Sharad Mehrota***, *University of California, Irvine*
- ***Elena Messina***, *National Institute of Standards and Technology*
- ***Susan Metros***, *University of Southern California*
- ***David Millard***, *Georgia Institute of Technology*
- ***Matt Minson***, *Texas Engineering Extension Service*
- ***Alexander Morrison***, *Ohio State University*
- ***Robin Murphy***, *Texas A&M University*
- ***Paul Oh***, *Drexel University*
- ***Leysia Palen***, *Colorado State University*
- ***Rebecca Pierce***, *University of Pennsylvania*
- ***David Pietrocola***, *National Science Foundation*
- ***Ramesh Rao***, *California Institute of Technology*
- ***Eric Rasmussen***, *San Diego State University*
- ***Kate Saenko***, *University of California, Berkeley*
- ***Shashi Shekhar***, *University of Minnesota*
- ***Ben Schneiderman***, *University of Maryland*
- ***Phil Smith***, *Ohio State University*
- ***Richard Smith***, *Texas A&M, Corpus Christi*
- ***Satoshi Tadokoro***, *Tohoku University*
- ***R. Uthurusamy***, *Association for Computing Machinery*
- ***Ganesh Venayagamoorthy***, *Clemson University*
- ***Ning Zhang***, *University of California, Berkeley*

REFERENCES

1. Blanchard, B.W. 2007. *Guide To Emergency Management And Related Terms, Definitions, Concepts, Acronyms, Organizations, Programs, Guidance, Executive Orders & Legislation*. FEMA Emergency Management Institute. Retrieved from http://training.fema.gov/EMIWeb/edu/termdef.asp.

2. President's Council of Advisors on Science and Technology (PCAST). 2007. *NIT for Resilient Physical Systems*. Office of Science and Technology Policy, Executive Office of the President, The White House.

3. Computing Research Association. 2003. *Grand Research Challenges in Information Systems*.

4. President's Council of Advisors on Science and Technology (PCAST). 2003. *The Science and Technology of Combating Terrorism*, Office of Science and Technology Policy, Executive Office of the President, The White House.

5. Committee on Environment and Natural Resources. 2008. *Grand Challenges for Disaster Reduction*. National Science and Technology Council, Executive Office of the President, The White House.

6. Sircar, J., et al. 2009. "*Catastrophe Bonds for Transportation Assets Feasibility Analysis for Bridges,*" Transportation Research Record: Journal of the Transportation Research Board, 2115: 12-19, The National Academies.

7. Mason, B., ed. 2007. *Citizen Engagement in Emergency Planning for a Flu Pandemic: A Summary of the October 23, 2006 Workshop of the Disasters Roundtable*. National Research Council. The National Academies.

8. Mason, B., ed. 2006. *Community Disaster Resilience: A Summary of the March 20, 2006 Workshop of the Disasters Roundtable*. National Research Council. The National Academies.

9. Kershaw, P.J. and B. Mason. 2006. *The Indian Ocean Tsunami Disaster: Implications for U.S. and Global Disaster Reduction and Preparedness -- Summary of the June 21, 2005 Workshop of the Disasters Roundtable*. National Research Council. The National Academies.

10. Committee on the Future of Emergency Care in the United States Health System. 2007. *Emergency Medical Services: At the Crossroads*. Board on Health Care Services (HCS) and the Institute of Medicine (IOM). The National Academies.

11. Committee on the Effective Use of Data Methodologies and Technologies to Estimate Subnational Populations at Risk. 2007. *Tools and Methods for Estimating Populations at Risk from Natural Disasters and Complex Humanitarian Crises*. Board on Earth Sciences and Resources. The National Academies.

12. Committee on Planning for Catastrophe: A Blueprint for Improving Geospatial Data Tools and Infrastructure. 2007. *Successful Response Starts with a Map: Improving Geospatial Support for Disaster Management*. Board on Earth Sciences and Resources and Earth and Life Studies. The National Academies.

13. Committee on Disaster Research in the Social Sciences: Future Challenges and Opportunities. 2006. *Facing Hazards and Disasters: Understanding Human Dimensions*. Board on Earth and Life Studies. The National Academies.

14. McClean, D., ed. 2010. *World Disasters Report 2010: Focus on urban risk. International Federation of Red Cross and Red Crescent Societies*.

15. Adam, N. 2009. *Science & Technology Directorate U.S. Department of Homeland Security Workshop on Emergency Management: Incident, Resource, and Supply Chain Management*. U.S. Department of Homeland Security.

16. Andersson, K.P. and E. Ostrom. 2008. "*Analyzing decentralized resource regimes from a polycentric perspective.*" Policy Science 41: 71-93.

17. Rogers, E.M. 2003. *Diffusion of Innovations*, 5th ed. Free Press.

18. Rao, R.R., J. Eisenberg, and T. Schmitt, eds. 2007. *Improving Disaster Management: The Role of IT in Mitigation, Preparedness, Response, and Recovery*. Computer Science and Telecommunications Board and Engineering and Physical Sciences The National Academies.

19. Eguchi, R., et al., eds. 2012. *2010 Canterbury and 2011 Christchurch New Zealand Earthquakes and the 2011 Tohoku Japan Earthquake: Emerging Research Needs and Opportunities*. A report from a Workshop held February 9 and 10, 2012. Earthquake Engineering Research Institute.

20. Subcommittee on Disaster Reduction. U.S. Department of Homeland Security. Retrieved from http://www.sdr.gov.

21. U.S. Department of Homeland Security. 2008. *Overview: ESF and Support Annexes Coordinating Federal Assistance in Support of the National Response Framework*.

22. Committee on Science and Technology for Countering Terrorism. 2002. *Making the Nation Safer: The Role of Science and Technology in Countering Terrorism*. National Research Council. The National Academies.

APPENDIX A: KEY STAKEHOLDERS AND FUNCTIONS

Critical Real-time Computing and Information Systems (CRICIS) requires efficient access to diverse information available from formal and informal stakeholders as well as a solid understanding of science and technology, rapid implementation of research information into disaster reduction programs and applications. A starting point for understanding the diversity of stakeholders is to consider the set of stakeholders in Table 1 identified by The Subcommittee on Disaster Reduction (SDR). While the list in Table 1 includes academics via funding from the NSF, it lacks direct engagement from industry partners, Citizens, and other informal stakeholders. Incorporating an industry liaison to various companies (e.g., CrisisMappers, Google, Twitter, etc.) could facilitate further cooperation with the tools that are frequently utilized during disasters and actively support emergency services.

The Subcommittee on Disaster Reduction (SDR) is an element of the White House Office of Science and Technology and Policy's National Science and Technology Council. The Council facilitates national strategies for reducing disaster risks and losses that are based on effective use of science and technology [20]. Table 1 shows the members of this subcommittee, representing a broad collection of federal agencies responsible for various aspects of disaster management.

Many of the key stakeholder agencies are also a component of the Department of Homeland Security/FEMA Emergency Support Function system [21] listed in Table 2 with the assigned support roles for emergency services. The Emergency Support Function system organizes the contribution of each stakeholder in reducing the harm disasters cause to society, the economy, and the lives of individuals and communities by function. Each of the 15 functions requires disaster managers to consider a different facet of the situation (e.g., transportation, communication, public works, public health, etc.) and to act within that scope to reduce uncertainty, to calculate and compare costs and benefits, and to manage resources, often on a much larger scale and at a much faster pace than are supported by methods and means for solving ordinary problems.

CRICIS will provide capabilities that can help the decision-makers for each of these functions grasp the dynamic realities of a disaster more clearly and help them formulate better decisions more quickly; it can also help keep better track of the myriad details across in all phases of disaster management [18]. An important element of a research program will be designing protocols to provide emergency-response organizations with authoritative knowledge and support regarding adoption of new technologies. In addition, designated emergency response agencies should use existing technology to achieve short-term improvements in the telecommunications and computing infrastructure for first responders [22].

FEMA	Dept. of Transportation	NOAA
US-DoD: Networks and Information	Economic Development Administration	US Public Health Services
Bureau of Land Management	EPA	US-AID
Centers for Disease Control	Federal Energy Regulatory Commission	US Army Corps of Engineers
Dept. of Homeland Security	NASA	Dept. of Agriculture
Dept. Housing and Urban Development	National Geospatial-Intelligence Agency	US Forest Service
Dept. of Energy	National Guard Bureau	US Geological Survey
Dept. of State	National Institutes of Health	NSF
Dept. of the Interior	NIST	American Red Cross

Table 1.
Members of the Subcommittee on Disaster Reduction (SDR).

Emergency Support Functions

Agency	Transportation	Communications	Public Works and Engineering	Firefighting	Emergency Management	Mass Care, Emergency Assistance, Housing, and Human Services	Logistics Management and Resource Support	Public Health and Medical Services	Search and Rescue	Oil and Hazardous Materials Response	Agriculture and Natural Resources	Energy	Public Safety and Security	Long-Term Community Recovery	External Affairs
USDA			◆		◆	◆	◆	◆		◆	● ✚ ◆	◆		✚	◆
USDA/FS	◆	◆	◆	● ✚		◆	◆	◆	◆	◆			◆		
DOC	◆	◆	◆	◆	◆		◆	◆	◆	◆	◆	◆	◆	◆	◆
DOD	◆	◆	◆	◆	◆	◆	◆	◆	✚	◆	◆	◆	◆	◆	◆
DOD/USACE	◆		● ✚	◆		◆	◆	◆	◆	◆	◆	◆	◆	◆	
ED					◆										◆
DOE	◆		◆		◆		◆	◆		◆	◆	● ✚	◆	◆	◆
HHS			◆		◆	◆	◆	● ✚	◆	◆	◆			◆	◆
DHS	◆	◆	◆		◆		◆	◆	◆	◆	◆	◆	◆	✚	●
DHS/FEMA	◆	✚	✚	◆	● ✚	● ✚ ◆	● ✚	◆	● ✚	◆	◆			● ✚	✚
DHS/NCS		● ✚					◆					◆			
DHS/USCG	◆		◆	◆				◆	✚	✚			◆		
HUD					◆	◆								✚	◆
DOI	◆	◆	◆	◆	◆	◆	◆	◆	✚	◆	● ✚	◆	◆	◆	◆
DOJ	◆				◆	◆		◆	◆	◆	◆		● ✚		◆
DOL					◆	◆	◆	◆	◆	◆	◆	◆		◆	◆
DOS	◆		◆	◆	◆			◆		◆	◆	◆			◆
DOT	● ✚		◆		◆	◆	◆	◆		◆		◆		◆	◆
TREAS					◆	◆							◆	◆	◆
VA			◆		◆	◆	◆	◆					◆		◆
EPA			◆	◆	◆			◆		● ✚	◆	◆	◆	◆	
FCC		◆			◆										◆
GSA	◆	◆	◆		◆	◆	● ✚	◆		◆	◆				◆
NASA					◆		◆		◆				◆		◆
NRC			◆		◆					◆		◆			◆
OPM					◆		◆								◆
SBA					◆	◆								✚	◆
SSA						◆							◆		◆
TVA			◆		◆							◆			◆
USAID								◆	◆						◆
USPS	◆				◆	◆		◆			◆		◆		◆
ACHP											◆				
ARC			◆		◆	◆		◆			◆			◆	
CNCS			◆			◆								◆	
DRA														◆	
HENTF											◆				
NARA											◆				
NVOAD														◆	

●=ESF coordinator ✚=Primary agency ◆=Support Agency

Table 2.
National Response Framework- Emergency Support Functions (ESF).

THE COMPUTING COMMUNITY CONSORTIUM

Established in 2006 through a Cooperative Agreement between the National Science Foundation (NSF) and the Computing Research Association (CRA), the CCC serves as a catalyst and enabler for the computing research community. Its goals are to unite the community to contribute to shaping the future of the field; provide leadership for the community, encouraging revolutionary, high-impact research; encourage the alignment of computing research with pressing national priorities and national challenges (many of which cross disciplines); give voice to the community, communicating to a broad audience the many ways in which advances in computing will create a brighter future; and grow new leaders for the computing research community.

www.cra.org/ccc/