Science, Engineering, and Education of Sustainability: The Role of Information Sciences and Engineering

Computing Community Consortium[1,2]
Version 18: May 10, 2011

Fundamental research in computer and information science and engineering (CISE) has led to technology that has transformed our modern world. Our ability to communicate with others, to access information, to benefit from modern medical care, to protect our nation's security, to improve and enhance classroom education, etc., stems from investments in CISE research over the past 60 years. In much the same way, CISE research has the potential to put the world on a path toward *sustainability*—allowing us to manage the use of natural resources and our impact on the environment such that future generations can enjoy a high quality of life. CISE technology can lead to energy, transportation, and water systems that achieve much higher degrees of efficiency by continuously engaging in a cycle of sensing environmental conditions, learning trends and patterns, and optimizing and controlling system operation. It can lead to systems that understand the needs and preferences of people and that support and encourage behavioral changes, both individually and collectively. It can enable scientists, policy makers, and citizens to better predict and understand the consequences of our actions through accurate and transparent modeling tools. Achieving these capabilities will require fundamentally new tools and approaches in such areas as large-scale data management and analysis, modeling and simulation, intelligent optimization and control, cyber-physical systems, human-centered and social computing, and systems engineering. It will also require advances in privacy preservation and cybersecurity to ensure that these sensor-rich, highly adaptive systems cannot be misdirected by malicious intruders or become unwitting conduits of sensitive information about individuals or critical infrastructure.

The CISE community has long engaged in research efforts in sustainability that are directed *inward*, focusing on the resource consumption and environmental impact of our own technology. Examples include the study of low-power electronics, power conservation in data centers, and improved recyclability of electronic devices. The information technology (IT) industry is already reducing the environmental impact of its own technology, and CISE research can support these efforts. We describe research opportunities in this area under the heading of *Green IT* (see Section 3.9). It is important to note, however, that even with the ever-growing number of large-scale data centers in the U.S., their total power consumption constitutes less than 2 percent of our nation's electricity usage. To truly put society on a path to sustainability, the CISE community must do more than just minimize its own environmental footprint. CISE research can have this much greater impact by looking *outward,* using CISE to fundamentally change how energy is generated, distributed, and consumed; how people and goods are transported; and how scientists

---

[1] Contact: Dr. Erwin Gianchandani, Director, Computing Community Consortium (202-266-2936; erwin@cra.org). For the most recent version of this report, as well as related reports, visit http://cra.org/ccc/initiatives.
[2] This white paper was co-authored by Randal Bryant (Carnegie Mellon University), Douglas Fisher (Vanderbilt University), Erwin Gianchandani (CCC), Carla Gomes (Cornell University), William Rouse (Georgia Institute of Technology), Prashant Shenoy (University of Massachusetts-Amherst), Robert Sproull (Oracle), and David Waltz (Columbia University). These individuals, along with Krishna Kant (National Science Foundation and George Mason University), were also co-organizers of a workshop on sustainability and IT; see details in [3] below.

analyze and manage environmental resources, particularly in the face of a changing climate. These outward-directed aspects of sustainability have received far less attention in the CISE research community, and hence they will be the major focus of this report.

Indeed, in recent years, several interdisciplinary teams led by CISE researchers have begun pursuing the outward direction, successfully exploring sustainability challenges in energy, transportation, and climate. For example, the National Science Foundation's (NSF) CISE Directorate has funded two large-scale, multi-investigator, multi-year "Expeditions" projects in the past three years that are pursuing cutting-edge research (1) across conservation and biodiversity, balancing socio-economic demands and the environment, and renewable energy (National Science Foundation, 2008); and (2) around data-driven approaches for understanding climate change (National Science Foundation, 2010).

These sustainability-focused Expeditions illustrate the desire by CISE researchers—and the need—to devise technology that will have impact on society and the environment. As we will describe in this report, **adequately addressing this set of problems requires long-term collaborations between CISE researchers and other scientists and engineers concerned with sustainability,** in such areas as electrical engineering, transportation systems, biology, environmental science, and social and behavioral science. **CISE innovations are critical for addressing core problems, and so CISE researchers must be central to these interdisciplinary teams**. At the same time, CISE researchers must also learn from their academic and industrial partners about the real-world problems faced by those seeking to achieve higher degrees of sustainability, such as evolving an aging electric infrastructure to a truly smart grid; transforming a huge, entrenched transportation system; or mitigating and adapting to climate change and its impact on the biosphere. **Most importantly, a sustained, multi-agency Federal investment in basic research and education, as well as cooperation by key industrial players, is critical in order to realize a sustainable vision for the future. Without such a long-term, integrated approach, the CISE research community is likely to fall far short of the profound capabilities that it can provide.**

In this report[3], we outline the vast challenges posed by sustainability, describe the critical role CISE research must play to adequately address these challenges, and emphasize core modalities for facilitating research advances in the short and long term. We conclude with a set of recommendations for achieving the vision for a sustainable future.

1. Defining sustainability

Sustainability **can be broadly defined as "meeting the needs of present and future generations while substantially reducing poverty and conserving the planet's life support**

[3] This report is the result of a workshop on the Role of Information Sciences and Engineering in Sustainability (RISES) co-funded by the NSF's CISE Directorate and the CCC. The workshop—comprising a diverse group of 60 leading computer scientists, systems engineers, power engineers, biologists, environmental scientists, etc.—took place in Washington, DC, on February 3-4, 2011. The goals were to understand new research opportunities in the information sciences and engineering that address sustainability objectives, particularly areas where gaps in current research lie or where transformational impact stands ready to be made. For more information about the workshop, including the complete agenda, videos of the plenaries, and short white papers submitted by some participants, etc., see http://cra.org/ccc/seesit.

systems" (*Proceedings of the National Academy of Sciences*, 2011). It requires reconciliation of environmental, social, and economic demands (United Nations General Assembly, 2005). Sustainability issues span the natural and built environments, including energy, transportation, climate, and biodiversity, and the science of sustainability thus comprises the fundamental character of interactions among humans, their technologies, and the environment, as well as the knowledge to advance sustainability goals relevant to water, food, energy, health, habitation, mobility, and ecosystems.

In the materially developed world, an often-implicit objective for sustainability is to preserve current living standards. People in the developing world, on the other hand, feel that they, too, should be able to enjoy the comfort, security, and longevity of people in the first world. Unfortunately, the world cannot sustain the rate of resource usage and environmental degradation that would be incurred by scaling our current means for supporting a middle-class lifestyle to 7 billion people, especially due to the high reliance on fossil fuels. To achieve a high quality of life at a global scale, we must therefore ensure that the world's energy needs can be satisfied by renewable resources, that safe and sufficient land and water supplies are available to meet our personal and agricultural needs, that houses and buildings can provide comfort and shelter to their occupants, and that we can transport people and goods efficiently.

Key goals for sustainability thus include (among others): decreasing overall energy consumption while also increasing use of renewable energy sources; improving transportation to minimize energy usage and environmental impact; reducing and adapting to climate change; conserving natural habitats; eliminating waste by designing products for full reuse; and moving toward zero loss of non-renewable resources. A sustainable ecosystem would be one in which decision making would not simply be based on current costs, but future costs and renewability as well.

Energy, transportation, and climate are key areas, but sustainability constitutes a broad-based discipline that is every bit as ubiquitous as computing is pervasive. It involves striking a well-reasoned balance between the near-term imperatives of energy and climate and the long-term requirements of a broader sustainable ecosystem. In addition to pursuing research on issues at the top of the national list of priorities, CISE researchers can find ample opportunities to have important impact across the entire spectrum of sustainability.

2. Key Mechanisms for Achieving Sustainability via CISE

Although computers will be used in almost every aspect of planning and operating systems designed to increase sustainability, CISE technology can have truly *transformative* impact along several different dimensions. The following is a sampling of the promising themes that CISE research can pursue in striving to go beyond the "routine use" of computers to one in which computing and information technology lead to fundamental improvements in efficiency and quality of life. Many of these ideas can be illustrated by their roles in developing more efficient and responsive transportation systems (see Box 1).

Systems should continuously adapt and optimize their operations based on measured data.

All systems—energy, transportation, water supply and treatment, etc.—can achieve fundamental improvements in efficiency while attaining higher reliability and better serving the needs of people by continuously monitoring and adapting their operations. Such systems must have rich networks of sensors to collect data on current states and environment conditions. Data collected from these sensors must be assimilated, managed, and analyzed to detect patterns and to generate predictive models. These models can then be used to optimize and tune system performance. Indeed, these different phases should operate at all times, so that the system continuously monitors, learns, and optimizes its behavior.

We can see some aspects of this idea being realized today, including the "smart grid" for electricity (see Box 3), as well as systems for optimizing urban traffic via centralized video monitoring and traffic light control, but this theme can be pushed well beyond the schemes under consideration today. Doing so will require major advances in sensor networks, data analytics, and intelligent control. In addition, advances in privacy preservation and cybersecurity will be necessary to prevent these feedback-rich systems from being hijacked by malicious agents or from becoming a source of unauthorized information about people, organizations, or governments.

Systems should be designed as large networks of loosely coupled agents to ensure robust operation.

Robust operation requires that a system have no single points of failure, that its control regimes remain stable, and that it handle all of the many contingencies that may arise. The only possible means of achieving robustness is by devising a system consisting of multiple interacting agents, with careful attention to fault tolerance. A distributed structure also helps limit the vulnerability of a system to attack, to natural disaster, or to unauthorized accessing of information. Achieving robust, reliable, and efficient operation with a network of millions of autonomous agents requires fundamental advances in distributed systems, auction technology, and systems engineering. Robustness is especially challenging for networked systems operating in the physical world, with unreliable sensors, anomalous events, and unpredictable human behavior.

People are important producers and consumers of information.

People increasingly produce information that systems—buildings, transportation networks, etc.—can use to learn and adapt to human needs and preferences. Some people are highly motivated to adopt sustainable lifestyles, buying hybrid cars, recycling whatever they can, and installing solar panels on their roofs. Such people may be willing to continuously monitor and optimize their own resource usage. Most people, however, do not want to be troubled with low-level monitoring and adaptation. Effective sensing and machine learning, along with expressive and convenient interfaces can make these tasks far more efficient and effective.

On the consumption side, achieving sustainability requires getting people to change their behaviors. Even highly motivated people have trouble understanding the potential impact of their choices. Advances in user interfaces and controls are needed to ease the use of complex adaptive systems. Social networking technology can be used to stimulate societal-scale efforts to adopt sustainable practices.

Scientists, policy makers, and citizens should have reliable and transparent mechanisms for predicting the future effects of choices we make today.

Scientists have long relied on large-scale simulations and other forecasting tools to model and predict phenomena ranging from weather and climate to the spread of contaminants in a water supply. Making full use of such tools requires better mechanisms for combining predictive modeling with sensor-based data collection and maximizing their accuracy and predictive power for a given computing budget. In addition, improved mechanisms for expressing and validating simulation models are required to enable scientists as well as non-specialists to fully trust the predicted behaviors.

Box 1: A sampling of CISE opportunities in transportation

Efficiency improvements in transportation contribute in several ways to sustainability by reducing energy used, greenhouse gases emitted, land seized for transportation uses, and raw materials used to build vehicles. Human acceptance and behavior will be key to achieving many of these savings, and there is always the risk that more efficient transportation will lead to greater use and override sustainability improvements (the rebound effect). Some approaches:

- Providing travelers with better, more accurate, and more connected information
- Making pricing and payment more convenient and efficient
- Control & scheduling of "personal rapid transit" vehicles
- Coordinated and collaborative transport across modes of traveling
- Real-time adjustment of schedules to minimize transfer waiting time
- Information about transfers to alternative transport modes
- Mining data to infer transportation behaviors
- Planning transportation systems and networks (modeling, simulation, optimization)
- Siting and operating sensors
- Using uncertain information from moving travelers and vehicles
- Wide use of vehicle-to-vehicle data exchange for everything from crash avoidance to routing
- Guidance of individual trips that promote "most efficient" path finding
- Integration and exchange of real-time data among disparate transportation systems and users
- Real-time monitoring of capacity (roads, parking) and consequent dynamic routing of trips
- Vehicle-to-command-center data flow to improve centralized response strategies
- Large-scale simulation to determine sound responses to disruptions and natural disasters
- Integration of freight transport with logistics and shippers/receivers
- Visualization techniques to effectively inform travelers
- Personal travel aids integrated into mobile devices
- Using people-movement data to inform transportation options
- Finding personalized travel options, such as sharing trips with neighbors
- Computer-mediated communications as a substitute for travel
- Integration of a heavily electric-dependent transportation system with the future energy “smart grid” (e.g., users need to know when and where to charge, etc.)

Adapted from workshop remarks of Michael Meyer. See also (Transportation for America, ITS America, Association for Commuter Transportation, and the University of Michigan’s SMART Initiative, 2010).

3. CISE research to further sustainability

We described in the preceding section guiding principles for achieving sustainability through CISE research. In this section, we expand upon that description by detailing how research in specific CISE areas can support fundamental improvements in sustainability (see Figure 1). The examples cited are not meant to be exhaustive, but only suggestive of how sustainability and CISE research can be coupled.

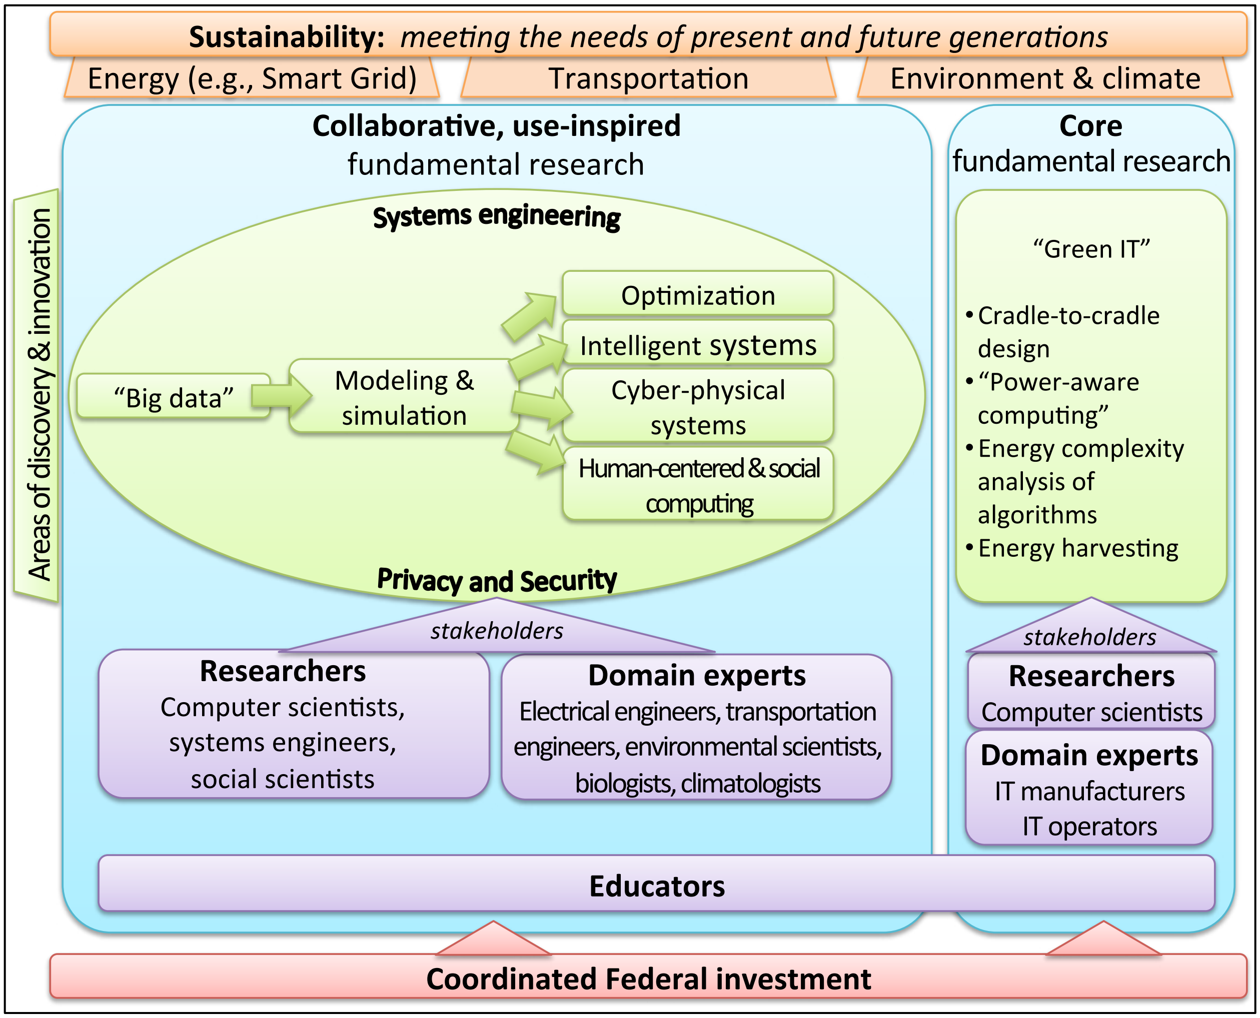

Figure 1. Toward an ecosystem for sustainable research.

3.1 "Big Data"

Like many other disciplines[4], the foremost challenge in sustainability science and engineering involves measuring, monitoring, and analyzing many aspects of the natural and built environments, yielding unprecedented amounts of data. We currently lack well-understood and thoroughly deployed techniques for addressing the particular needs associated with the data that are being amassed for sustainability research. For instance, large data sets already abound in

[4] See http://cra.org/ccc/dan.php for a series of white papers about data analytics in the context of core U.S. priorities.

areas of sustainability (e.g., climate modeling), but we still deal with these in an ad hoc manner.

Importantly, there are several unique aspects of sustainability—and, correspondingly, sustainability data—that make the "big data" problem especially challenging in this area. Sustainability data are particularly heterogeneous, as they can comprise components as varied as maps and overlays (e.g., from geographic information systems), sampled measurements (e.g., water pollution observations), imagery (e.g. satellite-based), extensive handwritten notes and comments (e.g., from geologists), data elements with potentially very large amounts of detailed measurements (e.g., DNA assays of individual organisms), and more. These data may be historical, such as handwritten observations, or real-time, such as social network-based traffic measurements (e.g., the recent work to locate earthquake epicenters based on cell phone traffic patterns[5]). They can be temporal as well as geospatial in nature, and the level of granularity, completeness, and/or confidence—particularly when considering historical data—can be highly variable from one collection or repository to another. For example, oftentimes researchers have to find ways to meaningfully analyze incomplete data sets by considering ranges of possible values and simulating the corresponding set of possible outcomes. Finally, related data may reside in several locations all over the world—owing to the varying locations where they must be captured—but will need to appear as a unified database to users at different locations.

Common infrastructures of techniques, software, and services (e.g., for climate modeling, tracking of pollutants, etc.) to support these data must therefore emerge, all the while enabling trust in the observations for data obtained from many different observers. Rather than have each research group create and support its own data repository, it may be most efficient to have one or more centers that combine sustainability research, research into the systems aspects of supporting sustainability research, and hosting particular databases. Research ideas and products from such centers could also benefit other researchers as well as industry in areas such as weather modeling, land-use planning and monitoring, and logistics planning for transportation and energy systems.

In addition to the basic need for facilities to store and access these data, several other approaches are necessary in order to exploit the sustainability data in meaningful ways that address the myriad challenges described above:

- Devise adequate metadata to describe data formats, meanings, and measurement conditions;
- Determine provenance of data and verify that it is genuine;
- Federate data collected from different sources, sensors, etc.;
- Curate data into trustworthy collections, fixing errors, and providing an audit trail, etc.
- Devise new visualization methods and tools;
- Adapt algorithms to work in parallel on large data sets (a la MapReduce, Hadoop, etc.);
- Improve data mining, to extract information and learn models from unstructured data;
- Manage large amounts of data, including by archiving ("sustaining sustainability data"); and
- Build applications and tools for real-time monitoring and control, combining streaming and historical data.

[5] http://irevolution.net/2010/10/26/earthquake-call-dynamics/

3.2 Modeling and simulation

Many of our attempts to chart the future and to design the built environment for sustainability depend on modeling and simulation of complex, large-scale problems. This will especially be the case as a result of the "big data" challenges described above; certain subsets of the unprecedented quantities of data about the natural and built environments will serve as the basis for constructing models, while other subsets will enable validation of the eventual simulation results. Climate modeling is of course a primary need, and a topic of considerable research and controversy within sustainability.

As models and simulations are used increasingly to inform public debate and policy choices, as well as to facilitate engineering decisions, it will be important to improve modeling and simulation technologies. Examples of computational problems that need to be addressed for sustainability science include:

- Simplifying scaling up simulations to the limit of what computing power that we have available (i.e., simulations can always be bigger and faster, but how do we maximize the information gained for a given computing budget?);
- Routinely conforming simulations to observed data, especially as sensors and sensor data mushroom; and
- Testing, debugging, and verifying models and simulations—since the point of simulation is to obtain a result we cannot verify independently, modeling and software errors are difficult to detect and rectify.

Moreover, as described previously, a key challenge in sustainability involves understanding the needs and preferences of people, as well as the knowledge and capabilities that support and encourage behavioral changes individually and collectively. Thus, unique to sustainability, we must:

- Find better ways to model and simulate human behavior and behavior changes, which will influence the effectiveness of many sustainability attempts; and
- Make models more transparent to permit domain experts—and even knowledgeable citizens—to inspect and debate properties of models without being overwhelmed by the complex software embedding the model in order to deploy a high-performance simulation.

Improvements in techniques for specifying and working with models and simulations could have enormous impact on the course of sustainability efforts around the world. Any proposal for a sustainability project will surely require modeling and simulation to predict outcomes. And since there are likely to be several competing proposals and debates about models and assumptions, a great deal of modeling and simulation work will precede decisions. Better ways to express models may have more impact than better ways to express computer programs!

3.3 Optimization

Efficiency drives many sustainability objectives, and optimization in turn drives efficiency. For instance, optimization is especially powerful for increasing energy efficiency. A ubiquitous example involves planning routes for delivering packages to households; optimizing these routes can save considerably on time and fuel. Very complex optimizers deal with even bigger logistics problems, such as scheduling trains, trucks, and airplanes to carry freight.

Though a very wide range of optimization applications is already in use today, large-scale efforts at sustainability may lead to even bigger optimization problems with requirements that have not yet been explored in great detail. For example, widespread ride-sharing or personalized traffic routing (that avoids using the same route to divert all traffic around a bottleneck) will require optimizations spanning tens or hundreds of thousands of vehicles. A useful system must also allow frequent entry of new vehicles and rapid determination of instructions to the drivers. To be acceptable to users, many more constraints may need to be considered, such as with which drivers a passenger is willing to ride, ranges of pickup times, and so on. Such user preferences may be particularly difficult to model, elicit, and incorporate into the optimization. Ultimately, it is likely that the optimization problems that arise as a part of sustainability efforts will require new algorithms or adaptations to address these new and specialized requirements.

Many projects to address sustainability will involve optimization in planning. For example, optimizing the location of wind turbine farms or individual towers within a farm requires tradeoffs of cost, energy output, and side effects (such as noise generated). Optimal siting can also be used to position sensors to obtain good coverage and low installation and operating costs. Similarly, there exists the complex problem of optimizing buildings (currently more than 70 percent of all electrical consumption), including systems that continuously learn from the actual behavior of the building and its occupants. These “smart buildings” must be even more capable if coupled to a “smart grid,” as described further below.

Often optimization will be part of a larger application-specific computing problem. Consider Box 2 below, in which we describe a planning problem with an objective function to optimize wildlife corridors to help ensure a sustainable wildlife population. Many of these planning problems cannot be tackled with existing algorithms, and instead require the formulation of new objective functions and new algorithms

In addition, the nature and scale of optimization required for many sustainability applications will favor decentralized schemes. For example, a household wishes to optimize its electricity usage within a smart energy grid that is optimizing energy delivery and consumption holistically. *Markets* and *auctions* are decentralized methods that can be run by combinations of humans and computers trying to optimize their outcomes. Such a distributed approach can avoid the bottlenecks, reliability risks, and privacy issues of a centralized optimizer. We currently face challenges in developing and operating distributed optimizations at very large scales.

Box 2: Modeling wildlife corridors for multiple species: The Steiner Multigraph Problem
Many endangered species are at risk because of habitat degradation and fragmentation due to deforestation, agriculture, urbanization, and other forms of human development. Habitat fragmentation creates small isolated communities that require efforts like translocation of similar species to revive the endangered population (U.S. Fish and Wildlife Service, 2006). One method to mitigate this problem is to set aside so-called wildlife corridors, or swaths of preserved land that connect the endangered species' remaining habitat. Wildlife corridors have been used successfully for various species like gray wolves (Shepherd and Whittington, 2006). Since funding is limited, deciding what land to buy is not an easy task. The number of land parcels under consideration can reach the thousands, and management agencies and conservation organizations are constrained both by their budget and a limited timeframe under which they must make their decisions. Motivated by this application, researchers have modeled the task of planning wildlife corridors as a Connection Subgraph problem, a variation of the classic combinatorial graph optimization problem of finding a minimum-cost Steiner tree (Conrad *et al.*, 2010; Dilkina *et al.*, 2010; Gomes *et al.*, 2008). A Steiner tree is a subgraph that connects a subset of the nodes called terminals. A feasible solution to the Connection Subgraph problem is a Steiner tree that is not necessarily of minimum cost but must cost less than some fixed budget, and an optimal solution has the maximum quality possible. In this work, the quality of a wildlife corridor is measured by the sum of the habitat suitability values of the purchased land parcels. While this model can design a wildlife corridor that caters to a single species' habitat requirements, there are oftentimes other species that need a corridor through the same landscape. As such, it would be even more economically efficient to design a single wildlife corridor that would serve all the species in the area. However, different species have different habitat preferences. More importantly, different landscape features (e.g. the steep side of a mountain) may act as barriers to some species but not to others. To address precisely this issue, some recent work generalizes the Steiner tree problem in a new way as the Steiner Multigraph Problem (SMP) (Lai *et al.*, 2011). In SMP, each species is represented by its own set of terminal nodes and the subset of the landscape that it can traverse. The objective of SMP is to minimize the cost of connecting each set of terminals while only using permissible nodes to connect them. Even though SMP was introduced to model corridors for multiple species, it is a source of interesting theoretical questions and is also relevant for applications outside computational sustainability such as wireless sensor networks and social networks. Finding a Steiner tree in a graph of potential sensor node locations corresponds to finding the cheapest backbone for multicast communication in the network. As an SMP instance, we can now model a network with different types of devices that may not all be compatible with the same communication protocol or may only have the required hardware for certain communication capabilities. This is but one example of **the importance of use-driven research, which we discuss in more length later: solutions motivated by a critically important problem require new science or engineering developments, which themselves can be applied to problem domains beyond those for which they were originally developed.**

3.4 Intelligent Systems

Intelligent systems promise to play a major role in enabling all aspects of sustainability. Approaches such as machine learning and intelligent agents exploit vast data sets to test hypotheses or uncover behaviors that serve sustainability. These results and insights can then drive intelligent control systems to optimize system behavior. Possible applications of intelligent systems to sustainability include:

- Making the smart grid even smarter (see Box 3 below), by learning to improve control to meet actual or localized behaviors of supply or consumption;
- Improving transportation systems by automatically monitoring traffic and then identifying sources of congestion and providing traffic-dependent driving directions;
- Monitoring and controlling building environment systems, based on learned patterns of usage and preferences;
- Optimizing the placement of wind farms based on factors such as weather patterns, geographic characteristics of energy usage and transmission, and land availability; and
- Detecting patterns of animal behavior or crop performance to detect ecosystem impairment or to inform climate change mitigation and adaptation efforts.

Intelligent systems are closely related to optimization; for example, the interplay between machine learning and intelligent systems constitutes an emerging area of work[6].

Some significant initiatives are already underway. For example, the Center for Computational Learning Systems at Columbia University (described in more detail in Box 5 below) is a largely industry-funded center applying these kinds of approaches to a variety of real-world needs. In general, however, there is a large gap between current practice and the most recent advances in intelligent systems.

One major research challenge is to ensure that systems employing intelligent control are truly robust. Machine learning and statistical control can improve the performance of a system operating under normal conditions, but other control mechanisms are required to handle anomalous and exceptional events. Given the major cost and safety issues involved in operating large-scale infrastructure, much greater attention must be paid to making intelligent systems completely trustworthy.

Box 3: Making the Smart Grid even smarter
Although the electric grid is considered one of the greatest engineering achievements of the early twentieth century, it is based on principles that yield fundamental inefficiencies and that prevent it from adapting to future energy sources and demands. Most significantly: 1) the utility companies attempt to meet the uncontrolled energy needs of their customers at all times, and 2) there is almost no ability to store energy as a way to achieve buffering between energy generation and consumption. As a result, utilities must build their facilities based on peak requirements, requiring significantly more infrastructure and more expensive energy sources

[6] http://jmlr.csail.mit.edu/papers/volume7/MLOPT-intro06a/MLOPT-intro06a.pdf

than if they could plan only for average demands.

The existing grid relies on power generation systems whose output can be readily increased or decreased according to demand. Renewable energy sources, such as wind and solar, do not fit into this scheme, and hence they are difficult to utilize beyond around 20 percent of the total energy supply[7]. The charging of electric and plug-in hybrid cars greatly increases the load on the electric system; charging the batteries in an electric car can draw as much power as that of a typical household. Imagine the demand that will be generated on a hot summer evening in the near future as people return from work, plug in their electric cars, turn on their air conditioners, do their laundry, and cook their meals, all within a short period of time.

The smart grid attempts to introduce a degree of *demand-response* into the electric system by providing mechanisms that will allow utility companies to lower energy consumption when demands cannot be met efficiently. Utility companies would like their customers to place their air conditioners under centralized control, so that they can be shut off when demands rise too high. Many utilities are installing "smart" meters, which allow them to continuously monitor power use. They will then introduce economic incentives, such as charging more for electricity during times of high overall demand, that will ideally motivate customers to time-shift some of their activities, e.g., laundry and electric car charging, to off-peak periods. However, notably these techniques are intrusive and rely heavily on customers carefully managing their energy usage.

Moving forward, a first major leap is to incorporate true *intelligence* into factories, offices, and homes. We can imagine that each building would have its own computerized energy resource manager that would understand the requirements and preferences of the occupants as well as current and near-term weather conditions (e.g., temperature, wind, sunlight, etc.), and continuously negotiate with the utilities to provide an optimal level of price, comfort, and convenience. This effort requires key computing advances:

- Research in the field of *cyber-physical systems* seeks means to implement highly adaptive, reliable, safe, and efficient sensor-based systems. For the case of the smart grid, these cyber-physical systems must scale to millions of loosely coupled autonomous agents, designed in such a way that their collective behavior achieves high reliability and efficiency.
- *Machine learning/data mining* can readily detect usage patterns and preferences automatically.
- *Agent-based systems* can negotiate complex contracts between consumers, suppliers, and transmission resources on a massive scale.
- The field of *human-computer interaction (HCI)* has studied how to create computer systems that provide appropriate levels of information and control to users in ways that are understandable and that minimize the chances of error.
- *Advanced optimization* can guide the adoption of renewable energy sources, such as wind and solar, based on projected macro-scale demand, grid capacity with anticipated upgrades, and consideration of the inherent intermittency of renewable power sources.

[7] http://www-fa.upc.es/personals/fluids/oriol/ale/eolss.pdf

Taking *a longer-term and more transformative view*, the Internet suggests alternative organizing principles for a twenty-first century decentralized peer-to-peer smart grid. The Internet succeeded by pushing intelligence to the network's edges while hiding the diversity of underlying technologies through well-defined interfaces. In other words, any device can be a source/sink of routable traffic, but the "intelligent" endpoints adapt their behavior to what the infrastructure can deliver in accordance with localized functions. Similarly, imagine an electric power system built on *packetized energy*: energy is stored where it is generated and conceptually "routed" to where it is needed. Achieving such a future presents even more fundamental challenges:

- Schemes are needed to efficiently store and release electric energy;
- The system must adapt to unexpected events;
- The system components must be able to cooperate with one another;
- The system must guarantee sufficient privacy;
- The system must be resilient to abuse or attack; and
- The system must learn and improve over time.

3.5 Cyber-physical systems

Cyber-physical systems (CPS) are those that use IT-based systems to monitor and control the physical world, including sensors and sensor networks, intelligent controllers, and robotics. We expect CPS will play a major role in all aspects of sustainability (see [8] for descriptions of future distributed energy and transportation systems, and how advances in CPS can contribute to these visions). Sensors and sensor networks are already used in limited ways for environmental monitoring (e.g., of species' populations) and for monitoring household and other building energy usage. In the long term, we imagine the entire world to be sensor-rich, monitoring our environment, transportation networks, and all forms of physical infrastructure. In addition, increasingly, the built environment is designed to be controllable (e.g., data-center cooling systems that use outside air when it is cool enough, vehicles whose engine controls are completely computerized, and room-by-room control of HVAC systems). Several important themes for advancing the role of CPS include:

- Having sensor networks use intelligence to automatically relocate, configure, tune, monitor, and repair themselves to make them easy to deploy and maintain;
- Using advanced machine learning and signal processing to harvest more information with fewer and lower cost sensors—including optimally placing sensors based on modeling and optimization algorithms;
- Making sensors easier to deploy by minimizing (or eliminating) battery power requirements (e.g., as described in Section 3.9, some low-power sensors already harvest solar or vibrational energy as their only power, and the European Commission has challenged the community to develop "zero-power" technology through energy harvesting at the nanoscale[9]);
- Making sensors practical and cost-effective to deploy, especially in scenarios such as

[8] http://varma.ece.cmu.edu/summit/CPS_Summit_Report.pdf
[9] http://cordis.europa.eu/fp7/ict/fet-proactive/2zerop_en.html

building retrofits or electric grid upgrades;
- Furthering networking to address the exploding number of sensors and controllers ("the network of things"), especially components that are mobile or have only fleeting encounters, such as vehicle-to-vehicle or vehicle-to-highway conversations that increase highway throughput or safety;
- Developing robots as mobile sensor and actuator platforms to be used for monitoring, e.g., autonomous underwater vehicles (AUVs) for ocean species tracking, as well as land, air, and sea robots for use in hazardous conditions such as the recent Gulf oil spill and Japanese nuclear reactor incident, storage facilities, and other repair and emergency response activities; and
- Continuing improvements in techniques for creating, testing, verifying and troubleshooting software that controls physical systems (many more such computerized controls will be embedded in devices that optimize the built environment, challenging a broader population of software developers to master these techniques).

One major research challenge is to ensure that systems employing intelligent control are truly robust. Not only must individual systems be robust, but the increasing interconnectedness of systems must not lead to instabilities or other robustness compromises.

3.6 Human-centered and social computing

Humans are arguably the most important (and problematic) players in achieving sustainability. Their attitudes about energy, transportation, and other systems, and their effectiveness at interacting with such systems will be extremely important. People will be asked to operate increasingly complex systems, telling computers their preferences for ride-sharing, deferring the laundry to save energy, understanding complex electricity pricing, and making many other tradeoffs. Moreover, these systems will try to induce people to modify their behavior to advance sustainability. If people resist because they do not understand what is asked of them, because the interfaces are too difficult or cumbersome to use, or because they feel the systems unduly intrude upon their lives, the sustainability efforts will be thwarted.

To date, research in human-computer interaction offers only partial solutions. We are able to make interfaces *usable* at a syntactic level—but we know very little about devising systems that induce cooperation when motivation is weak. We have seen a dramatic rise in computer systems for social networking, which tap into people's inclination to socialize—to find and interact with friends, people like themselves, and affinity groups. But we do not yet know what properties of the conceptual system design or user interface lead to success, or how these effects might be harnessed to societal needs, as in the case of sustainability.

Moreover, further research into human-centered computing can address the key problem of trust. Many surveys have shown that customers do not really trust utility companies. They are suspicious that new technologies, such as smart meters, are a ploy to make more money[10].

[10] The rollout of smart electric meters in Northern California is encountering resistance from customers concerned about erroneous bills, glitches in measurement, communications outages, and the like. As of March 2011, the utility has been forced to offer "opt out" alternatives that collect less data, but of course the meters still record total watt-hour usage.

Providing better information to consumers may help alleviate some of these concerns. Ultimately, trust is developed through series of interactions that prove convenient, valuable, and trouble-free.

Popular consumer computing devices today are not viewed as computers, but as smartphones, media players, information appliances, cameras, and the like. Their appeal is due to a complex combination of factors, but good design—and good user interfaces—is essential. A key issue thus entails understanding how this kind of product appeal can be harnessed to promote sustainability, both in the short and long terms.

3.7 Privacy and security

Data sensed about people, their possessions, movements, and other behaviors will be essential for optimizations to improve energy efficiency, traffic flow, and other sustainability-related purposes as outlined above. Because much of these data will be collected and stored by services not controlled by those providing the data, the privacy of the data must be protected. Also critical is the issue of who owns or has rights to the data; for example, can customers obtain data an electric company has collected about them?

Additionally, many government and industry efforts to measure and optimize resource use will want access to data collected about people as they pursue their daily lives. Aggregation and other ways to "anonymize" data must be used to avoid identifying individuals in the data and thus violate their privacy. Recent research has shown flaws in these techniques (Narayanan and Shmatikov, 2008), and consequently this is an area of essential research that could have an enormous impact on the efficacy of resource-optimization efforts.

An optimized, efficient world will be far more interconnected even than today's systems, which means that security will also be important for communications that link operational systems, e.g., different components of the smart electric grid, or transportation systems that are trying to coordinate schedules. If these communications are safety-critical [as in supervisory control and data acquisition (SCADA) controls in use today], it will be especially important that they be secure and robust against cyberattack.

3.8 Systems engineering and system integration

A collection of systems striving for efficiency—for example, to conserve resources such as energy or water—will require previously disparate systems to interact with one another to achieve jointly optimal outcomes. Systems that manage different sources of electricity, for instance, must communicate in real time to coordinate their separate efforts to meet demand. If a wind farm predicts declining winds, gas turbine generators managed by another system may need to be brought online in time to avoid a dip in power. Indeed, part of the inefficiency of the current electric grid is due to its overdesign to sustain peak loads, with relatively little effort to dynamically balance supply and demand. By contrast, the smart grid envisions real-time interaction with electricity consumers—cars, houses, buildings, trains, etc.—and with sensors and controls throughout the grid, to manage and schedule demand. To achieve this vision, the smart grid will require communications and protocols to "integrate" its various systems.

Moreover, many of the systems working toward sustainability will involve humans as system components. The "human component" is quite different from computerized control components: humans make different kinds of errors, but they may also be able to detect and avert unforeseen disasters. Human components operating inside complex control systems usually do so via human-computer interfaces, which themselves may be critical to human performance in the system.

Systems engineering is a body of technique and practice to cope with such large, complex engineering projects. Complexity usually arises from the interrelation of large numbers of components and subsystems. Systems engineering is brought to bear when a system has fluid boundaries or requirements or when the "straightforward" decomposition of a system into separate components is not advisable because the resulting assembly would lack important systemic properties such as robustness, affordability, flexibility, maintainability, and the like. Robustness—the property that a system must not fail despite the failure of a small number of its components—is usually easier to achieve with various systemic redundancy techniques than by making each component super-reliable.

Many tough problems face engineering complex interconnected systems. Some of these problems fall within CISE disciplines precisely because the systems being interconnected are operated by computer controls, and aspects of the design of those controls and the protocols that interconnect them are CISE engineering problems. A sampling of the problems that arise is presented below; while all these problems are addressed today, the solutions are not ideal and may improve with further research.

- How does a system detect when a system it's communicating with is "down" or has failed? Is every such failure handled gracefully, i.e., can a single failed component bring down a very large system?
- How can a new version of a system be deployed without disrupting those it interacts with? How are versions of communications protocols managed?
- How does a system defend against errors or misoperation by its peers?
- What kinds of system-wide defects can be detected, e.g., instability, "hunting," livelock, deadlock? Some attempts to detect or avoid these problems can lead to far worse problems.
- How can the operation of large complex system be understood? How is trouble detected and diagnosed?
- What if a component's behavior changes over time, such as a "learning" system whose behavior changes based on historical records?
- In the face of a changing external environment and perhaps changing performance of component parts, how is a systemic performance objective met?
- Before a new component is introduced into the system, how do we know that it operates correctly with its neighbors? The practice today for networked components is to organize "connectathons" in which new components are tested[11]. But these are usually crude tests and fall short of building confidence that components will work correctly in all settings.

[11] http://www.connectathon.org

A critical organizing principle for an interconnected system is its *information architecture*, i.e., the definitions of objects and actions that form a "language" for communication and control among systems and components. Current design efforts for the smart grid are creating its information architecture (IEEE Standards Association, 2010). While such a design is essential, we have few tools for validating such designs, for insuring conformance, or for managing changes over time.

3.9 Green IT

Most of the research directions described up to this point have outward-looking components, i.e., using CISE to fundamentally address sustainability challenges. However, the CISE community has long engaged in research efforts in sustainability that are directed *inward*, focusing on the resource consumption and environmental impact of our own technology. Indeed, "Green IT"—the collection of efforts to manage power consumption and improve overall sustainability of IT equipment—has become a focus for both manufacturers and operators of computing equipment. The Green IT Council[12] is a group of companies committed to "the study and practice of designing, manufacturing, using, and disposing of computers, servers, and associated subsystems—such as monitors, printers, storage devices, and networking and communications systems—efficiently and effectively with minimal or no impact on the environment." Manufacturers are beginning to use materials that can be sustainably harvested or more easily reused or recycled. Computer operators, principally those who operate large data centers, are striving to conserve energy by methods such as more efficient cooling (in some cases avoiding electrically-driven chillers) and turning computers on and off to respond to variations in demand. More and more, attention is being placed upon "cradle-to-cradle" approach, i.e., developing tools and devices with sustainability in mind over the entire lifecycle of the products.

There are several research thrusts that contribute to Green IT, including power-efficient microprocessors, dynamic data-center management software to optimize energy efficiency, etc. "Power-aware computing," which trades power for the timeliness or precision of the answer, is also being explored. This area has been part of the CISE research agenda for over 10 years, targeting all layers of system design and operation, but ongoing research is required to adapt to changes in IT technology and the IT environment. A related, nascent, and potentially transformative addition to the repertoire of tools is energy complexity analysis of algorithms (Bingham and Greenstreet, 2008; Albers, 2010; Jain *et al.*, 2005).

As described previously, one particular case of energy-efficient computing, given the name "energy harvesting," refers to self-contained computing elements with built-in sustainable energy sources. Sensors in a cornfield, for example, powered by solar cells and communicating by radio, can be made to operate using whatever energy is available. Sensors embedded in some structures and vehicles can harvest energy from vibration or acceleration to power their computing and communication. Clearly, these sensors must do their job using very little power, often powering down for 90 percent or more of the time. In addition to being sustainable, these designs use radio communications to avoid all external wired infrastructure. Although energy-harvesting devices are in their infancy, the technique holds great promise for deploying the enormous number of sensors required to drive the sustainability methods described elsewhere in this report.

[12] http://www.greenitcouncil.org

Importantly, many of these research areas are already being undertaken by industry, and academic researchers must consider partnerships with their industrial counterparts in order to have tangible impact.

4. The power of use-inspired (collaborative) fundamental research

As described above, "curiosity-driven" fundamental computing research is necessary to further sustainability—and there have been many instances of this type of research that have yielded enormous practical payoffs. For example, those who did seminal work on algorithms that find the shortest paths in graph structures could not have foreseen modern versions of these being embedded into today's GPS software to help drivers navigate toward their destinations. By contrast, **use-inspired fundamental research** stands to directly address the hardest problems posed by sustainability while at the same time yielding results that apply to other problems and also further CISE generally. This is the kind of research and innovation described in Stokes' famous *Pasteur's Quadrant*: solutions to a critically-important problem require new science or engineering developments, which themselves can be applied to solve a wider range of additional problems than the problems that initially motivated the research (Stokes, 1997). Starting with a concrete, applied problem can yield results that apply more broadly. A good recent example is the role of CISE research in decoding the genome: fundamental science and engineering in both CISE and molecular biology contributed essentially (see Box 4 below).

Some CISE researchers are reluctant to work on applied problems out of concern for the time and effort to become familiar with other domains, as well as the fear that they will be perceived as simply being the "coders" for other scientists. Others attempt to enter a new domain and find "low-hanging fruit" to which they can apply their tools and techniques, and then move on to the next application area. However, use-inspired fundamental research (as defined by Pasteur's quadrant) is where lots of great computing research takes place—and the computing research community cannot shy away from it in the context of sustainability. After all, computer science exists to solve problems and change the world.

The most successful outcomes occur when CISE researchers find collaboratively-minded domain scientists and invest significant time and energy in learning new application areas. Simply trusting the domain scientists to know in advance what CISE capabilities they require is seldom effective. Once engaged in the application area, however, CISE researchers must define suitable abstractions of the problems to be solved to be able to either find existing solution techniques or to devise new ones. These abstracted problems and their solution techniques can then be "exported" to the rest of the CISE community to stimulate further research and applications.

These kinds of collaborative teams require remarkable leadership and effort to build and maintain. In addition, there are remaining (though arguably declining) impediments to such structures—academic recognition, funding sizes and durations, intellectual-property claims, and others. However, we can turn to existing examples of multi-disciplinary research centers devoted to the pursuit of sustainability goals, such as the two NSF Expeditions in Computing cited above as well as the Center for Computational Learning Systems at Columbia, noteworthy for its strong industrial collaborations (see Box 5 below).

Box 4: Transformative research: The genome sequencing example

CISE can enable transformative breakthroughs in sustainability, much as it has already done in other fields. Consider, for example, how CISE radically altered the course of genome sequencing. Traditional genome sequencing was a linear process: a genome was separated into countless smaller sequences, and these were sequenced one by one. In 1981, a computer scientist named Gene Myers—working alongside geneticist Jim Weber—proposed a novel approach: why not break the genome into multiple, differently-sized fragments, sequence those pieces in parallel, and then develop computational algorithms to fit them back together? While the rest of the community considered the huge percentage of the human genome comprising repeating non-coding DNA as a hurdle to sequencing, Myers and Weber sought to use these segments to their advantage. The approach involved sequencing the unique DNA sequences, and using the non-coding repeats as markers to rebuild/reconnect the fragments.

By the 1990s, Myers, Weber, and others at Celera Genomics were developing an algorithm—500,000 lines of code – implementing the computational component, i.e., aligning non-repeating sequences paired with repeating sequences with other sets of such sequences. Coupled with high-end gene sequencers and massive computing power, they were able to sequence and assemble—with five-fold coverage—a 14.8 billion-base-pair human genome in under 9 months, a feat an international effort reliant on traditional sequencing methods could not accomplish in over 15 years. Ultimately, rather than pursuing a divide-and-conquer strategy, Celera *brought biology and CISE together to transform genome sequencing*.

In the decade since this research was completed, there have been a series of technological advances, resulting in what is today referred to as next-generation sequencing (NGS). In NGS, the fragments are much shorter than before (roughly 100 nucleotides, as opposed to over 1,000), but the number of fragments that can be sequenced rapidly in parallel is vastly increased. NGS is advancing by leaps and bounds, and new sequencing algorithms are being developed to handle the multitude of short reads. The applications are advancing as well; as the cost of sequencing an individual's genome plummets (today it is under $10,000 and continuing to drop rapidly), the possibility of personalized genetic testing and counseling arises.

Adapted from (Yung, 2006).

Box 5: Developing collaboration between Columbia's Center for Computational Learning Systems and Con Edison
The Center for Computational Learning Systems (CCLS), directed by Prof. David Waltz, was established to conduct research in machine learning and data mining, eyeing the explosion of data collected by industries, governments, and universities and the need to understand the data and its implications. It has built interdisciplinary ties within Columbia University as well as strong collaborations with industry with the objective of doing "Pasteur's Quadrant research." Collaboration with Consolidated Edison (Con Ed) was championed by Arthur Kressner, then Director of R&D for Con Ed of New York. He invited a small group of students and staff from CCLS to spend almost a year inside Con Ed, learning about their operations and problems. Ten research proposals resulted, and Con Ed chose to explore the use of machine learning to guide predictive maintenance choices—field repairs estimated to be most cost-effective at preventing future outages. Predictions were based on previous maintenance records and fragmented records of the equipment installed in Manhattan's grid. Records of previous repairs were often terse, missing information, or wrong. The result of the research is a system that works for Con Ed; it has also produced research papers. CCLS has melded government funding from NSF and DoE with industrial contributions from partners and collaborators like Con Ed, GE, and Federal Express. About half of the work of the center is on smart grid projects, including a recent GE Ecomagination award to work with Con Ed to develop ways to optimize the recharging of electric vehicles for a new facility FedEx is building in Manhattan. Waltz points to several essential characteristics of their Con Ed collaboration. A strong champion within the partner company, Center staff and students with persistence and with social and technical skills to work for and with the partner, and a sense of mission—in this case, to keep New York's lights on. The CCLS team worked hard, especially at first, to establish its value to Con Ed—on the company's terms, not as academic researchers. When presenting work at conferences—and even receiving some "best paper" awards—the CCLS team is eager to invite the Con Ed partners along. Although this collaboration has taken several years to develop, both Waltz and Kressner are confident that fascinating problems and research lie ahead.

5. Education and workforce development

The efforts described in this report cannot succeed without education of a next-generation workforce, and this in turn requires embedding computing in sustainability education *and* embedding sustainability in computing education. In the former case, sustainability programs, departments, and schools are forming at colleges and universities (e.g., Arizona State University, Indiana University, Oregon State University). The ubiquity of computing and its substantial relevance across the sustainability landscape, as this report highlights, speaks to the desirability of its inclusion in sustainability programs. In fact, students motivated to learn about and do laboratory work in sustainability topics may also learn to use computing in their work: to collect and analyze data, to build models, to run simulations, and even to develop new software. (This requires careful monitoring, in order to ensure that the sustainability students learn appropriate computing principles.)

The best way to promote computing in sustainability education is to bring sustainability into the computing curriculum, where students who want to tackle *real* problems can come to grips with challenges and approaches. The opportunities fill the wide range of computing and sustainability topics, ranging from low power computing for the developing world in the context of a computer organization course, to reverse engineering the backend database of a dorm-energy monitoring system in a database course, to solving problems in wildlife reserve design in an optimization course. In most cases, educational embedding will correspond to "routine use" of computing for sustainability that we have characterized in this report. Nonetheless, such infusion will help grow a next generation of computing professionals, including researchers and educators, who are motivated and capable of tackling tougher sustainability challenges.

The infusion of sustainability into the computing curriculum needs to be coupled with (a) substantial and well-designed outreach plans that will pass the benefit beyond the walls of the developing institution, and (b) the evaluation of curricular changes with behavioral and educational science expertise so that judgments can be made on the effectiveness of the changes, perhaps along measurable indices of future orientation and environmental stewardship.

6. The importance of collaborative Federal investment

Achieving long-term sustainability requires addressing a broad array of challenges, from near-term issues of climate change and energy efficiency to more long-term questions about how to make sense of a deluge of heterogeneous sustainability data, revamp and rebuild entire systems and infrastructures, and encourage human behavior change. Moreover, given the pervasiveness of information and communications technologies in our daily lives, it is clear that research in nearly every facet of CISE is critical to facilitating the necessary transformations for moving toward a sustainable ecosystem.

These advances require basic research in computer science, in tandem with advances in areas of systems engineering, the social and behavioral sciences, and relevant sustainability science and engineering domains (e.g., biology, climate science, etc.). As the mission of the NSF is to support discovery, innovation, learning, and infrastructure in each of these disciplines, NSF's investment in sustainability is essential in order for the opportunities outlined in this paper to be realized. We therefore laud NSF's Science, Engineering, and Education for Sustainability (SEES) initiative[13], focused on addressing challenges in climate, energy, and transportation through basic research and education. By including all of the NSF directorates, SEES seeks to foster an interdisciplinary, systems-based approach to understanding, predicting, and reacting to change in the linked natural, social, and built environment. It is critical for this initiative to remain highly interdisciplinary, with CISE researchers playing a core role on equal footing with other scientists and engineers.

---

[13] http://www.nsf.gov/sees

At the same time, the basic research cannot be undirected. To achieve advances in the sustainability sciences, the research must be informed from the outset by the various intended domains of use, that is, by expertise and progress in all aspects of sustainability science and engineering. Thus, collaboration with a number of mission-oriented government agencies is essential. For example, the U.S. Department of Energy (DoE) is focused on a comprehensive strategy for reducing energy consumption and making better use of renewable energy sources. Of particular note are DoE's Office of Energy Efficiency and Renewable Energy (EERE) and Advanced Research Projects Agency-Energy (ARPA-E). The Department of Transportation (DoT) oversees an aging transportation infrastructure. The Departments of Agriculture and Interior, along with the Environmental Protection Agency, have oversight on issues of climate change mitigation and adaptation, habitat preservation, etc. Working with these departments to coordinate and sponsor research can increase the likelihood that the research will lead to impact.

Ultimately, **mechanisms that support collaboration between computing researchers, systems engineers, social scientists, *and* sustainability scientists and engineers are the single most important component of the research and education agenda we have prescribed above**. That research interaction can be achieved even without the collaboration of every Federal agency delineated above, but the combined effort of these agencies will go far in accelerating this work.

7. Conclusion

The CISE research community sits in a unique position to make vital contributions to efforts to achieve a sustainable world. This report—and the workshop that it covers—describes areas in which computing research is likely to have strong impact. Given the ubiquitous nature of CISE technologies—from scientific study via data analysis, modeling, and simulation; to the development of optimized, intelligent systems; to the use of human-centered and social networking tools that account for individuals' security and privacy; to cyber-physical systems that bridge the computational and physical worlds—research in CISE areas can make significant contributions to sustainability efforts. Indeed, there are projects and groups already conducting research to address sustainability topics, often with academic computer scientists working in collaboration with industry or colleagues in other disciplines.

The research we have described in this report as necessary to address our sustainability challenges requires a strong commitment by multiple Federal funding agencies to support multi-disciplinary approaches, with the consequent need to invest time and talent to build deep connections across disciplinary and institutional boundaries. Importantly, as research programs are crafted, whether by funders or by researchers, they should address the need to build large and highly collaborative teams, working together to solve difficult, varied problems.

8. References

Albers, S. (2010). Energy-efficient algorithms. *Communications of the ACM* **53**(5): 86-96.

Bingham, B.D. and Greenstreet, M.R. (2008). Computation with energy-time trade-offs: Models, algorithms, and lower-bounds. In *Proceedings of the International Symposium on Parallel and Distributed Processing and Applications (ISPA-08)*, Sydney, Australia, pp. 143-152.

Chubin, D.E., DePass, A.L. and Blockus, L. (2009). *Understanding interventions that broaden participation in research careers: Embracing a breadth of purpose* (a report on the third annual conference, held in Bethesda, MD, May 7-9). Washington: AAAS.

Conrad, J., Gomes, C.P., van Hoeve, W.-J., Sabharwal, A. and Suter, J (2010). Connections in networks: Hardness of feasibility versus optimality. In *Proceedings of the 7th International Conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR)*, Bologna, Italy, pp. 16-28.

Dilkina, B. and Gomes, C.P. (2010). Solving connected subgraph problems in wildlife conservation. In *Proceedings of the 7th International Conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR-10)*, Bologna, Italy, pp. 102-116.

Ezell, S. (2010). Bringing U.S. roads into the 21st century. *Issues in Science and Technology* **26**(4). Retrieved from http://www.issues.org/26.4/ezell.html; last accessed April 4, 2011.

Gomes, C.P. (2009). Computational sustainability: Computational methods for a sustainable environment, economy, and society. *The Bridge* **39**(4): 5-13. Retrieved from http://www.nae.edu/File.aspx?id=17673; last accessed April 4, 2011.

Gomes, C.P., van Hoeve, W.-J., and Sabharwal, A. (2008). Connections in networks: A hybrid approach. In *Proceedings of the 5th International Conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR-08),* Berlin, Germany, pp. 303-307.

IEEE Standards Association (2010). *IEEE P2030 Draft Guide for Smart Grid Interoperability of Energy Technology and Information Technology Operation with the Electric Power System (EPS), and End-Use Applications and Loads.* Retrieved from http://grouper.ieee.org/groups/scc21/2030/2030_index.html; last accessed April 4, 2011.

Interagency Climate Change Adaptation Task Force (2010). *Progress Report of the Interagency Climate Change Adaptation Task Force: Recommended Actions in Support of a National Climate Change Adaptation Strategy*, The White House Council On Environmental Quality. Retrieved from http://www.whitehouse.gov/administration/eop/ceq/initiatives/adaptation; last accessed March 29, 2011.

Jain, R., Molnar, D. and Ramzan, Z. (2005). Towards a model of energy complexity for algorithms. *Wireless Communications and Networking Conference* 3: 1884-1890.

Katz, R. (2008). "Smart Grid": R&D for an intelligent 21st century electrical energy distribution infrastructure. Computing Community Consortium. Retrieved from http://www.cra.org/ccc/docs/init/Energy_Grid.pdf; last accessed April 4, 2011.

Lai, K.J., Gomes, C.P., Schwartz, M., McKelvey, K., Calkin, D. and Montgomery, C. (2011).

The Steiner Multigraph Problem: Wildlife corridor design for multiple species. *Proceedings of the 25th Conference on Artificial Intelligence, Special Track on Computational Sustainability and AI (AAAI-11)*, to be held in San Francisco, California, Aug. 7-11.

McDonough, W. and Braungart, M. (2002). *Cradle to Cradle: Remaking the Way We Make Things*. New York: North Point Press.

Miller, C.A. and Edwards, P.N. (2001). Representing the global atmosphere: Computer models, data, and knowledge about climate change. In *Changing the atmosphere: Expert knowledge and environmental governance*. Cambridge: MIT Press.

Narayanan, A. and Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *Proceedings of the 29th IEEE Symposium on Security and Privacy*, Oakland, California, pp. 111-125.

National Science Foundation (2008). NSF Announces Expeditions in Computing Awards. Retrieved from http://www.nsf.gov/news/news_summ.jsp?cntn_id=112075&org=CISE; last accessed March 2, 2011.

National Science Foundation (2010). NSF Announces New Expeditions in Computing Awards. Retrieved from http://www.nsf.gov/news/news_summ.jsp?cntn_id=117560&org=CISE&from=home; last accessed March 2, 2011.

Petersen, J.E., Shunturov, V., Janda, K., Platt, G., and Weinberger, K. (2007). Dormitory residents reduce electricity consumption when exposed to real-time visual feedback and incentives. *International Journal of Sustainability in Higher Education* **8**(1): 16-33.

President's Council of Advisors on Science and Technology (2010). Designing a digital future: Federally funded research and development in networking and information technology. The White House Executive Office of the President. Retrieved from http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-nitrd-report-2010.pdf; last accessed April 4, 2011.

Proceedings of the National Academy of Sciences (2011). Sustainability Science. Retrieved from http://www.pnas.org/site/misc/sustainability.shtml; last accessed March 2, 2011.

Service, R.F. (2011). Coming soon to a lab near you: Drag-and-drop virtual worlds. *Science* **331**: 669-671.

Shepherd, B. and Whittington, J. (2006). Response of wolves to corridor restoration and human use management. *Ecology and Society* **11**(2): 1.

Stokes, D.E. (1997). *Pasteur's Quadrant*. Brookings Institution Press, http://www.brookings.edu/press/Books/1997/pasteur.aspx.

Thrun, S. and Kelly, H. (2009). Surface transportation 3.0. Computing Community Consortium. Retrieved from http://www.cra.org/ccc/docs/init/Surface_Transportation_3.0.pdf; last accessed April 4, 2011.

Transportation for America, ITS America, the Association for Commuter Transportation and the University of Michigan's SMART Initiative. Smart Mobility for a 21st Century America. Retrieved from http://t4america.org/wp-content/uploads/2010/10/ITS-White-Paper-100710-FINAL.pdf; last accessed April 3, 2011.

United Nations General Assembly (2005). 2005 World Summit Outcome. Resolution A/60/1, adopted by the General Assembly on September 15, 2005. Retrieved from http://data.unaids.org/Topics/UniversalAccess/worldsummitoutcome_resolution_24oct2005_en.pdf; last accessed March 20, 2011.

U.S. Fish and Wildlife Service (2006). Florida Panther and the Genetic Restoration Program. Retrieved from http://www.fws.gov/southeast/news/2006/images/PantherGeneticFactSheet.pdf; last accessed March 29, 2011.

Waltz, D., and King, J. (2009). Information technology and America's energy future. Computing Community Consortium. Retrieved from http://www.cra.org/ccc/docs/init/itandenergy.pdf; last accessed April 4, 2011.

Yung, E. (2006). Celera: A unique approach to genome sequencing. University of California, Berkeley, Biocomputing. Retrieved from http://www.ocf.berkeley.edu/~edy/genome/celera.html; last accessed March 21, 2011.