Researchers in computing, information science, and many other disciplines are working together to support sustainable development.

Computational Sustainability

Computational Methods for a Sustainable Environment, Economy, and Society

Carla P. Gomes

Carla P. Gomes is associate professor, Faculty of Computing and Information Science, Department of Computer Science and Department of Applied Economics and Management, Cornell University.

The dramatic depletion of natural resources in the last century now threatens our planet and the livelihood of future generations. *Our Common Future*, a report by the World Commission on Environment and Development published in 1987, introduced for the first time the notion of "sustainable development: development that meets the needs of the present without compromising the ability of future generations to meet their needs" (UNEP, 1987). The concerns raised in that report were reiterated by the Intergovernmental Panel on Climate Change (IPCC, 2007). In the fourth Global Environmental Outlook report published later that same year the authors concluded, "there are no major issues raised in *Our Common Future* for which the foreseeable trends are favorable" (UNEP, 2007).

Key issues in the development of policies for sustainable development will entail complex decisions about the management of natural resources and more generally about balancing environmental, economic, and societal needs. Making such decisions optimally, or nearly optimally, presents significant computational challenges that will require the efforts of researchers in computing, information science, and related disciplines, even though environmental, economic, and societal issues are not usually studied in those disciplines.

In this author's opinion, it is imperative that computer scientists, information scientists, and experts in operations research, applied mathematics, statistics, and related fields pool their talents and knowledge to help find efficient and effective ways of managing and allocating natural resources. To that end, we must develop critical mass in a new field, computational sustainability, to develop new computational models, methods, and tools to help balance environmental, economic, and societal needs for a sustainable future.

Examples of computational sustainability problems presented in this short paper range from wildlife preservation and biodiversity to balancing socio-economic needs and the environment to the large-scale deployment and management of renewable energy sources.

Computational sustainability will give us the tools to balance environmental, economic, and societal needs.

Biodiversity and Species Conservation

The reduction and fragmentation of natural habitats as a result of deforestation, agriculture, urbanization, and land development is a leading cause of species decline and extinction. One strategy for improving the chances of species viability is to protect habitats by creating biologically valuable sites or reserves. Examples include the National Wildlife Refuge System, managed by the U.S. Fish and Wildlife Service, national parks, and conservation reserves established by private groups, such as the Nature Conservancy and the Conservation Fund.

Given the limited resources available for conservation, these sites must be carefully chosen. From a mathematical point of view, the site-selection or reserve-design problem involves optimizing certain criteria, such as habitat suitability for species, while simultaneously satisfying one or more constraints, such as limited budgets (e.g., Ando et al., 1998; Moilanen et al., 2009; Polasky et al., 2008).

In recent years biologists attempting to combat habitat fragmentation have promoted so-called "conservation corridors," continuous areas of protected land that link biologically significant zones. The design of conservation corridors is a special aspect of the site-selection problem, and the objective is to create connected corridors made up of parcels of land that will yield the highest possible level of environmental benefit ("utility") (Onal and Briers, 2005; Williams et al., 2005).

At the Institute for Computational Sustainability (ICS) at Cornell University, we recently formulated this problem mathematically as a so-called "connection sub-graph problem" (Conrad et al., 2007; Dilkina and Gomes, 2009; Gomes et al., 2008). The goal was to design wildlife corridors for grizzly bears in the U.S. northern Rockies to enable movement between three core ecosystems—Yellowstone, Salmon-Selway, and Northern Continental Divide Ecosystems—that span 64 counties in Idaho, Wyoming, and Montana. This large-scale optimization problem places significant demands on current computational methods.

To scale up solutions, we needed a deeper understanding of the underlying structure of the problem. To that end, we developed a budget-constrained, utility-optimization approach using hybrid constraint-based mixed-integer programming that exploits problem structure. Our results showed that we can dramatically reduce the cost of large-scale conservation corridors by provably finding corridors with minimum cost. If more than minimum funding for a corridor is available, this approach guarantees optimal utility. For example, for the grizzly bear problem our solutions are guaranteed to be within 1 percent of the optimal solution for budget levels above the minimum cost.

Complexity in site-selection and corridor-design problems increases when different models for land acquisition over different time periods (e.g., purchase, conservation easements, auctions), dynamic and stochastic environments, and multiple species must be considered. For example, preserving bird habitats and designing bird corridors requires a good understanding of hemispheric-scale bird migrations with complex population dynamics across different climate and weather systems and geographic topologies.

Thus modeling complex species distributions and developing conservation strategies requires new large-scale stochastic-optimization methods. Moreover, to obtain the right model parameters and determine current species distribution, machine learning and statistical techniques must be used to analyze large amounts of raw data (Dietterich, 2009; Elith et al., 2006; Kelling et al., 2009; Phillips et al., 2004).

Gathering biological, ecological, and climatic data is essential to studying complex systems, and the deployment of large-scale sensor networks is becoming a key tool for environmental monitoring (e.g., Polastre et al., 2009; Werner-Allen et al., 2006). The National Science Foundation (NSF) supports several cyber-infrastructure initiatives for massive data collection and data analysis based on large-scale autonomous sensor networks, such as the National Ecological Observatory Network (NEON) and the Long-Term Ecological Research Network (LTER).

Designing a large-scale sensor network also presents computational challenges (e.g., network architecture, operating system and programming environments, data collection, analysis, synthesis, and inference) (Akyildiz et al., 2007). For example, when using sensor networks to monitor spatial phenomena, selecting the best placement of sensors to maximize information gain while minimizing communication costs is a complex problem that requires new techniques (Krause and Guestrin, 2009).

Citizen observation networks have several benefits. They help in collecting data and, at the same time, enable the general public to engage in scientific investigation and develop problem-solving skills. Galaxy Zoo,[1] for example, provides access to a large collection of images and engages the general public in classifying galaxy shapes to improve our understanding of their formation. eBird,[2] a joint initiative of the Cornell Laboratory of Ornithology and the National Audubon Society, engages citizen-scientists in observing birds using standardized protocols. Since eBird was released in 2002, it has been visited by more than 500,000 users and has collected more than 21 million bird records from more than 35,000 unique users in more than 180,000 locations across the Western Hemisphere and New Zealand (Sullivan et al., 2009).

Management of Natural Resources

This example concerns the state of marine fisheries. The biomass of top marine predators is estimated to be one-tenth of what it was half a century ago and is still declining (Worm et al., 2006). As a result of overfishing, pollution, and other environmental factors, many important marine species are extinct, with dramatic consequences for the filtration of nutrients by the ocean. Researchers believe that the collapse of major fisheries is primarily the result of mismanagement (Clark, 2006; Costello et al., 2008). Therefore, we must find sustainable ways of managing fisheries.

One approach that has been shown to be effective for counterbalancing the overharvesting of fisheries involves both placing limitations on total allowable catches per species and requiring permits for harvesting specific quantities of fish (individual transferable quotas) (Costello et al., 2008; Heal and Schlenker, 2008; Worm et al., 2009). Complex dynamical models, originally developed as part of dynamical systems theory, can be used to identify the optimal amount of fish that can be harvested annually in a certain fishery, taking into consideration re-generation rates, carrying capacity of the habitat, discount rates, and other parameters.

Dynamical systems theory, which provides tools for characterizing the dynamics and long-term behavior of systems as a function of the system parameters, provides insights into nonlinear system dynamics and identifies patterns and laws, particularly bifurcations (Ellner and Guckenheimer, 2006; Strogatz, 1994). A bifurcation occurs when small changes in the parameter values of a system (e.g., the rate of harvesting fish) lead to an abrupt qualitative change (e.g., the collapse of a fishery). Decisions (e.g., the amount of fish to be harvested) are often based on combinations of continuous and discrete variables. This leads to hybrid dynamical optimization models, which, in principle, provide information on optimal harvesting strategies (Clark, 1976; Conrad, 1999). However, finding such strategies is computationally difficult, especially when considering multiple species.

The biomass of top marine predators is about one-tenth of what it was 50 years ago.

Balancing Socioeconomic and Environmental Needs

Chris Barrett of ICS has studied the socioeconomic interrelationship between poverty, food security, and environmental stress in Africa, particularly links between resource dynamics and the poverty trap in small-holder agrarian systems (Barrett et al., 2007). Barrett's focus has been on pastoral systems in East Africa that involve herds of cattle, camels, sheep, and

[1] Available online at *http://www.galaxyzoo.org/*.
[2] Available online at *http://ebird.org/content/ebird*.

goats (Luseno et al., 2003). Due to high variability in rainfall, pastoralists must migrate with their herds looking for water and forage, sometimes traveling as much as 500 kilometers.

The purpose of our studies is to develop a predictive model of the migratory patterns and decision models of these pastoralists. To do that, we use machine-learning methods to determine the structure and estimate the parameters of the models, based on field data about households, water sources, and climate patterns.

Ultimately, these models will help policy makers predict the effects of potential policy interventions and environmental changes, with the goal of improving the livelihoods of thousands of pastoralists. The project involves new technical approaches to large, structural-dynamic, discrete-choice problems that will lead to the development of computational models to support both descriptive studies and predictive policy analyses (Toth et al., 2009).

Other computational sustainability topics in this context include automated decision-support tools for providing humanitarian aid in response to catastrophes, famines, and natural disasters in developing countries. The design of such systems will require the development of intuitive, user-friendly interfaces for use by aid workers.

Data centers emit more CO_2 than Argentina and the Netherlands.

Energy-Efficient Data Centers

The implications of climate change for environmental, economic, and social systems have led to major changes in energy policy in many industrial countries, including incentives for increasing energy efficiency. These incentives present tremendous computational opportunities for helping to increase energy efficiency through the design of intelligent or "smart" control systems for energy-efficient buildings, vehicles, and appliances.

According to the World Business Council for Sustainable Development (2008), buildings account for as much as 40 percent of energy use in industrialized countries. Data centers (i.e., computing facilities with electronic equipment for data processing, storage, and communications networking) are especially inefficient users of energy.

In recent years the shift to digital services has led to a major increase in demand for data centers. The Environmental Protection Agency estimates that in the next decade the demand for data-center capacity will grow at a 10 percent compounded annual growth rate (EPA, 2007). In addition, the costs of data centers in the information technology (IT) sector are estimated to increase at an annual rate of 20 percent, compared to an overall increase in IT of 6 percent (Kaplan et al., 2008).

Data centers also have negative environmental impacts. According to a recent report, the amount of carbon dioxide emissions produced by data centers worldwide exceeds the total emissions of both Argentina and the Netherlands (Kaplan et al., 2008). Thus the IT industry is looking to advanced power-management hardware, smart cooling systems, virtualization tools, and dense server configurations to reduce energy consumption (Katz, 2009).

These new approaches rely heavily on large amounts of data provided by large-scale sensor networks (e.g., Bodik et al., 2008; Hoke et al., 2006; Patnaik et al., 2009; Shah et al., 2008). Some companies are using containers that integrate computing, power, and cooling systems in one module for data centers, instead of raised-floor rooms. Several IT companies are committed to using alternative energy sources, such as hydropower, solar power, and wind power, to bring the carbon footprint of data centers to zero.

On a larger scale, data centers can contribute to reductions in energy use and carbon emissions by facilitating e-commerce and telecommuting, for example, which can eliminate some of the need for paper printing and for freight and passenger transportation.

The Smart Grid

Under the Energy Independence and Security Act (EISA) of 2007, the U.S. Department of Energy was charged with modernizing the nation's electricity grid to improve its reliability, efficiency, and security, a concept known as the Smart Grid. Ideally, the Smart Grid will radically transform the industry's business model from a largely non-digital, electromechanical grid to a network of digital systems and power infrastructure and from a centralized, producer-controlled network to a more decentralized system with more interaction between consumers and local producers.

The objectives for the Smart Grid include: enabling active participation by consumers; making possible the easy integration of a variety of generation options (with a focus on renewable sources) and storage options; enabling new products, services, and markets; providing quality power for the digital economy; optimizing assets and operating efficiently; automatically anticipating and responding to system disturbances; and operating resiliently in the event of attacks or natural disasters.

To realize these objectives, the Smart Grid will include smart sensors and controls throughout the transmission and distribution system and a broad communication platform for two-way communications to move data and electricity between utilities and consumers. For example, consumers will have smart meters that can track energy consumption, monitor individual power circuits in the home, control smart appliances, and actively manage energy use.

Planning and operating such a large, complex digital ecosystem will require technological advances in computing and information science related to sensing and measuring technologies, advanced control methods, monitoring and responding to events, support for dynamic pricing, computational aspects of game-theory models and mechanism design, multi-agent based models, improved interfaces, decision-support and optimization tools, and security and privacy tools.

Renewable Energy

The development of renewable energy can have an even greater environmental impact than increasing energy efficiency. In recent years technological progress has been made (partly in response to government incentives) in renewable energy sources, such as biofuels and biomass, geothermal, solar, and wind power. For example, EISA set fuel economy standards for vehicles that will require the production of 36 billion gallons of renewable fuels per year by 2022, a fivefold increase over current ethanol production levels.

The logistics and planning of this large-scale domestic-based biofuels production system raise complex stochastic optimization problems—variants of the so-called "facility-location problem"—that must take into consideration feedstock and demand and the dynamics of demand and capacity (Shmoys, 2004). And the stakes are high. Finding good solutions to these problems can make the difference between economic viability and failure. Overall, we will need complex computational models to find the best mix of energy generation and storage technologies.

A larger project will be the development of computational models (Figure 1) that show interactions between different energy sources and the agents directly or indirectly involved (e.g., households, landowners, farmers, ethanol producers, gasoline refiners, food producers) and impacts on the environment (e.g., greenhouse gas emissions, water, soil erosion, biodiversity, etc.).

To begin with, the overall impact of biofuels is not well understood. Take, for example, their impact on land use. Traditional life-cycle studies do not take into account emissions from changes in land use, which are difficult to quantify (Seager et al., 2009; Searchinger et al., 2008; Tilman et al., 2009).

Another example is the impact of wind power, a promising renewable energy source that has raised concerns about damage to bird and bat populations. Research will be necessary to provide guidelines for the

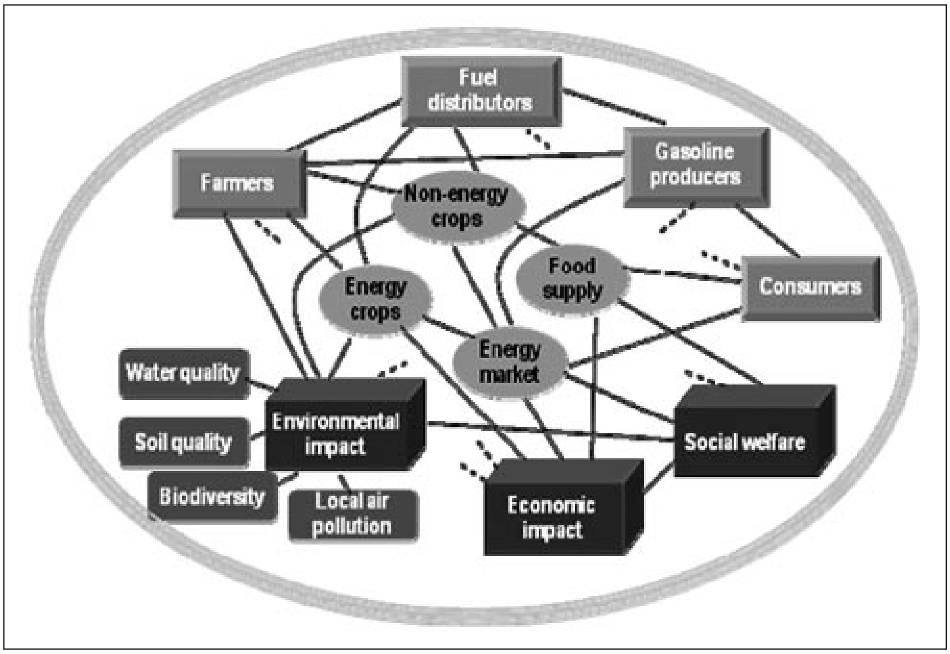

FIGURE 1 Interacting components for biofuel analysis.

location of wind farms, especially because most areas with favorable winds are associated with important migratory pathways.

The research challenge is to develop realistic models that capture multiple impacts and interdependencies without imposing strong (unrealistic) assumptions. In traditional approaches, convexity assumptions force unique equilibria, or at the very least, the set of equilibria are themselves convex (Codenotti et al., 2005; Heijungs and Suh, 2002; Ye, 2008). This has made their algorithmic solution possible, but such models do not capture key aspects of systems. Researchers will have to develop more complex decision models through collaboration with resource economists, environmental scientists, and computer scientists.

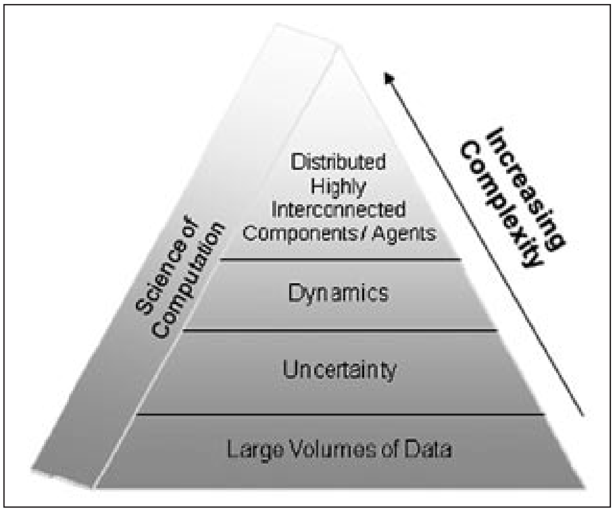

FIGURE 3 Increasing levels of complexity in computational sustainability problems.

Individual Interests vs. the Common Good

A key issue in environmental policy is balancing individual interests and the common good (e.g., Hardin, 1968). In this area, game-theory models can model the interactions of multiple agents and show the effects of competing interests. In the context of natural resources or climate change on the international level, for example, economic incentives may influence whether a country is motivated to enter an agreement and then abide by it.

Incentive-based policies can also facilitate sustainability challenges on a smaller scale (e.g., the establishment of novel markets for land-conservation activities). To be useful, multi-agent models will have to explore mechanisms and policies for the exchange of goods.

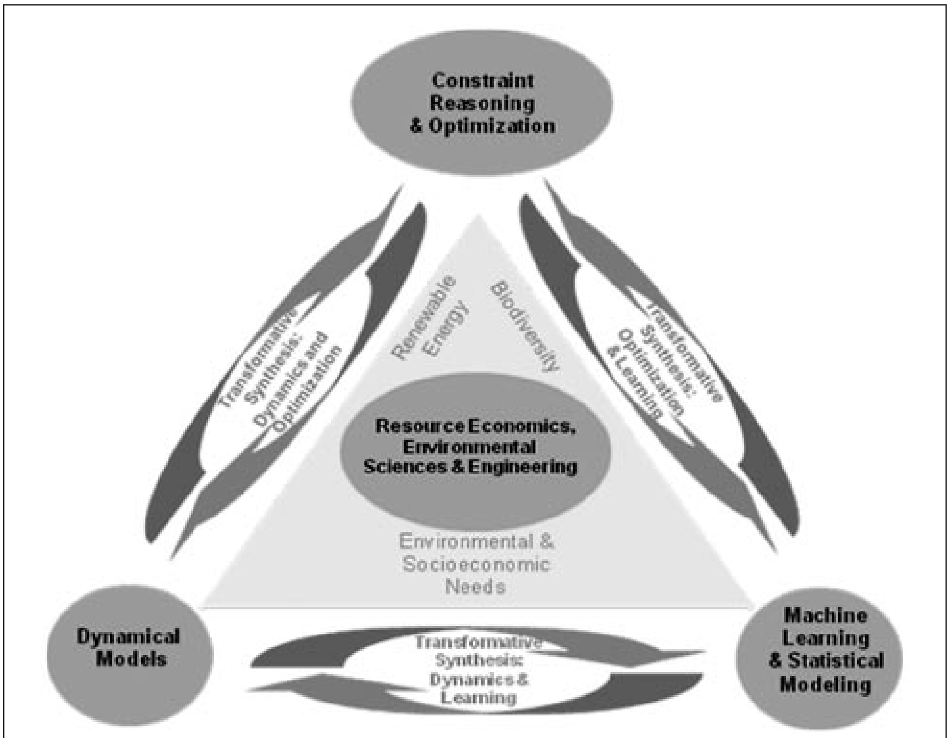

FIGURE 2 Examples of research themes and interactions in computational sustainability that are closely aligned with the research agenda of the Institute for Computational Sustainability at Cornell University.

The Research Challenges

Research in computational sustainability involves many different areas in computing, information science, and related disciplines. Figure 2 shows some of the areas that are closely related to examples in this article and to the ICS research agenda (ICS, 2010). Figure 3 shows the levels of complexity in computational sustainability, which often addresses large-scale problems based on large volumes of data in highly dynamic and uncertain environments with many interacting components.

Given these complexities, the study of computational sustainability problems requires a fundamentally new approach that is unlike the traditional computer science approach (i.e., the science of computation), which is driven mainly by worst-case analyses. From the perspective of computational sustainability, problems are considered "natural" phenomena that are amenable to scientific methodology, rather than purely mathematical abstractions or artifacts. In other words, to capture the structure and properties of complex real-world sustainability problems, principled experimentation is as important as formal models and analysis (Gomes and Selman, 2005, 2007).

Summary

The development of policies for a sustainable future presents unique computational problems in scale, impact, and richness that will create challenges, but also opportunities, for the advancement of the state of the art of computer science and related disciplines. The key research challenges are developing realistic computational models that capture the interests and interdependencies of multiple agents, often involving continuous and discrete variables, in a highly dynamic and uncertain environment.

Research in this new field is necessarily interdisciplinary, requiring that scientists with complementary skills work together. In fact, collaboration is an essential aspect of the new science of computational sustainability, an interdisciplinary field that applies techniques from computer science, information science, operations research, applied mathematics, statistics, and related fields to help balance environmental, economic, and societal needs for a sustainable future.

The focus is on developing computational and mathematical models, methods, and tools for making decisions and developing policies concerning the management and allocation of resources for sustainable development. The range of problems encompasses computational challenges in disciplines from ecology, natural resources, economics, and atmospheric science to biological and environmental engineering. Computational sustainability opens up fundamentally new intellectual territory with great potential to advance the state of the art of computer science and related disciplines and to provide unique societal benefits.

Acknowledgments

The author is the director of the Institute for Computational Sustainability (ICS) at Cornell University and the lead principal investigator of an Expedition in Computing grant on computational sustainability from the National Science Foundation (NSF) (Award Number: 0832782). The author thanks NSF for its support and ICS members for their many contributions to the development of a vision for computational sustainability.

References

Akyildiz, I.F., T. Melodia, and K.R. Chowdhury. 2007. A survey on wireless multimedia sensor networks. Computer Networks: The International Journal of Computer and Telecommunications Networking 51(4): 921–960.

Ando, A., J. Camm, S. Polasky, and A. Solow. 1998. Special distributions, land values, and efficient conservation. Science 279(5359): 2126–2128.

Barrett, C., P. Little, and M. Carter, eds. 2007. Understanding and Reducing Persistent Poverty in Africa. Florence, Ky.: Routledge.

Bodik, B., M.P. Armbrust, K. Canini, A. Fox, M. Jordan, and D.A.. Patterson. 2008. A Case for Adaptive Datacenters to Conserve Energy and Improve Reliability. Technical Report No. UCB/EECS-2008-127. Berkeley, Calif.: EECS Department, University of California, Berkeley.

Clark, C. 1976. Mathematical Bioeconomics: The Optimal Management of Renewable Resources. New York: Wiley.

Clark, C. 2006. The Worldwide Crisis in Fisheries. New York: Cambridge University Press.

Codenotti, B., S. Pemmaraju, and K. Varadarajan. 2005. On the polynomial time computation of equilibria for certain exchange economies. Pp. 72–81 in Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms. Philadelphia: SIAM.

Conrad, J. 1999. Resource Economics. New York: Cambridge University Press.

Conrad, J., C. Gomes, W.-J. van Hoeve, A. Sabharwal, and J. Suter. 2007. Connections in Networks: Hardness of Feasibility versus Optimality. Pp. 16–28 in Proceedings of the Fourth International Conference on the Integration of AI and OR Techniques Constraint Programming, Brussels, Belgium. Berlin, Heidelberg: Springer-Verlag.

Costello, C., S.D. Gaines, and J. Lynham. 2008. Can catch shares prevent fisheries collapse? Science 321(5896): 1678–1681.

Dietterich, T. 2009. Machine Learning in Ecosystem Informatics and Sustainability. Pp. 8–13 in Proceedings of the 21st International Joint Conference on Artificial Intelligence. Pasadena, Calif.: IJCAI.

Dilkina, B., and C. Gomes. 2009. Wildlife Corridor Design:

Connections to Computer Science. Pp. 3–4 in Proceedings of the Workshop on Constraint Reasoning and Optimization for Computational Sustainability (CROCS-09). Lisbon: CP-09.

Elith, J., C.H. Graham, R.P. Anderson, M. Dudik, S. Ferrier, A. Guisan, R.J. Hijmans, F. Huettmann, J.R. Leathwick, A.Lehmann, J. Li, L.G. Lohmann, B.A. Loiselle, G. Manion, C. Moritz, M. Nakamura, Y. Nakazawa, J.McC. Overton, A.T. Peterson, S.J. Phillips, K. Richardson, R. Scachetti-Pereira, R.E. Schapire, J. Soberón, S. Williams, M.S. Wisz, and N.E. Zimmermann. 2006. Novel methods improve prediction of species' distributions from occurrence data. Ecography 29(2): 129–151.

Ellner, S., and J. Guckenheimer. 2006. Dynamic Models in Biology. Princeton, N.J.: Princeton University Press.

EPA (Environmental Protection Agency) 2007. Report to Congress on Server and Data Center Energy Efficiency. Washington, D.C: EPA. Available online at: *http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency_study*.

Gomes, C., and B. Selman. 2005. Computation science: can get satisfaction. Nature 435(7043): 751–752.

Gomes, C., and B. Selman. 2007. Science of constraints. Constraint Programming Letters 1: 15–20.

Gomes, C., W.-J. van Hoeve, and A. Sabharwal. 2008. Connection in Networks: A Hybrid Approach. Pp. 303–307 in Proceedings of the Fifth International Conference on the Integration of AI and OR Techniques Constraint Programming, Paris, France. Berlin, Heidelberg: Springer Verlag.

Hardin, G. 1968. The tragedy of the commons. Science 162(3859): 1243–1248.

Heal, G., and W. Schlenker. 2008. Sustainable fisheries. Nature 455(7216): 1044–1045.

Heijungs R., and S. Suh. 2002. The Computational Structure of Life Cycle Assessment (Eco-Efficiency in Industry and Science). Boston: Kluwer Academic Publishers.

Hoke, E., J. Sun, and C. Faloutsos. 2006. InteMon: Intelligent System Monitoring on Large Clusters. Pp. 1239–1242 in Proceeding of Very Large Databases (VLDB'06). Available online at *http://www.vldb.org/conf/2006/p1239-hoke.pdf*.

Luseno, W.K., J.G. McPeak, C.B. Barrett, P.D. Little, and G. Gebru. 2003. Assessing the Value of Climate Forecast Information for Pastoralists: Evidence from Southern Ethiopia and Northern Kenya. World Development 31(9): 1477–1494.

ICS (Institute for Computational Sustainability). 2010. Agenda. Available online at *http://www.cis.cornell.edu/ics/*.

IPCC (Intergovernmental Panel on Climate Change). 2007. Fourth Assessment Report (AR4). Technical Report, United Nations IPCC. Available online at *http://www.ipcc.ch/publications_and_data/publications_and_data_reports.htm#1*.

Kaplan, J., W. Forrest, and N. Kindler. 2008. Revolutionizing the Data Center Energy Efficiency. Technical report produced by McKinsey & Company.

Katz, R.H. 2009. Tech Titans Building Boom. IEEE Spectrum Online. Available online at *http://www.spectrum.ieee.org/feb09/7327*.

Kelling, S., W.M. Hochachka, D. Fink, M. Riedewald, R. Caruana, G. Ballard, and G. Hooker. 2009. Data intensive science: a new paradigm for biodiversity studies. BioScience 59(7): 613–620.

Krause, A., and C. Guestrin. 2009. Optimizing sensing from water to the Web. IEEE Computer Magazine 42(8): 38–45.

Moilanen, A., K. Wilson, and H. Possingham, eds. 2009. Spatial Conservation Prioritization. New York: Oxford University Press.

Onal, H., and R. Briers. 2005. Designing a conservation reserve network with minimal fragmentation: a linear integer programming approach. Environmental Modeling and Assessment 10(3): 193–202.

Patnaik, P., S.R. Marwah, and N. Ramakrishnan. 2009. Sustainable Operation and Management of Data Center Chillers using Temporal Data Mining. In Proceedings of Knowledge Discovery and Data Mining (KDD'09), June 28–July 1, 2009, Paris. France. New York: Association for Computing Machinery.

Phillips, S., M. Dudik, and R.E. Schapire. 2004. A Maximum Entropy Approach to Species Distribution Modeling. P. 83 in Proceedings of the Twenty-First International Conference on Machine Learning, Banff, Alberta, Canada. New York: Association for Computing Machinery.

Polasky, S., E. Nelson, J. Camm, B. Csuti, P. Fackler, E. Lonsdorf, C. Montgomery, D. White, J. Arthur, B. Garberyonts, R. Haight, J. Kagan, A. Starfield, and C. Tobalske. 2008. Where to put things? spatial land management to sustain biodiversity and economic returns. Biological Conservation 141(6): 1505–1524.

Polastre, J., R. Szewczyk, A. Mainwaring, D. Culler, and J. Anderson. 2009. Analysis of Wireless Sensor Networks for Habitat Monitoring. Pp. 399–424 in Wireless Sensor Networks, edited by C.S. Raghavendra. K.M. Sivalingam, and T. Znati. Dordrecht: Springer Netherlands.

Seager, T., S.A. Miller, and J. Kohn. 2009. Land Use and Geospatial Aspects in Life Cycle Assessment of Renewable Energy. Pp. 1–6 in Proceedings of IEEE International Symposium on Sustainable Systems and Technology. Available online at *http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5156724&isnumber=5156678*.

Searchinger, T., R. Heimlich, R.A. Houghton, F. Dong, A. Elobeid, J. Fabiosa, S. Tokgoz, D. Hayes, and T.-H. Yu. 2008. Use of U.S. croplands for biofuels increases greenhouse gases through emissions from land-use change. Science 319(5867): 1238–1240.

Shah, A., V. Carey, C. Bash, and C. Patel. 2008. Energy analysis of data center thermal management systems. Journal of Heat Transfer 130(2): 021401 (10 pages).

Shmoys, D. 2004. The Design and Analysis of Approximation Algorithms: Facility Location as a Case Study. Pp. 85–98 in Trends in Optimization, edited by S. Hosten, J. Lee, and R. Thomas. AMS Proceedings of Symposia in Applied Mathematics, Vol. 61. Providence, R.I.: American Mathematical Society.

Strogatz, S. 1994. Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering. Jackson, Tenn.: Perseus Books Group.

Sullivan, B., C. Wood, M. Iliff, R. Bonney, D. Fink, and K. Kelling. 2009. eBird: a citizen-based bird observation network in the biological sciences. Biological Conservation 142(10): 2282–2292.

Tilman, D., R. Socolow, J. Foley, J. Hill, E. Larson, L. Lynd, S. Pacala, J. Reilly, T. Searchinger, C. Somerville, and R. Williams. 2009. Beneficial biofuels—the food, energy, and environment trilemma. Science 325(5938): 270–271.

Toth, R., C. Barrett, Y. Guo, and C. Gomes. 2009. Pastoral Systems in East Africa. Manuscript in preparation.

UNEP (United Nations Environment Programme). 1987. Our Common Future. Published as annex to the General Assembly document A/42/427, Development and International Cooperation: Environment. Report of the World Commission on Environment and Development. Nairobi, Kenya: UNEP.

UNEP. 2007. Global Environment Outlook 4 (GEO4). Technical Report. Available online at *http://www.unep.org/geo/*.

WBCSD (World Business Council for Sustainable Development). 2008. Energy Efficiency in Buildings: Facts and Trends—Full Report. Available online at *http://www.wbcsd.org/DocRoot/qUjY7w54vY1KncL32OVQ/EEB-Facts-and-trends.pdf*.

Werner-Allen, G., K. Lorincz, M. Welsh, O. Marcillo, J. Johnson, M. Ruiz, and J. Lees. 2006. Deploying a wireless sensor network on an active volcano. IEEE Internet Computing 10(2): 17–25.

Williams, J.C., C.S. ReVelle, and S.A. Levin. 2005. Spatial attributes and reserve design models: a review. Environmental Modeling and Assessment 10(3): 163–181.

Worm, B., E.B. Barbier, N. Beaumont, J.E. Duffy, C. Folke, B.S. Halpern, J.B.C. Jackson, H.K. Lotze, F. Micheli, S.R. Palumbi, E. Sala, K.A. Selkoe, J.J. Stachowicz, and R. Watson. 2006. Impacts of biodiversity loss on ocean ecosystem services. Science 314(5800): 787–790.

Worm, B., R. Hilborn, J.K. Baum, T.A. Branch, J.S. Collie, C. Costello, M.J. Fogarty, E.A. Fulton, J.A. Hutchings, S. Jennings, O.P. Jensen, H.K. Lotze, P.M. Mace, T.R. McClanahan, C. Minto, S.R. Palumbi, A.M. Parma, D. Ricard, A.A. Rosenberg, R. Watson, and D. Zeller. 2009. Rebuilding global fisheries. Science 325(5940): 578–585.

Ye, Y. 2008. A path to the Arrow-Debreu competitive market equilibrium. Mathematical Programming 111(1–2): 315–348.