DOI:10.1145/3183582

The future of computing research relies on addressing an array of limitations on a planetary scale.

BY BONNIE NARDI, BILL TOMLINSON, DONALD J. PATTERSON, JAY CHEN, DANIEL PARGMAN, BARATH RAGHAVAN, AND BIRGIT PENZENSTADLER

Computing within Limits

COMPUTING RESEARCHERS AND practitioners are often seen as inventing the future. As such, we are implicitly also in the business of predicting the future. We plot trajectories for the future in the problems we select, the assumptions we make about technology and societal trends, and the ways we evaluate research.

However, a great deal of computing research focuses on one particular type of future, one very much like the present, only more so. This vision of the future assumes that current trajectories of ever-increasing production and consumption will continue. This focus is perhaps not surprising, since computing machinery as we know it has existed for only 80 years, in a period of remarkable industrial and technological expansion. But humanity is rapidly approaching, or has already exceeded, a variety of planet-scale limits related to the global climate system, fossil fuels, raw materials, and biocapacity.[28,32,38]

It is understandable that in computing we would not focus on limits. While planetary limits are obvious in areas such as extractive capacity in mining or fishing, or the amount of pollution an ecosystem can bear, limits are less obvious in computing. Many believe the only limit worth considering is human ingenuity, and that we can surpass any and all other limits if we, as a global community, pool our creative resources. But we collectively face new global conditions that warrant our attention.

In this article we explore the relationship between these potential futures and computing research. What hidden assumptions about the future are embedded in most computing research? What possible or even probable futures are we ignoring? What work should we be doing to respond to fundamental planetary limits, and to the ecological and energy constraints that global society faces over the coming years and decades? Confronting such limits is likely to present challenges that we—humanity—have never before faced.

Given that computing underlies virtually all the infrastructure of global society—in commerce, communication, transportation, agriculture, manufacturing, education, science, healthcare, and governance—computing has an enormous role to play in responding to global limits and in shaping a society that meaningfully adapts to them. We contend that the root of much of computing research has been driven predominantly by growth-oriented visions

» key insights

- **Most computing work is premised on industrial civilization's default worldview in which ongoing economic growth is both achievable and desirable.**
- **This growth-focused worldview, however, is at odds with findings from many other scientific fields, which see growth as deeply problematic for ecological and social reasons.**
- **We proposed that the computing field transition toward "computing within limits," exploring ways that new forms of computing supported well-being while enabling human civilizations to live within global ecological and material limits.**
- **Computing underlies virtually all the infrastructure of global society, and will therefore be critical in shaping a society that meaningfully adapts to global limits.**

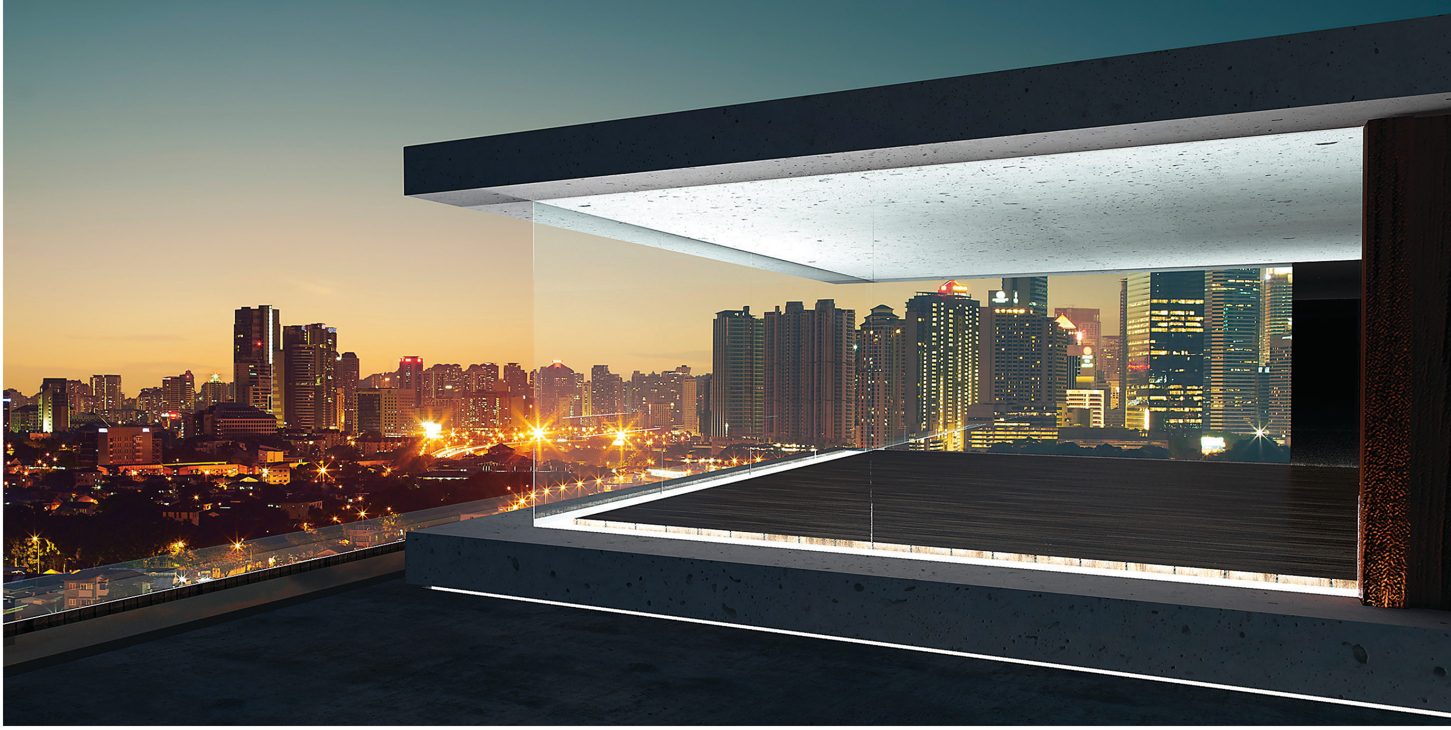

of society's future.[26,34,39] If we broaden our view to a more diverse set of possible futures, including non growth-reliant futures, the societal challenges of ecological and energy limits can shape concrete technical challenges in computing research and practice.

In order to consider these futures, we have been building a community of scholars from computer science and engineering, information science, social science, ecology, agriculture, and earth sciences to explore what we call "computing within limits" or "LIMITS" for short. The LIMITS research community integrates three topics: current and near-future ecological, material, and energy limits; the ways new forms of computing may help support well-being while living within these limits; and the impact these limits are likely to have on the field of computing. LIMITS is concerned with the material impacts of computation itself, but, more broadly and more importantly, it engages a deeper, transformative shift in computing research and practice to one that would use computing to contribute to the overall process of transitioning to a future in which the well-being of humans and other species is the primary objective.

The LIMITS perspective is related to Green IT,[17] sharing an interest in improvements in efficiency and other traditionally "green" research topics. However, LIMITS research questions Green IT's implicit assumption that we can "engineer around" the finiteness of the Earth's resources and waste capacity. LIMITS sees ecological and environmental issues as a "predicament"— that is, a situation for which there are not likely to be clear-cut "solutions" but rather a constellation of complex issues that requires broad new assumptions and approaches. We seek to engage this predicament by adopting a new framing for computing research. We question the focus on ongoing economic growth that lies at the heart of industrial civilization and propose a shift from emphasis on standards of living and material productivity to an emphasis on long-term well-being. LIMITS research looks ahead to future scenarios cognizant of work such as that of Rockström et al.[28] that draws attention to "planetary boundaries that must not be transgressed." Each of these topics will be discussed in greater detail.

Here, we present background literature in ecological economics and archaeology that has informed LIMITS research, and then review computing research in sustainable human computer interaction, crisis informatics, and information and communication technology for development (ICTD). Although LIMITS researchers come from many subfields of computing including networking and software engineering, research in these three areas in particular is closely related to LIMITS with potential for deeper fu-

PHOTOS BY (TOP) JAMES TEOH ART; (BOTTOM) QUICK SHOT ture connections. We then briefly summarize the three annual workshops on LIMITS that began in 2015. Finally, we discuss several key principles that have arisen from LIMITS work to guide future research. We see work in this area as a subfield that is an important alternative to traditional growth-oriented computing research.

Computing has an enormous role to play in responding to global limits and in shaping a society that meaningfully adapts to them.

Background

Since the beginning of computing, all research and development has taken place against a backdrop of exponential growth of, for example, transistors per integrated circuit (Moore's Law), disk storage density (Kryder's Law), bandwidth capacity (Nielsen's Law), and fiber-optic capacity (Keck's Law). These developments have led to the establishment of a "cornucopian paradigm"[23] where the design of new services stimulates demand, which drives growth of increased infrastructure capacity, which then cycles back to enable the design of new services in a self-perpetuating cycle. The idea that exponential growth of computing capacity and an ever-expanding infrastructure for computing will continue into the future is usually taken for granted. We draw from research in ecological economics and the historical record in archeology to question this assumption.

This research suggests that other futures are not just possible but probable. While most economists sidestep questions of finite resources,[6] economists in the subfield of ecological economics have grappled with these questions for decades. How can we maintain or increase well-being while staying within ecological limits? How can we promote well-being and not exceed the assimilative and regenerative capacities of the Earth's biochemical life-support systems? We have already exceeded many such limits through, for example, overfishing, deforestation, soil depletion, falling water tables, rising temperatures, and emitting CO_2 and other greenhouse gases at rates that dangerously increase their concentrations in the atmosphere.[28,32,38] Ecological economist Herman Daly has proposed that we abandon the idea of striving for economic growth in favor of a steady-state economy (in line with classical economist Adam Smith's idea that the economy would eventually reach a "stationary state"). A steady-state economy would maintain material throughput at a rate that is largely stable across time and that remains within ecological limits.[7] At the same time, Daly notes that culture and society need not be static: "Not only is quality free to evolve, but its development is positively encouraged in certain directions. If we use 'growth' to mean quantitative change, and 'development' to refer to qualitative change, then we may say that a steady-state economy develops but does not grow, just as the planet Earth, of which the human economy is a subsystem, develops but does not grow." Daly suggests that a single-minded focus on growing the economy comes at the eventual cost of decreasing human well-being and quality of life. Such growth results in, for example, charging for things that used to be free, the health consequences of polluting the environment, and decreasing long-term possibilities to produce food or earn a livelihood.

Looking at societal trends through the lens of human history, archaeologist Joseph Tainter's book *The Collapse of Complex Societies* argues that civilizations eventually collapse, declining over a period of decades or centuries.[33] Analyzing extensive historical and archaeological materials, Tainter presented collapse as a process that arises from increasing societal complexity, which, over time, creates burdens for systems that they eventually cannot sustain.

Decline will result in less material abundance as we push the limits of the Earth's resources necessary for economic activity. But it is not necessary for our society to end in abject collapse. The societies that Tainter studied—the Maya, the Mesopotamians, the Minoans, the Inca, the Romans, the Egyptians, and others—did not possess the resources of science, history, and technology that we have amassed in the last 500 years. These resources have the potential to be usefully deployed to fashion a transition from the current, unsustainable system to a new system based on today's realities. We optimistically assume that with advances in science and progress in philosophies of human rights, we have a good chance of transformative change to a system more like the steady-state economy Herman Daly envisions. The implication of the work in ecological economics and archaeology is that we should endeavor to build computer systems that aim at increasing well-being and quality of life while contributing to staying within ecological limits. Foregrounding human well-being is supported by the ACM Code of Ethics and Professional Conduct, the first imperative of which states: "As an ACM member I will contribute to society and human well-being." (https://www.acm.org/aboutacm/acm-code-of-ethics-and-professional-conduct)

We turn now to a review of computing literature that has been foundational for the development of computing within LIMITS perspectives.

SCHI: Sustainable Human-Computer Interaction

The Sustainable Human-Computer Interaction community is about a decade old, and a number of LIMITS researchers have roots in this area. Eli Blevis's "Sustainable Interaction Design"[3] is a primary source, offering a rubric to identify how interaction designs lead to material effects, as well as several principles for engaging in sustainable interaction design. Early papers that sparked interest among LIMITS researchers were Jeff Wong's "Prepare for Descent: Interaction Design in Our New Future"[40] and Silberman and Tomlinson's "Precarious Infrastructure and Postapocalyptic Computing."[31] Several high-profile CHI papers drew attention to the challenges of sustainability and the shortcomings of SHCI work in failing to address questions of physical, material, and energy limits. DiSalvo et al.'s "Mapping the Landscape of Sustainable HCI"[8] sought to provide structure to the array of papers in SHCI, and identified gaps in the areas being studied, such as the need to focus on collectives and broader contexts, not just individuals, the importance of engaging with policy issues, and stronger connections to sustainability work in fields outside of computing.

From this context, Tomlinson et al.'s "Collapse Informatics"[35] was the first full treatment of LIMITS topics in the SHCI community. This paper explored "the study, design, and development of sociotechnical systems in the abundant present for use in a future of scarcity." This work helped lay the groundwork, along with papers from other subfields of computing[24,37] for LIMITS research.

LIMITS has drawn heavily from collapse informatics but shifts emphasis to planetary limits rather than societal decline. LIMITS focuses on exposing basic processes of resource use and waste management in complex human systems. The metrics used to assess sustainability must shift correspondingly. As examples, Pargman and Raghavan's "Rethinking Sustainability in Computing: From Buzzword to Non-negotiable Limits"[20] and Raghavan and Pargman's "Means and Ends in Human-Computer Interaction: Sustainability through Disintermediation,"[25] offer major contributions, arguing that "sustainability" must be grounded in rigorous metrics arising from planetary limits, and that the complexity of societal systems might be reduced, easing resource use and waste production. The forthcoming edited collection *Digital Technology and Sustainability: Engaging the Paradox*[10] incorporates influences from LIMITS research. Several of the papers mentioned here as well as Preist et al.[23] have won best paper awards, signaling interest in the issues.

Crisis Informatics

We are often asked if computing within LIMITS is the same as crisis informatics. Crisis informatics is concerned with technology-based studies of disaster planning and response, and constitutes an important subfield of human-computer interaction.[19] There are some key differences between crisis informatics and LIMITS, although we think that in the future the two may increasingly mutually inform one another. At present, crisis informatics research generally assumes an external entity that enacts a rescue when a disaster, such as a flood or earthquake, occurs. Events are conceived as localized, describing a space into which the surrounding society can pour resources to alleviate the resulting disorder and disruption. These scenarios accurately describe an important subset of possible issues confronting human civilizations. LIMITS, however, assumes long time frames and a global spatial scale. There is no external entity to provide relief. LIMITS emphasizes phenomena such as climate change, soil erosion, water pollution, civic instability, mass migration, reduced infrastructure, and an economy that requires continuous growth.[4,5,14,20,21,24,30,36]

Potentially there is a strong link between LIMITS and crisis informatics. Some crisis informatics researchers are beginning to examine long-term processes underlying crises, suggesting that when looked at more broadly, "crises" are often more than acute events of short duration, with roots in underlying processes that may have been developing over decades.[1] This understanding provides a bridge for future development and crossfertilization between the two subdisciplines.

PHOTO BY VALENTIN VALKOV

PHOTO BY GENCHO PETKOV

ICTD: Information and Communication Technology for Development

ICTD is a relatively young field that has explored the potential of computing for improving the socioeconomic situation of the poor. While computing within LIMITS typically focuses on the future, Tomlinson et al.[35] note that our imagined "future" LIMITS scenarios may already exist today in the conditions in which poor communities live around the world. However, few studies within the ICTD literature consider global ecological, material, and energy limits. Most research is situated in resource-constrained contexts and assumes the constraints will be relaxed in the future after sufficient economic growth has occurred.[12,15] The only paper so far that explicitly makes the link between LIMITS and ICTD in an ICTD venue is Tomlinson et al.'s DEV paper, "Toward alternative decentralized infrastructures."[36] The vacuum regarding the implications of phenomena such as climate change in the ICTD literature could be filled by a LIMITS perspective.

There is, however, a tension between economic development in poor countries—the focus of ICTD—and sustainability. As Herman Daly points out, the total resource footprint of the Global North and the Global South combined together must stay within the boundaries of a global steady state economy that is sustainable in the long run. To ameliorate the problem of unequal distribution of wealth and the consequent problem of poverty in the Global South, the Global North must shrink its resource footprint enough that countries in the Global South are afforded some space for necessary economic growth. However, everyone—North and South—must operate within some absolute global limits. The ethical argument for improving the quality of life of the poor is easy to make, but reducing the Global North's consumptive (and exploitative) practices to afford the Global South opportunities to grow, especially in the face of mounting resource and climate pressures, remains an enormous challenge, and one computing should be cognizant of.

Despite differing perspectives, LIMITS and ICTD have much in common and potential for integration and collaboration.[4] For example, LIMITS work has studied the use of digital technology to design habitations in refugee camps,[29] problems of networking in rural populations in Zambia and Guatemala[30] and infrastructure in conditions of scarcity in Haiti.[21] While these are classic ICTD topics, the authors in each case considered ecological, material, and energy limits in their analyses, unlike typical ICTD studies. The papers engage models of scarcity, examining the cases as possible future global LIMITS scenarios. Drought, flooding, environmental disasters, infrastructure disruption, mass migration, and permanent settlement in refugee camps in low-resource environments are seen as highly relevant to global futures, not just as problems that will be solved through economic growth.

Computing Within Limits Workshops

LIMITS ideas have been developed through three workshops (2015–2017) convened by the LIMITS community (the latter two in cooperation with ACM). The first two were held at the University of California, Irvine, and the third at Westmont College in Santa Barbara, with funding from the two universities as well as from Facebook and Google. Participants came from institutions in Abu Dhabi, Canada, Hong Kong, Pakistan, Spain, Sweden, Switzerland, the U.K., and the U.S., consistent with the global nature of LIMITS concerns and research. The 2018 workshop was held in Toronto, co-located with the Fifth International Conference on Information and Communication Technology for Sustainability (ICT4S). Sparked by discussions at the workshops, LIMITS participants have co-authored several papers published in mainstream conferences and a research grant. The LIMITS workshop papers are available at computingwithinlimits.org

Three Key Principles

We propose three principles that can help frame computing research and practice in a way that is consistent with the ideas described in this paper and the literature we have surveyed.

Question growth. The industrialized world's current economic system, capitalism, is predicated on growth. Economic growth has brought more than an order of magnitude rise in per capita income from $3 a day in 1800 to $100 in the early 2000s for most of Europe and North America.[16] However, despite such unprecedented prosperity, global income inequality is increasing. Wealth is accumulating in the hands of fewer and fewer astoundingly rich persons.[22] Poverty is widespread. Such social dysfunction, along with the burdens on ecosystems produced by economic activity,[28,32,38] suggest we must rethink the growth paradigm. The ubiquity and power of computing make it well positioned to act as an agent of change to influence proposals for transformative economic systems and methods of governance. While discussion of specific proposals is beyond the scope of this article, we point to the work of, for example, Daniel O'Neill,[18] Peter Frase,[9] and Tim Jackson[13] as thoughtful responses to current problems that might inform the ways we practice computing.

Daly's notion of promoting development rather than economic growth suggests a sound mechanism for moving civilization forward, deploying our creativity and capacity for innovation in LIMITS-compliant ways. An economy that demands endless growth entails a cycle of consequences that must be interrupted if we are to address massive problems such as climate change and resource depletion.[20] Exploring relations between computing and the economy will be an important direction for future development of the computing community and a considerable challenge.

Currently, the implicit organizing framework for a great deal of computing work puts a focus on increasing the proximate financial value of companies. Even when particular products, from a narrow perspective, are seeking to make people's lives better through new technology, these products are typically embedded in a rapid churn of objects and services that foster runaway consumption.[23,27] By shifting the explicit focus, first and foremost, to the pursuit of long-term well-being, we may finally escape the growth paradigm and build systems that more effectively lead to sustainable improvements in the quality of life for humans and other species.

To make this principle actionable, we encourage researchers and practitioners to consider whether their work is a) reliant on growth, b) seeking to make growth happen, c) contributing to growth. We encourage those working in computing to build systems and envision worlds that are neither reliant on nor contributing to runaway growth. A number of existing LIMITS relevant papers have addressed this principle.[24,31,35]

Consider models of scarcity. Clever technological fixes may help us defer catastrophes for some time, but not indefinitely, and especially not if events such as wildfires, hundred-year storms, and Category 5 hurricanes become more numerous and more powerful as outcomes of global environmental changes. Our track record of being prepared for dealing with unpredictable catastrophic events is not encouraging. We would benefit from seriously considering LIMITS-related scenarios rather than blithely denying their possibility or treating their foreshocks as isolated incidents. Engaging with these difficult scenarios before they occur, rather than only in their aftermath, will help us evaluate our level of preparedness and perhaps prevent certain undesirable future scenarios from happening.[21]

We encourage those working in computing to build systems and envision worlds that are neither reliant on nor contributing to runaway growth.

To speak of LIMITS-scenarios only in the future tense, however, is misleading. These events are here now, as several climate-related catastrophes in the U.S. and Europe have shown, even during the writing of this article. Science fiction author William Gibson famously said, "The future is already here—it's just not evenly distributed." We see this future currently on display in places such as Flint, Michigan where toxic wastes have poisoned the water supply. It is thus possible to frame LIMITS scenarios (including, for example, heat waves, drought, rising sea levels, and floods) not in terms of random irregularities or threats that might afflict us in the future, but in terms of an increasing incidence of phenomena arising from intensive economic activity.

A concrete research strategy is to develop case studies of current changes that may model futures of relative scarcity. For example, a study of the continuing impact of the 2010 earthquake in Haiti found that the regrowth of infrastructures was occurring in a more distributed fashion than would be typical for countries with more resources.[21] Distribution networks for clean water, electricity, Internet, and gasoline were severely damaged in the earthquake. Corporate and government responses were hampered by political and financial obstacles. In many cases, survivors themselves began to rebuild the infrastructures in a bottom up manner. For example, large private water tanks were installed on local properties. Wealthier residences allowed adjacent poorer households to tap into power lines via jerry-rigged extension cords without paying for the service—a generous if somewhat precarious arrangement.

Such a re-arrangement certainly went against existing building codes, but recognized the low cost of alleviating some resource deprivation in exchange for neighborhood stability.

A case study such as this can be generative by revealing opportunities for developed and less developed regions to transfer technologies and schemes of sociotechnical organization that present a different set of economic incentives for actors. Being aware of the wide diversity of current and future potential contexts in which humans may find themselves, more than a few of which are characterized by scarcity, may help computing researchers and practitioners design technology that promotes global well-being.

Several other LIMITS-relevant papers have focused on aspects of this principle, including work found in Refs.[4,14,29,30,40]

Reduce energy and material consumption. Sticking to the dominant narrative of growth is riskier than just making a bad guess. It is dangerous because it creates a possibility that we will reach a point at which resources have precipitously dwindled and we may not have enough remaining resources to make the necessary corrections to avert catastrophic outcomes. Therefore, it is important to acknowledge that computing uses energy and material resources. If, as we have argued, these resources are declining, a threshold that LIMITS research should meet is that it is worth the resources it consumes. Put another way, LIMITS research, once applied, should reduce energy expenditures and material consumption. This reduction is difficult to assess, but not something we can sidestep.

LIMITS research, once applied, should reduce energy expenditures and material consumption. This reduction is difficult to assess, but not something we can sidestep.

More broadly speaking, attempts to limit resource usage in any human system are notoriously challenging. Most of us are well aware of the problems of CO_2 emissions, but less aware of more subtle dynamics such as the Jevons paradox, that is, that more efficient technologies often encourage greater use of a resource, reducing or eliminating savings. A more efficient gas engine may reduce fuel consumption by half, but stimulate more than twice as much driving (as well as more cars). A more efficient cryptocurrency mining chip effectively increases electricity consumption through competitive pressure. Mitigating the Jevons paradox requires creative approaches that may include substitution of goods by services and dematerialization, for example, by virtualization.[11] Such changes have the potential to entail a drop in absolute consumption, although so far, most approaches have tended to focus on increasing efficiency, which may or may not result in absolute reductions.[13] However, there is scope for significant change; for example, the energy costs of a virtual meeting that transmits data to a large number of remote participants is tiny compared to the energy cost of a single airplane trip for a single participant. The energy needed for data transmission is decreasing at a fast pace, unlike the energy costs of air travel. Aslan et. al.[2] estimate that data transmission costs decrease by 50% every two years.

Accounting for resource use must be done thoughtfully, with long-term goals in mind, in view of the big picture. There is justification for spending resources during a time of relative abundance to prepare for a future of scarcity.[12] Not all investments need to pay off immediately. There is a place for experimenting when we don't know for sure if savings will be accrued. But such experimentation should fail fast, and have a plausible hope of saving resources. In this regard, we need to be cognizant of the power of capital markets in deciding what is a success and what is a failure. While markets are very good at optimizing the delivery of the goods and services that they incentivize, they tend not to be organized in such a way that promotes long-term returns or incorporates the costs of the externalities that push limits. Structural changes such as cap-and trade markets, taxes, fees, rationing, and quotas are needed, in concert with technological changes, to address these issues.

Another key approach involves finding energy savings through disintermediation, that is, the process of leveraging technology to supplant "middleman" actors in resource chains.[25] Traditionally, in the absence of information technology, such middlemen provided value and extracted costs by creating markets and distribution centers for goods. For example, systems to directly connect small-scale worker/producer owned facilities with consumers could be of value in a new economy. Such simplification is responsive to Tainter's argument that increasing complexity leads to increasing burdens for systems which at some point they cannot bear.[33] Technologies that provide services while reducing complexity at the same time, square conceptually with what we know from the historical and archaeological record about the relationship between increasing societal complexity and eventual societal decline. This and other efforts at disintermediation[3,36] could help reduce energy and material consumption.

Conclusion

While we do not know for certain what the future holds, scientists from disciplines such as climate science and ecology have made evidence-based predictions about directions the future will likely take if current trends continue. However, what many computing researchers and practitioners do in practice is to assume there is only one possible likely future—that current trajectories of increased growth and consumption will continue. The burden of our message in this article is that science is telling us the kinds of growth we have recently experienced are unsustainable. Consequently, we believe the field of computing should be paying serious attention to futures in which we encounter planetary limits.

LIMITS thinking emphasizes incentivizing long-term returns. It seeks to align its efforts with the scientific disciplines documenting global transformations through climate change and numerous other global effects. LIMITS seeks to explore ways that computing may support long-term well-being. We see significant cause for concern in many science-based projections of the future, and we want to enable our work to be relevant and useful with respect to these potential realities.

Acknowledgments

The authors thank several anonymous reviewers for their cogent comments, as well as the entire LIMITS community for helping shape the ideas in this article. This material is based in part on work supported by the NSF under Grants No. CCF-1442749 and IIS-0644415.

References

1. Anderson, J. et al. Far far away in Far Rockaway: Responses to risks and impacts during Hurricane Sandy through first-person social media narratives. In *Proceedings of the 13th International Conference on Information Systems for Crisis Response and Management*, (2016).
2. Aslan, J, Mayers, K., Koomey, J. and France, C. Electricity intensity of Internet data transmission: Untangling the estimates. *J. Industrial Ecology*, (2017).
3. Blevis, E. Sustainable interaction design: Invention & disposal, renewal and reuse. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (2007), ACM, 503–512.
4. Chen, J. Computing within limits and ICTD. In *Proceedings of the 1st Workshop on Computing within Limits*, (2015). ACM.
5. Chen, J. A strategy for limits-aware computing. In *Proceedings of the 2nd Workshop on Computing within Limits*, (2016). ACM.
6. Costanza, R. *Ecological Economics: The Science and Management of Sustainability.* Columbia University Press, 1992.
7. Daly, H. *Steady State Economy.* Island Press, 1977.
8. DiSalvo, C, Sengers, P. and Brynjarsdóttir, H. Mapping the landscape of sustainable HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (2010). ACM, 1975–1984.
9. Frase, P. *Four Futures.* Verso, London, U.K., 2016.
10. Hazas, M. and Nathan, L. (Eds.). *Digital Technology and Sustainability: Engaging the Paradox.* Routledge, New York, 2018.
11. Hilty, L.M. and Aebischer, B. (Eds.). *ICT Innovations for Sustainability.* Vol. 310. Springer International Publishing, 2015.
12. Houston, L. and Jackson, S. Caring for the next billion mobile handsets: Opening proprietary closures through the work of repair. In *Proceedings of the 8th International Conference on Information and Communication Technologies and Development*, (2016). ACM.
13. Jackson, T. *Prosperity without Growth.* Routledge, London, U.K., 2017.
14. Jang, E., Johnson, M., Burnell, E. and Heimerl, K. Unplanned obsolescence: Hardware and software after collapse. In *Proceedings of the 3rd Workshop on Computing within Limits*, (2017). ACM.
15. Masinde, M., Bagula, A. and Muthama, N. The role of ICTs in downscaling and up-scaling integrated weather forecasts for farmers in sub-Saharan Africa. In *Proceedings of the 5th International Conference on Information and Communication Technologies and Development*, (2012). ACM, 122–129.
16. McCloskey, D. *Bourgeois Dignity.* University of Chicago Press, Chicago, 2010.
17. Murugesan, S. Harnessing green IT: Principles and practices. *IT Professional 10*, 1 (2008), 24–33.
18. O'Neill, D. Measuring progress in the degrowth transition to a steady state economy. *Ecological Economics 84* (2011), 1–11.
19. Palen, L., Starbird, K., Vieweg, S. and Hughes, A. Twitter-based information distribution during the 2009 Red River Valley flood threat. *Bulletin of the American Society for Information Science and Technology 36*, 5 (2010), 13–17.
20. Pargman, D. and Raghavan, B. Rethinking sustainability in computing: From buzzword to non-negotiable limits. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction*, (2014). ACM, 638–647.
21. Patterson, D. 2015. Haitian resiliency: A case study in intermittent infrastructure. In *Proceedings of the 1st Workshop on Computing Within Limits*. ACM, 111–117.
22. Piketty, T. *Capital in the 21st Century.* Belknap Press, Cambridge, MA, 2014.
23. Preist, C., Schien, D. and Blevis, E. Understanding and mitigating the effects of device and cloud service design decisions on the environmental footprint of digit al infrastructure. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 1324–1337).
24. Raghavan, B, and Ma, J. Networking in the long emergency. *Proceedings of the 2nd ACM SIGCOMM Workshop on Green Networking*, (2011). ACM, 37–42.
25. Raghavan, B. and Pargman, D. Means and ends in human-computer interaction: Sustainability through disintermediation. In *Proceedings of the 2017 SIGCHI Conference on Human Factors in Computing Systems.* ACM, 786–796.
26. Ramchurn, S., Vytelingum, P., Rogers, A. and Jennings, N. Putting the 'smarts' into the smart grid: A grand challenge for artificial intelligence. *Commun. ACM 55*, 4 (Apr. 2012), 86–97.
27. Remy, C. and Huang, E. Addressing the obsolescence of end-user devices. In *ICT Innovations for Sustainability*. Springer, New York, 2014.
28. Rockström, J. et al. A safe operating space for humanity. *Nature 461* (2009), 472–475.
29. Sabie, S., Chen, J., Abouzied, A., Hashim, F., Kahlon, H. and Easterbrook, S. Shelter dynamics in refugee and IDP camps: Customization, permanency, and opportunities. In *Proceedings of the 3rd Workshop on Computing Within Limits*, (2017). ACM, 11–20.
30. Schmitt, P. and Belding, E. Navigating connectivity in reduced infrastructure environments. In *Proceedings of the 2nd Workshop on Computing Within Limits*. ACM.
31. Silberman, M.S. and Tomlinson, B. Precarious infrastructure and postapocalyptic computing. In *Proceedings of CHI 2010 Workshop on Examining Appropriation, Re-use, and Maintenance for Sustainability*.
32. Steffen, W. et al. Planetary boundaries: Guiding human development on a changing planet. *Science 347* (2015), 6223.
33. Tainter, J. *The Collapse of Complex Societies.* Cambridge University Press, Cambridge, U.K., 1990.
34. Thrun, S. et al. Stanley: The robot that won the DARPA Grand Challenge. *J. Field Robotics 23*, 9 (2006), 661–692.
35. Tomlinson, B., Six Silberman, M., Patterson, D., Pan, Y., and Blevis, E. Collapse informatics: augmenting the sustainability & ICT4D discourse in HCI. In *Proceedings of the 2012 SIGCHI Conference on Human Factors in Computing Systems*. ACM, 655–664.
36. Tomlinson, B., Nardi, B., Patterson, D., Raturi, A., Richardson, D., Saphores, J-D. and Stokols, D. Toward alternative decentralized infrastructures. In *Proceedings of the 2015 Annual Symposium on Computing for Development*. ACM, 33–40.
37. Vardi, M. The financial meltdown and computing. *Commun. ACM 52*, 9 (Sept. 2009), 5.
38. Vitousek, P., Mooney, H., Lubchenco, J. and Melillo, J. Human domination of Earth's ecosystems. *Science 277* (1997), 494–499.
39. Weiser, M. The computer for the 21st century. *Scientific American 265*, 3 (1991), 94–104.
40. Wong, J. Prepare for descent: Interaction design in our new future. In *Proceedings of the 2009 CHI Workshop on Defining the Role of HCI in the Challenges of Sustainability*.

Bonnie Nardi (nardi@ics.uci.edu) is a professor in the Department of Informatics at University of California, Irvine, USA.

Bill Tomlinson (bill.tomlinson@vuw.ac.nz) is a professor in the Department of Informatics at the University of California, Irvine, USA, and an adjunct professor in the School of Information Management, Victoria University of Wellington, New Zealand.

Donald J. Patterson (dpatterson@westmont.edu) is a professor in the Department of Math and Computer Science at Westmont College, Santa Barbara, CA, USA.

Jay Chen (jay.chen@nyu.edu) is an assistant professor in the Department of Computer Science at NYU Abu Dhabi, U.A.E.

Daniel Pargman (pargman@kth.se) is an associate professor in the Department of Media Technology and Interaction Design at KTH Royal Institute of Technology, Stockholm, Sweden.

Barath Raghavan (barath.raghavan@usc.edu) is an assistant professor of computer science at the University of Southern California, Los Angeles, USA.

Birgit Penzenstadler (Birgit.Penzenstadler@csulb.edu) is an assistant professor in the Department of Computer Engineering and Computer Science at California State University, Long Beach, USA.