

# Networking in the Long Emergency

Barath Raghavan and Justin Ma  
ICSI and UC Berkeley

## ABSTRACT

We explore responses to a scenario in which the severity of a permanent energy crisis fundamentally limits our ability to maintain the current-day Internet architecture. In this paper, we review why this scenario—whose vague outline is known to many but whose consequences are generally understood only by the scientists who study it—is likely, and articulate the specific impacts that it would have on network infrastructure and networking research. In light of this, we propose a concrete research agenda to address the networking needs of an energy-deprived society.

## Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design

## General Terms

Design, Economics, Management

## Keywords

Internet Architecture, Energy, Resource Limits

## 1. INTRODUCTION

*“If I had my finger on the switch, I’d keep the juice flowing to the Internet even if I had to turn off everything else...The Net is the one solvent we can still afford; jet travel can’t be our salvation in an age of climate shock and dwindling oil, so the kind of trip you can take with the click of a mouse will have to substitute.”* And thus eminent science journalist Bill McKibben sums up his sobering survey of our current energy and climate predicament and identifies the importance of the Internet to our future [29]. We think he’s right, and present in this paper a challenge to the research community that may be among the hardest we have faced.

Computer networking is rightly a linchpin of modern society—it has given us the ability to exchange information at a speed and

scale that is unprecedented in human history. However, networking is the product of an energy-intensive economy and industrial system, and if that economic system faces a sudden and permanent energy shock, then networking must change to adapt. If we wish to preserve the benefits of the Internet, then a crucial opportunity emerges for networking researchers to reevaluate the overall research agenda to address the needs of an energy-deprived society. Due to the near-term depletion of oil this scenario is not hypothetical; it marks the beginning of an era of contraction.

Oil is the foundation of our industrial system—its energy enables the creation and transportation of goods and services that were unimaginable a century ago. Unfortunately, we face worldwide oil depletion—an era when oil production declines. Although it is unknown precisely when this era will begin, recent studies suggest it will be underway soon, as we survey later. A study commissioned in 2005 by the U.S. Department of Energy on the peaking of world oil production found the following [23]:

“The problems associated with world oil production peaking will not be temporary, and past ‘energy crisis’ experience will provide relatively little guidance. The challenge of oil peaking deserves immediate, serious attention, if risks are to be fully understood and mitigation begun on a timely basis.”

A severe, permanent energy crisis would have far-reaching consequences.<sup>1</sup> Though we are far from the first to consider such a scenario, we are unaware of prior work that considers the implications that this challenge will have for networking or computing generally. In this paper we issue a call to re-evaluate the broader networking research agenda to face this unprecedented predicament.

At first glance, the connection between the global energy system—oil in particular—and networking research may be unclear. However, our oil-depleted, post-peak future will bring about energy and financial pressures that will make developing and maintaining network infrastructure much more difficult. We explore this connection and touch on the direct and indirect impacts on networking infrastructure and networking research as this foundation of oil crumbles, consider practical constraints that we are likely to face, suggest design principles that might aid in future research, and propose a small list of research questions for the community to tackle. We begin our discussion with a brief overview of near-term global energy constraints.

<sup>1</sup>In this paper we primarily consider impacts on industrial nations, though any society that uses oil, or interacts with those that do, will be affected in some way.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*GreenNet’11*, August 19, 2011, Toronto, Ontario, Canada.  
Copyright 2011 ACM 978-1-4503-0799-4/11/08 ...\$10.00.

## 2. BACKGROUND

Our society faces long-term energy scarcity. Although the exact timing is uncertain, the emerging consensus is that this transition will begin soon. This section provides a brief overview of the coming energy challenges, of why existing alternatives are inadequate substitutes, and of economic impacts. While modern networking exists in the zeitgeist as the seemingly abstract movement of information, we cannot ignore the depletion of physical and financial resources that make it possible. Although we focus on oil and energy due to Liebig's Law<sup>2</sup>, the global limits we face are all-encompassing [30] and thus our survey of these issues, and of the post-peak era sometimes called the Long Emergency, is necessarily truncated; we refer interested readers to the many works on the subject [10, 12, 13, 18, 19, 21, 23, 26, 29, 30, 35, 36, 38].

### 2.1 Fossil Fuel Dependence

Industrialized economies require oil for vital functions such as transportation, manufacturing, and agriculture; they similarly depend on coal and natural gas for the production of electricity.<sup>3</sup> Combined, these three fossil fuels account for 81% of global energy consumption [1], and oil is the plurality of that share. Since oil is the most malleable and energy-dense of these fossil fuels, it is also the most highly prized (especially for transportation, which consumes 61% of global oil production). Oil is also an essential input in the extraction, production, and transportation of other energy resources, from coal to natural gas to alternative technologies.

Because fossil fuels are non-renewable, there is a point at which the rate of resource extraction reaches a peak and begins to decline. A persistent decline in annual energy yields can significantly disrupt the economy. Thus a pressing matter is the question of timing—when do we expect to face energy constraints?

### 2.2 Near-Term Energy Shortage

In 2005 the U.S. Department of Energy commissioned a study known as the Hirsch report, which analyzed the timing, consequences, and mitigation of a peak and subsequent decline in world oil production [23]. Nine out of twelve of its forecasts predicted the date of peak oil production to occur by 2016. Similar studies were commissioned last year, projecting peak production by 2015: a consortium of British corporations (peak: 2014–15) [38], the German military (peak: 2010) [36], Kuwait University (peak: 2014) [31], the U.S. Defense Department (peak: 2012) [2], and Lloyd's of London (peak: 2013) [12]. The peak may even be behind us—the International Energy Agency suggests that conventional oil production may have peaked in 2006, though they anticipate no peak in all-liquids production [3].

None of these studies claims to be or can in fact be definitive—oil reserves and production capacity are closely guarded secrets in many countries, and a production peak can only be confirmed years after the fact because of year-to-year production variability. Also, as Heinberg notes, disagreement regarding the date of peak production is due to varied definitions of “oil”, ranging from solely conventional crude to all liquid fuels; his study in 2003 concluded that an all-liquids peak would occur between 2006–2015 [18]. These studies represent a wide range of expert analysis that suggests that we will soon live in a world that has less oil available each year than the year before.

<sup>2</sup>Liebig's Law of the Minimum states that growth is limited not by total resources but by the most constrained resource [8].

<sup>3</sup>Nuclear provides a select few nations most of their electricity.

### 2.3 Inadequate Alternatives

Replacing oil is extremely difficult. Industrially feasible alternatives are environmentally-unfriendly whereas environmentally-friendly alternatives are industrially infeasible. Also, the scale of the transition requires a multi-decade effort, yet energy companies may be reluctant to immediately convert their current fossil-fuel infrastructure, which represents a \$10 trillion investment that will require 10–50 years of continued operation to recover [29].

The Hirsch report surveyed a number of responses to the challenge of oil depletion: enhanced oil recovery, heavy oil / tar sands production, coal liquefaction, gas-to-liquids, and fuel-efficient vehicles; the report shows that such a successful mitigation response would not only have to begin at least *20 years before* the peak date, but also would have to be a *crash program*—one that takes on singular national importance and focus. Moreover, the report did not consider the environmental or climate impact of the above mitigation approaches. Because coal, heavy oil, and tar sands are highly polluting sources, increasing their use would have significant, negative consequences [16].

In a broader analysis, Heinberg evaluates the full spectrum of energy alternatives: he concludes that while it is possible that some alternatives may be viable in the future, they cannot easily replace oil in its role in transportation or agriculture [18]. Ignoring substitutability, Griffith constructs a hypothetical global energy profile (for the 2030s) using an array of alternatives [14]; however, for these alternatives (solar, wind, nuclear, etc.) current industrial capacity is inadequate and would require a WWII-like manufacturing effort. For example, we would have to construct nuclear power plants at roughly *5 times* the peak rate ever achieved globally—and sustain it for 20 years—to supply just the nuclear contribution to the alternative energy profile [14, 28].

While we can imagine optimistic future energy profiles, there exist no practical (i.e., within even generous assumptions about industrial capacity) and acceptable (i.e., no worse than oil environmentally) energy alternatives to address a *near-term peak* in oil production. While we believe a multi-decade transition to alternatives may be achievable, we expect that it will be extraordinarily difficult in the face of the economic consequences of peak oil.

### 2.4 Economic Consequences

The predominant way that oil depletion is expected to impact developed nations is via oil price shocks and consequent economic hardship. We omit a detailed survey of possible economic impacts in this paper, but note that most conventional economic analyses ignore the issue of oil depletion. It is however well established that high oil prices have been a factor in most recent recessions in the United States [15]. In addition, the Hirsch report provides a warning of the economic consequences of peak oil production, and that only via timely mitigation (at least 20 years before peak in a crash program) can these consequences be avoided [23]:

“The long-run impact of sustained, significantly increased oil prices associated with oil peaking will be severe. Virtually certain are increases in inflation and unemployment, declines in the output of goods and services, and a degradation of living standards. Without timely mitigation, the long-run impact on the developed economies will almost certainly be extremely damaging, while many developing nations will likely be even worse off.”

Said another way [18]:

“Energy scarcity will cause a recession of a new kind—one from which anything other than a temporary, partial recovery will be impossible.”

This will mark a new era for industrial society, and we believe computer networking may be eclipsed by more pressing needs such as food and transportation. Nevertheless, like McKibben we believe that the Internet is a key resource for the post-peak future, and therefore it is important for us to consider how to respond to this predicament.

### 3. NETWORKING RECONSIDERED

Although the challenges we face are dire, they also offer tremendous intellectual and practical prospects. In this section we consider what the future of networking might look like and how an energy descent will change the way we navigate systems constraints. After establishing this baseline panorama, we propose a set of guiding principles that may be valuable when developing research responses to these new challenges. Finally, we pose a series of concrete open questions on research challenges in post-peak networking that we believe are well worth answering. We hope this provides a starting point for discussion so that our community will consider future scenarios that differ markedly from conventional, optimistic projections of increasingly plentiful bandwidth, powerful cloud platforms, and rapid development and adoption of new technology. We begin by establishing several key premises.

#### 3.1 Premises

**Non-uniform / volatile descent.** We expect that energy and financial constraints will affect different regions and nations at different times and with different degrees of severity, and we expect that both resource availability and prices are likely to exhibit dramatic volatility in the years ahead. As part of this variability, we expect the economy to proceed in cycles of contraction and partial recovery, with an overall downward trajectory [13, 18, 22, 23].

**Economic impacts.** We expect that economic decline (rather than direct impacts of energy scarcity) will be the primary challenge for networking in the initial post-peak years [18, 22, 23].

**Liquid fuels, then electricity.** We expect that the initial true energy constraint will be in liquid fuels and not in electricity [23], but that over time constraints and challenges will become more apparent with electricity and the grid [18].

**Stalling trends.** Much of systems and networking research depends upon the ability to project into the future and meet that future with a research solution. These projections often depend on a) the extension of current trends into the future, b) the expectation that the future will resemble the present, c) the analysis of technology stair-steps that are expected, or some combination of these approaches. We expect that such conventional projections are likely to be overly optimistic in a post-peak world; we expect current technology trends to slow or stall [30].

**Relocalization.** Those who have studied this predicament have by and large concluded that societal relocalization is a crucial (though imperfect) response [18, 26, 29, 35], since the cost of moving goods long distances will become prohibitive. We expect that relocalization of social and governing structures will begin to occur, though it may be slow at first.

**Shrinking user bases.** We expect the societal discontinuity we face—from growth to contraction—to eventually cause a decrease in the overall use of computing to meet societal needs. That is, low-energy or low-complexity alternatives may come to replace current technology [13, 18, 26, 29], though this process may take decades to even begin. Thus Metcalfe’s Law will be forced to run in reverse: as user bases shrink, the value of certain technologies or systems may fall quite quickly.

#### 3.2 A Future Scenario

To provide a concrete basis for discussion, we begin with a possible glimpse into our collective future. To ground this forecast, we first look to oil production. In the coming years, it appears that the best case for annual post-peak oil production decline is as follows: 1% for all-liquids production, 2% for crude production, and 4% for available net exports of crude [5, 9, 22]. Using Hirsch’s estimate of a one-to-one correspondence between production decline and GDP decline, this roughly yields a *best-case of 1-4% annual global economic decline* as long as no crash program is pursued [22].<sup>4</sup> A decline of 1-4% may not seem significant—it is on the order of an ordinary recession—but a consistent, long term decline in that range yields a best-case 5-19% decline over 5 years, 10-34% decline over 10 years, and 19-56% decline over 20 years.

Given the number of different complicating factors that will interact in the coming post-peak years—from agriculture to climate to finance to geopolitics—it is extremely difficult to make projections of any sort. Nevertheless, we attempt to provide one here based upon simple analyses of the impacts of past declines or shortfalls in oil production and compounding factors [15, 22]. In Table 1 we translate the decline rate into quantitative guesstimates of costs and other constraints for dates 5, 10, and 20 years post-peak.<sup>5</sup>

Since small decreases in the supply of commodities such as oil can yield large increases in prices, we expect that prices for transportation and shipping will go up significantly. Eventually we expect that electricity prices will follow due to the falling net energy of present and future sources and the increased input cost of oil in their production.<sup>6</sup> We present our projections as ranges since we expect there to be significant variation between regions and over time. Thus in our view, the challenge for systems and networking researchers is in coping with the full range of these radically new constraints, though in the near term they will first be felt by those who operate today’s networks.

The constant upgrading and maintenance of the Internet and network infrastructure today is not a cheap undertaking: telecoms spent about \$3 trillion over the last decade on capital expenditures [24]. Thus some operations may be downscaled or canceled for lack of funds; we suspect that R&D and maintenance may be the first to go, as the former requires the promise of future financial growth and the latter requires the constant use of expensive

<sup>4</sup>Hirsch’s best-case scenario projection is a slightly more pessimistic 2-5% annual decline [22]. Also, as we stated earlier, we expect the decline to be uneven across regions and time, with periods of partial recovery.

<sup>5</sup>Our guesstimates are roughly in line with similar analyses such as that of the recent IMF World Economic Outlook, which projects that a 2% annual decline in oil production would result in a price increase of 8x over 20 years (not considering declining purchasing power or other factors) [4].

<sup>6</sup>For example, we may be at peak energy from coal [34] or soon at peak cheap coal [20]; also, other analyses put all future potential sources at far below the net energy of past sources [19].

Category	Years Post-Peak		
	5	10	20
Transportation	3-5x	5-15x	10-25x
Electricity	2-4x	2-10x	5-20x
Grid Reliability	98-99%	95-99%	75-98%
Discretionary Income	5-15%	2-10%	1-5%

**Table 1: A set of possible post-peak constraints for industrial nations given a 1-4% annual decline and no crash program. Transportation and electricity are expressed as costs relative to peak in terms of purchasing power, grid reliability is in uptime, and discretionary income is in % of median household budget.**

resources. Just as recent studies have indicated that the aviation industry is likely to discontinue service to smaller, rural airports in response to energy descent [29], we expect telecoms will have a financial incentive to focus on their networks in major cities at the expense of less populated regions. This will likely have a downstream effect on hardware vendors, who rely on network infrastructure operators to regularly purchase new equipment and service contracts. Since we expect end users to have less discretionary income, network services and device vendors that depend upon consumer spending will face threats to their existence as the ubiquity of networking declines.

### 3.3 Principles

To design for a post-peak world, we will have to revisit the doctrines that guide present-day networking. Below, we discuss six principles that we believe are worth considering, all of which are already known and applied in specific contexts.

**P1. Target Absolute (not Relative) Consumption.** For an energy descent to be controlled and intelligent rather than uncontrolled and haphazard, our energy constraint must be self-imposed rather than exogenous [13, 19, 29]. Thus rather than aiming to improve status quo energy consumption by some relative fraction, resource-efficient networking should aim to work within a specific energy and resource budget. By targeting an absolute consumption goal—which may be orders of magnitude lower than current-day consumption—it may be necessary to discard components, services, etc. that cannot be sustained within that constraint.

**P2. Account for All Energy Inputs.** A common myth is that technological innovation alone can solve our energy challenges—this ignores that all technology depends upon and is distinct from energy. Fundamentally, any energy technology must derive its energy from one of the sources available on Earth. All new energy technologies require some resources and some energy expenditure to construct (e.g. the energy and resources to build a power plant or a solar panel), and quantifying this expenditure is crucial—this is captured in the well-known concepts of Energy Returned on Investment (EROI) and net energy [19]. Embodied energy (*emergy*) or life-cycle energy—the amount of energy used to mine, assemble, transport, use, dispose, etc. a system or device—is similar to EROI and applies to all systems, including networking technology [32].<sup>7</sup> In particular, we should not focus exclusively on the operational power cost of a networked system. There is no doubt

<sup>7</sup>As Odum noted in his classic work, “The natural conversion of sunlight to electric charge that occurs in all green-plant photosynthesis after 1 billion years of natural selection may already be the highest net *emergy* possible.” [32]

that recursively accounting for all energy inputs for a system (from manufacture, to installation, to operation, to disposal, etc.) will be challenging. However, accounting for these inputs will be important for determining whether existing and future networking and energy technologies are truly “green” or whether they use more energy to create than they will save.

**P3. Reuse Hardware and Software.** In a scenario where energy and other resources are scarce, creating new energy-efficient network devices to replace old network devices and/or their software potentially increases the energy deficit—the embodied energy costs of new devices would accumulate quickly. Thus we must be willing to reuse existing hardware and software. A future Internet may consist of partially-disconnected networks cobbled together using a range of different protocol stacks that suit each local region. As a result, we may have to account for a higher degree of hardware and software diversity than for present-day standards, which assume that perpetual, periodic upgrades are the norm.

**P4. Design Resilient Systems.** An old urban legend is that the Internet was designed to survive a nuclear war. However, the story is apocryphal and we have in fact a limited understanding about the resilience of the Internet. Today’s Internet has grown organically; the evolution of its interconnection has been a disorderly process and the complete topology is unknown. Beyond the basic structure of the network, years of research have suggested that it is important to design robust and resilient network protocols [6]. Designing and studying resilient protocols and systems will be especially important since economic decline, energy shortages, and an erratic climate can only make network infrastructure less reliable.

**P5. Become Multidisciplinary.** To conduct networking research in the context of an energy-deprived society, we must gain a deeper understanding of the environment and the constraints. For example, if we want to assess the energy inputs and resilience of a networked system (including its materials, assembly, transportation, maintenance, impacts, etc.), then collaborating with geologists, mechanical and civil engineers, ecologists, and others will be crucial. It will become increasingly difficult to think about networking in isolation because the multi-faceted nature of the post-peak scenario will no doubt impact networking.

**P6. Build Self-Sustaining Systems.** Current networks are composed of a complex array of hardware and software assembled around the world with materials, energy, skills, and designs also from a global resource base. We expect that today’s approach to the design, creation, and deployment of networking technology is likely, in time, to become too costly or simply physically infeasible. Thus networking technology should follow the principles of Appropriate Technology [17, 37]: be designed to be a) simple, b) locally reproducible, c) composed of local materials / resources, d) easily repairable, e) affordable, and f) easily recyclable.

### 3.4 A Research Agenda

Next we propose a research agenda of open questions that range in flavor from specific to broad. The list of questions is far from exhaustive, and the questions themselves may require significant refinement; we offer them as concrete starting points.

#### 3.4.1 Network Structure

While today’s networks and services are generally oligarchic in nature, we expect that a post-peak Internet will need to distribute all of its functions in the service of relocalization. The “core” may

cease to be a clear group of providers, and the hierarchy might break down. In this model, services and control might be pushed further to what we think of today as the edge of the network.

**Open Question 1.** *As networks become more localized, the cost and latency of communicating with far away nodes will be higher than it is today. How will we cope with this?*

**Open Question 2.** *How might we carefully guide this structural transition (transferring management from the core to the edges), instead of allowing it to descend haphazardly?*

**Open Question 3.** *What do standards look like post-peak? What role do standards bodies such as ICANN and IETF play?*

**Open Question 4.** *What cost sharing mechanisms can be feasibly deployed to offload a substantial portion of the true cost of a network service onto its user?*

**Open Question 5.** *What does the programming model for a fully-distributed datacenter-less cloud look like?*

### 3.4.2 Reevaluation

We are moving into the post-peak future with a huge existing base of networking hardware, software, protocols, and metrics. We must begin to reevaluate networks, networking components, and network research in a new light, with changed constraints, goals, metrics, and circumstances—in particular, accounting for the emergy of devices, as well as coping with greater network volatility. In doing so, we might find that some protocols and systems can readily be adapted or salvaged, while others must be discarded.

**Open Question 6.** *Can we develop a common methodology for calculating the emergy of a network device? Can hardware vendors develop a standard for describing how much energy went into producing a device?*

**Open Question 7.** *Can we measure which existing projects in energy-efficient networking are well-suited to the post-peak world and which are not?*

**Open Question 8.** *When do free network services become infeasible due to energy costs? In other words, at what price per KWh (and auxiliary costs in real terms) will Google no longer be able to offer free searches?*

**Open Question 9.** *How can network protocols be best redesigned to cope with post-peak volatility?*

**Open Question 10.** *How can existing software implementations of network protocols be re-purposed without modification?*

**Open Question 11.** *When is it the case that software upgrades—while using old hardware—are preferable to upgrading to a more resource-efficient hardware platform?*

**Open Question 12.** *Will wireless deployments continue to be a good bet? Or will the emergy and power cost of wireless devices and infrastructure eventually make them cost prohibitive vs. reusing old wire-line infrastructure?*

**Open Question 13.** *Which services provided by the Internet are least essential to its continued value and operation? Which are most essential?*

### 3.4.3 Integration

Coping with the complexities and vicissitudes of a post-peak world is likely to require knowledge from more than one domain or field of study. Old models and expertise in one area may apply to the new circumstances in another. Systems may need to be more than the sum of their parts for their continued survival.

**Open Question 14.** *Can we encourage more video conferencing adoption by incorporating computer vision or similar techniques into video streaming protocols to augment the video?*

**Open Question 15.** *Can computer network protocols and algorithms be applied to transportation networks (or vice versa) so as to improve their overall efficiency?*

**Open Question 16.** *Using today's architecture, how can we enable and promote a systematic way of leveraging cross-layer and network-internal knowledge at end points?*

**Open Question 17.** *What are the economic models for a demand-based pricing system for a post-peak Internet?*

**Open Question 18.** *How will the economics of network misbehavior (spam, DoS, malware, etc.) change post-peak?*

**Open Question 19.** *How can a secure, peer-to-peer localized microlending system be built?*

### 3.4.4 Components and Tools

Ultimately, the network is only going to be of value to post-peak society if it meets basic needs. For this reason, specifically engineering components and tools for new constraints will be of great value.

**Open Question 20.** *What is the minimal set of mining facilities, hardware manufacturing, software tools, etc., all capable of self-bootstrapping, that can provide all the needed components of the Internet? How small and cheaply can it be made?*

**Open Question 21.** *How can network switches and routers be built to passively perform forwarding?*

**Open Question 22.** *How might technology costs and energy trends change with respect to in-network storage, and when will it become unviable?*

**Open Question 23.** *How can a long-term network-attached data archival service be designed to provide persistence and proof of storage?*

**Open Question 24.** *Can we develop a "currency" for local network bandwidth sharing?*

**Open Question 25.** *How might networks leverage older, simpler technologies such as packet radio?*

## 4. RELATED WORK

In the last decade, and especially in the last few years, networking researchers have turned their attention to improving the energy efficiency of network devices, protocols, and systems and to reevaluating and integrating prior approaches. Some recent efforts have looked at large-scale challenges, such as applying networking ideas to the design of smart grids [25, 27]. We believe that these strands of research are valuable—and too broad to itemize here—and can be repurposed to target the future we describe in this paper.

The work that may in fact be closest in spirit to what we describe herein is networking research for developing regions [7, 33]. This body of work has typically encompassed low-cost, low-energy, and low-infrastructure solutions to networking problems. While the constraints in developing-regions research are somewhat different—the regions in question typically have little existing infrastructure or computing resources and likewise have little dependence upon them—once again its techniques can likely be repurposed to meet these new challenges. However, for bootstrapping—manufacturing, deployment, and repair—such work typically relies upon the globalized, fossil-fuel based industrial system.

More broadly, research into resilient network models such as delay-tolerant networking [11] and older network technologies such as packet radio are highly appropriate for further exploration as approaches for post-peak networking. The key properties that these and similar approaches possess is a key ability to function in a wide range of operating conditions during an irregular post-peak decline over time and across space.

## 5. CONCLUSIONS

“And so we turn to the essentials of our future. In order: food, energy, and—yes—the Internet” says McKibben [29]. Modern society has never before faced a predicament like this. Experts in food and energy and myriad other areas are beginning to take appropriate action and we believe that it is incumbent upon us, as researchers in one of these most crucial areas of study, to do our part.

## 6. ACKNOWLEDGEMENTS

We thank Scott Shenker, John Michael Greer, Harsha Madhyastha, Stuart Staniford, Jenny Trujillo, and the anonymous reviewers for their encouragement and thoughtful feedback.

## 7. REFERENCES

- [1] Key world energy statistics. *International Energy Agency*, 2010.
- [2] The Joint Operating Environment. *U.S. Joint Forces Command*, 2010.
- [3] World energy outlook. *International Energy Agency*, 2010.
- [4] World economic outlook. *International Monetary Fund*, 2011.
- [5] K. Aleklett, M. Höök, K. Jakobsson, M. Lardelli, S. Snowden, and B. Söderbergh. The Peak of the Oil Age-Analyzing the world oil production Reference Scenario in World Energy Outlook 2008. *Energy Policy*, 38(3), 2010.
- [6] T. Anderson, S. Shenker, I. Stoica, and D. Wetherall. Design guidelines for robust Internet protocols. In *Proceedings of HotNets*, 2002.
- [7] E. Brewer, M. Demmer, B. Du, M. Ho, M. Kam, S. Nedeveschi, J. Pal, R. Patra, S. Surana, and K. Fall. The case for technology in developing regions. *IEEE Computer*, 38(6), 2005.
- [8] R. Brewer. *Principles of ecology*. Saunders, 1979.
- [9] J. J. Brown and S. Foucher. Egypt, a Classic Case of Rapid Net-Export Decline and a Look at Global Net Exports. *ASPO-USA*, Feb. 2011.
- [10] K. Deffeyes. *Beyond oil: The view from Hubbert’s Peak*. Farrar, Straus, and Giroux, 2005.
- [11] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of ACM SIGCOMM*, 2003.
- [12] A. Froggatt and G. Lahn. Sustainable energy security: Strategic risks and opportunities for business. *Chatham House-Lloyd’s 360 Risk Insight White Paper*, 2010.
- [13] J. M. Greer. *The Long Descent*. New Society Publishers, 2008.
- [14] S. Griffith. Climate change recalculated. In *The Long Now Foundation Seminars*, 2009.
- [15] J. Hamilton. Causes and Consequences of the Oil Shock of 2007–08. *Brookings Papers on Economic Activity*, 2009.
- [16] J. Hansen. The threat to the planet. *The New York Review of Books*, Jul. 13, 2006.
- [17] B. Hazeltine and C. Bull. *Appropriate Technology; Tools, Choices, and Implications*. Academic Press, 1998.
- [18] R. Heinberg. *The Party’s Over*. New Society Publishers, 2003.
- [19] R. Heinberg. Searching for a Miracle: “Net Energy” Limits and the Fate of Industrial Society. *The International Forum on Globalization and the Post Carbon Institute*, 2009.
- [20] R. Heinberg and D. Fridley. The end of cheap coal. *Nature*, 468(7322), 2010.
- [21] R. Heinberg and D. Lerch. *The Post Carbon Reader: Managing the 21st Century’s Sustainability Crises*. University of California Press, 2010.
- [22] R. Hirsch. Mitigation of maximum world oil production: Shortage scenarios. *Energy Policy*, 36(2), 2008.
- [23] R. Hirsch, R. Bezdek, and R. Wendling. Peaking of World Oil Production: Impacts, Mitigation, & Risk Management. *U.S. Department of Energy NETL*, 2005.
- [24] Infonetics Research. <http://www.infonetics.com/>.
- [25] S. Keshav and C. Rosenberg. How internet concepts and technologies can help green and smarten the electrical grid. In *Proceedings of the ACM SIGCOMM Workshop on Green Networking*, 2010.
- [26] J. Kunstler. *The Long Emergency*. Grove Press, 2005.
- [27] LoCal Project. <http://local.cs.berkeley.edu/>.
- [28] D. MacKay. *Sustainable Energy — without the hot air*. UIT, 2009.
- [29] B. McKibben. *Eaarth: Making a life on a tough new planet*. Henry Holt and Company, 2010.
- [30] D. Meadows, J. Randers, and D. Meadows. *The limits to growth: the 30-year update*. Chelsea Green, 2004.
- [31] I. Nashawi, A. Malallah, and M. Al-Bisharah. Forecasting world crude oil production using multicyclic Hubbert model. *Energy & Fuels*, 24(3), 2010.
- [32] H. Odum. *Environmental accounting: emery and environmental decision making*. John Wiley and Sons, 1996.
- [33] T. Parikh and E. Lazowska. Designing an architecture for delivering mobile information services to the rural developing world. In *Proceedings of ACM WWW*, 2006.
- [34] T. Patzek and G. Croft. A global coal production forecast with multi-Hubbert cycle analysis. *Energy*, 35(8), 2010.
- [35] J. Rubin. *Why your world is about to get a whole lot smaller*. Random House of Canada, 2009.
- [36] S. Schultz. Military study warns of a potentially drastic oil crisis. *Der Spiegel*, Sep. 1, 2010.
- [37] E. Schumacher. *Small is Beautiful: A study of economics as if people mattered*. Abacus, 1974.
- [38] W. Whitehorn, S. Roberts, J. Miles, S. Behling, I. Marchant, J. Chandler, J. Leggett, S. Stewart, N. Fox, and A. Knight. The Oil Crunch - a wake-up call for the UK economy. *Second report of the UK Industry Taskforce on Peak Oil and Energy Security*, 2010.