

# Taming Limits with Approximate Networking

Junaid Qadir  
Information Technology  
University (ITU)-Punjab,  
Lahore, Pakistan  
junaid.qadir@itu.edu.pk

Liang Wang  
Computer Laboratory,  
University of Cambridge,  
United Kingdom  
liang.wang@cl.cam.ac.uk

Arjuna Sathiaseelan  
Computer Laboratory,  
University of Cambridge,  
United Kingdom  
arjuna.sathiaseelan@cl.cam.ac.uk

Jon Crowcroft  
Computer Laboratory,  
University of Cambridge,  
United Kingdom  
jon.crowcroft@cl.cam.ac.uk

## ABSTRACT

Internet is the linchpin of modern society, which the various threads of modern life weave around. But being a part of the bigger energy-guzzling industrial economy, it is vulnerable to disruption. It is widely believed that our society is exhausting its vital resources to meet our energy requirements, and the cheap fossil fuel fiesta will soon abate as we cross the tipping point of global oil production. We will then enter the long arc of scarcity, constraints, and limits—a post-peak “long emergency” that may subsist for a long time. To avoid the collapse of the networking ecosystem in this long emergency, it is imperative that we start thinking about how networking should adapt to these adverse “undeveloping” societal conditions. We propose using the idea of “*approximate networking*”—which will provide *good-enough* networking services by employing *contextually-appropriate* tradeoffs—to survive, or even thrive, in the conditions of scarcity and limits.

## CCS Concepts

•Networks → Network design principles;

## Keywords

Approximate Networking; Computing Under Limits; ICTD

## 1. INTRODUCTION

Human beings today enjoy a standard of living unparalleled in human history. This age of abundance is driven largely by technology and in particular through information and communication technology (ICT). The Internet—which has impacted all facets of human life (business, governance, education, leisure) through its ability to connect people and facilitate communication—is widely believed to be a gateway to human prosperity and opportunities.

But due to the non-sustainable<sup>1</sup> trajectory adopted by the indus-

<sup>1</sup>The sustainability of a system such as the Internet measures its

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

LIMITS '16, June 08-10, 2016, Irvine, CA, USA

© 2016 ACM. ISBN 978-1-4503-4260-5/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2926676.2926678>

trialized civilization, the ICT-enabled societal progress has been achieved at a big cost. The fundamental physical limits of a finite world are being taxed by, and will inevitably stall, the exponential trends displayed in human/ society demographics (e.g., population) and ICT (e.g., Moore’s law, big data). It is now widely believed that modern civilization is close to exhausting non-trivial resource limits (such as the non-renewable fossil fuels [1]). As this depletion will start to show its effects in the post-peak environment, even developed countries will face degradation of their social and economic systems due to societal collapse (and will “undevelop” [2]—i.e., despite having some infrastructure, these countries will regress and will become economically and politically unstable).

### 1.1 Post-Peak Future of Limits and Scarcity

A characteristic reason for the “*collapse*”<sup>2</sup> of a society is that its citizens commit “*ecocide*”—i.e., they overuse and exhaust their vital resources [3]. Through an over-reliance on non-renewable fossil fuels, the industrialized world is also risking a future collapse. It is estimated that fossil fuels comprising oil, coal, and natural gas together account for approximately 80% of global energy consumption, with oil servicing the major bulk [4]. Many experts are predicting that the production of oil—which is the foundation of the industrial system—will deplete in the near term [5]. After the peaking of the world’s oil production, the decline of oil production will likely lead to a “long emergency” whose effects will be wide reaching and long lasting (that may last for decades).

A US government commissioned study (called the Hirsch report) attempted to analyze the timing and consequences of oil peaking—the results indicated that peak oil production would occur by 2016. Similar studies elsewhere indicate that this peak may already be behind us—although oil reserves and production capacity are often closely guarded secrets, and the collapse consequences of the post-peak may take some time to manifest itself, some of the signs (e.g., global climatic change, economic slowdown) have already started to emerge. Since fossil-fuels are non-renewable (i.e., they exist in finite non-replenishable quantity), a point will inevitably come where the rate of resource extraction peaks.<sup>3</sup> While there is some ongoing work on using renewable energy for networking [6],

capacity to endure diverse exogenous pressures (such as resource depletions) while remaining productive indefinitely.

<sup>2</sup>Collapse happens over a long time—decades or sometimes even centuries and should not be confused with an apocalypse that occur suddenly and instantaneously.

<sup>3</sup>Some experts indicate that we may already be in this post-peak era [4].

it is doubtful, extrapolating current trends, that we will be able to match our inflated requirements with renewable energy in the short term. Such a likely persistent decline in energy will provide a permanent shock to the energy-guzzling industrial ecosystem which would likely lead to a societal collapse.

It is instructive to note some other reasons for societal collapse noted in literature [3], such as: (1) reliance on trading partners; (2) self-inflicted environmental damage; and (3) inflexibility of institutions when change is needed. The latter two reasons are directly relevant to our subject topic. The first reason is also related but indirectly: it shows how dependence on entities can make systems less resilient to disruptions. To be sure, the emergence of an *oil-depleted and post-peak future poses fundamental constraints and limits on the Internet architecture and infrastructure*.

## 1.2 Networking in the Long Emergency

To cope with the impending sudden and potentially long-lasting energy shock, it is necessary for the networking architecture to adapt. The key questions in the networking context are:

1. *how should we adapt Internet technology so that it becomes sustainable?;*
2. *how will the Internet users and applications fare if the Internet turns out not to be sustainable?.*

The answers to these questions can motivate the adaptation of our network architecture and applications so that they become “collapse compliant” [7].

Researchers have only started to look at how should networking adapt to deal with a “long emergency” that can arise in an undeveloping environment. In the seminal paper that looked at networking issues for the long emergency [4], a number of premises or assumptions were stated. It was assumed that energy and financial constraints will be non-uniform in different regions and will iterate between contraction/partial recovery but with an overall downward trend. It was also assumed that economic decline will also be a major challenge for networking (apart from the direct impacts of energy scarcity). It was also assumed that user bases will likely shrink as there may be a decrease in the overall use of computing to meet societal needs (and non-digital alternatives may be increasingly deployed).

## 1.3 The Approximation Tradeoff

*“Although this may seem a paradox, all exact science is dominated by the idea of approximation.”—Bertrand Russell*

We necessarily employ approximation in many technologies and sciences. We use approximation in measurement and in digital computing. We use approximation when the problem is too intractable to solve optimally: in such cases, we lower down our targets to *satisficing* (i.e., producing “good enough” answers) rather than optimizing [8].

While ideally speaking, we will like an Internet that is perfect, and has extremely high capacity, bandwidth, and reliability in addition to extremely low or negligible delays, errors, and congestion.<sup>4</sup> We call such networks “ideal networks”. In contrast, we consider “approximate networks” that are networks that make some design tradeoffs to deal with varying levels of challenges and impairments. We note here that ideal networks and approximate networks do not

<sup>4</sup>For practical purposes, the modern fiber-based broadband high-speed networks available in select places (mostly in advanced countries) come close to this ideal.

define a binary divide but a spectrum of options. We can also define approximate networks as networks that come close to ideal networks in quality, nature, and quantity<sup>5</sup>. We need “approximate networking” when the imperfections of the real world preclude an “ideal networking” solution. In particular, an approximate networking solution is appropriate when any of the ideal networking assumptions—e.g., that there is 24x7 connectivity; an end-to-end path is always available; the end-to-end delay, and the link propagation delay, is never too high (i.e., is less than (half) a second); the networks should not be congested; the networks should not have high error rates—are not met.

Approximate networking is inspired in part from the emerging architectural trend of “approximate computing” [9] that leverages the capability of many computing systems and applications to tolerate some loss of quality and optimality by trading off “precision” for “efficiency” (typically energy efficiency). Approximate computing works by relaxing the convention of exact equivalence between specification and hardware implementation. By employing approximations at the hardware level, the energy efficiency of systems can be boosted significantly without too much loss in quality. As a concrete example, consider that it has been reported that using approximate computing techniques for the k-means clustering algorithm can provide an astounding 50x energy efficiency gain with an accuracy loss of only 5 percent [10].

Approximate networking, in a similar vein, enables a network architecture that allows networking protocols and applications to trade off service quality for efficiency in terms of cost/ affordability/ accessibility. Approximate networking is closely aligned to the philosophy of “appropriate technology” since it aims to match the user and the need in complexity and scale. Appropriate technologies are defined as “small scale, energy efficient, environmentally sound, labor-intensive, and controlled by the local community” [11]. Approximate networking also fits well as a collapse informatics solution since it can be used as a “tool for the study, design, and development of sociotechnical systems in the abundant present for use in a future of scarcity” [7].

## 1.4 Comparison of ICTD and LIMITS

While there is some overlap in the focus of ICT for development (ICTD) and the research on the problems of the undeveloping world with limits (we refer to this setting in this paper as LIMITS), these areas are fundamentally different. A number of standard ICTD assumptions [12] (e.g., Moore’s law will hold into the future; there will be increased access to capital and improved business environments; the use of widespread ICT is desirable<sup>6</sup>) may be invalid in the context of LIMITS. In addition, in contrast to the ICTD research that emphasizes a developing/developed world split, *the limits considered in our context are more global*. Thus it is anticipated that virtually all nations will experience a reduction in their material standards of living with developed countries also facing the crunch (perhaps even more so since they are more heavily reliant on ICT and infrastructure). The “undeveloping countries” focused on in the LIMITS context is a generalization of the narrow class of developing regions (focused mostly in ICTD research [12]) to also include developed countries that have some infrastructure but also a regressing economical and political climate.

For sure, there can be significant diffusion and reuse of ideas

<sup>5</sup>Oxford Dictionary: Approximate (v): come close or be similar to something in quality, nature, or quantity.

<sup>6</sup>In fact in the context of LIMITS, less ICT may be more desirable than more—e.g., previous research has focused on developing self-obviating systems designed to make themselves superfluous through their use so that ICT reliance is reduced [13].

across LIMITS and ICTD. LIMITS solutions can produce innovations that are broadly useful for localized collapse solutions, emergency response, and ICTD [7]. Similarly, the ICTD research focusing on dealing with scarcity will be valuable in the LIMITS context. The technique of approximate networking is relevant for both ICTD and LIMITS research. Approximate networking can be used in the LIMITS settings to wean our reliance on energy-inefficient overengineered “perfect” ICT while still attaining contextually appropriate service.

## 1.5 Contributions of this paper

The main contribution of this paper is to position approximate networking as a suitable framework for systematically thinking about networking tradeoffs that will inevitably arise in the era of post-peak world burdened with limits and societal collapse. The resource crunch in such an environment will necessitate a move away from overengineered “perfect products” towards *contextually-appropriate* “good enough” solutions. The challenge entailed in approximate networking is in determining what these context-appropriate network tradeoffs should be. Approximate networking is relevant both for ICTD research (that focuses on developing countries) and for research in the LIMITS context [2] [4] that focuses more broadly on “undeveloping countries” that have regressed due to global limits.

## 2. WHY USE APPROXIMATE NETWORKING UNDER LIMITS?

*“And so we turn to the essentials of our future. In order: food, energy, and—yes—the Internet”*—McKibben.

In the post-peak era of decline of the industrial society, a number of needs such as food and transportation will become prioritized, but Internet will also likely be a key resource for the post-peak future and thus developing solutions to retain its functionality (even if certain approximation tradeoffs are adopted) will be essential. In this era of scarcity, providing “ideal networking” service will not be economically feasible and some sort of tradeoff would be inevitable and an inescapable recourse for network designers. This is not new since we have learnt through decades of experience with the Internet that invariably there are Internet design tradeoffs and there is no single one-size-fits-all solution. Approximate networking is a useful way to think of tradeoffs that users and applications should consider for sustainable and collapse-compliant networking in the grim situation where we will run out of many of the essential resources (such as cheap energy through fossil fuels).

Some important reasons we should seriously consider approximate networking for dealing with limits are described next.

### 2.1 Coping With Resource Constraints

In many developing parts of the world, resource constraints (such as limited power, unstable government) are a norm of life. Even at a global level, it is anticipated that the modern fossil fuel based industrial system is not sustainable, and the impending depletion of these resources will probably give rise to a sudden and permanent shock that may lead to economic instability and infrastructural challenges [2]. Such a severe permanent energy crisis can have far-reaching consequences on the economy and lead developed countries towards being “undeveloping countries” [2]. Approximate networking insights can be used to reorient the design of the Internet’s algorithms, protocols, and infrastructure to better manage the overarching energy, societal, material, and economic limits that this looming scarcity-based future will impose.

## 2.2 Need of Energy Efficiency

Information and communication technology (ICT) is a big consumer of world’s electrical energy, using up to 5% of the overall energy (2012 statistics) [14]. The urgency of developing an energy efficiency manifesto is reinforced when we consider the impending decline of non-renewable energy resources as well as the increased demand of ICT (as more and more people get online and use ICT for exchanging greater and greater amounts of data traffic) [4]. This strongly motivates the need for energy efficient Internet-networking [15]. The approximate networking trend can augment the hardware-focused approximate computing trend to ensure that the energy crisis is managed through the ingenious use of approximation.

### 2.3 The Pareto Principle (80-20 Law): The Power of “Good Enough”

*“Among the factors to be considered there will usually be the vital few and the trivial many.”*—Turan.

To help manage the approximate networking tradeoffs, it is instructive to remember the *Pareto principle*, alternatively called the *80-20 rule* [16]—which states that roughly speaking that 20% of the factors result in 80% of the overall effect. This principle has big implications for approximate networking since this allows us to provide adequate fidelity to ideal networking by only focusing on the most important 20% of the effects. Alternatively put, this theory states that 80% of what goes into creating the ideal networking experience provides little cosmetic benefits to the user. The key challenge in approximate networking then becomes the task of determining these all-important essential non-trivial factors. Through this exercise, we can create “ideal networking” solutions by identifying which factors are costly to implement but provide little gains allowing these resources to be used more efficiently.

## 3. IMPLEMENTING APPROXIMATE NETWORKING UNDER LIMITS

In this section, we propose some concrete building blocks, or loosely speaking principles, that can support approximate networking solutions. We think that the following 5 principles can be useful for implementing approximate networking under limits: (1) *adopt context-appropriate tradeoffs*; (2) *adopt resource pooling & bottom-up networking*; architect a (3) *failure-cognizant network design* and (4) *scarcity-inspired network design*; and finally (5) *design for intermittency*. These principles are discussed in turn next.

In deriving these basic building blocks, we have utilized insights from previous works that have proposed principles for computing within limits [15] [17] [18], robust networking [19], and frugal innovation under scarcity and austerity [20] [21]).

### 3.1 Context-Appropriate Tradeoffs

*“Wisdom is intelligence in context.”*—Unknown.

A *tradeoff* refers to the fact that a design choice can lead to conflicting results in different quality metrics. The performance of computers networks depends routinely on multiple parameters. Since these multiple objectives often conflict with each other, it is rare to find one-size-fits-all solution and tradeoffs have to be necessarily employed. We can borrow concepts from economics to study scarcity and choice. The concept of *opportunity cost*—which is the “cost” incurred by going with the current choice and not adopting any other choice—is a key idea that can be used to ensure efficient usage of scarce resources.

In approximate networking, the task for optimizing for one ex-

explicit parameter is easier than optimizing for multiple optimization variables (such as throughput, delay, energy, etc.). The main problem arises when the various objectives—e.g., jointly minimizing both the BER and the transmit power—conflict and have to compete for dominance (see Figure 1). Another important concept is that of *Pareto optimality*, which refers to a state of resource allocation in which it is not possible to make any one individual better off without making at least one individual worse off.

An interesting approach to understanding tradeoffs is to use visualization techniques. One approach to solving such a problem is to look for a solution on the so called *Pareto Front* (see Figure 1) that defines the set of input parameters that define non-dominated solutions in any dimension. The use of a *tradeoff curve* [22] can be used to visualize bi-objective problems. The problem of visualizing high-dimensional tradeoffs is more challenging. One approach that has been proposed is to utilize Pareto front, which defines the set of values that are each Pareto optimal.

### 3.1.1 Performance vs. Cost Efficiency

We can tradeoff performance (measured in metrics such as *resilience, reliability, throughput*) to gain on cost efficiency. It is said that one of the easiest way to gain cost efficiency is often by sacrificing resilience and reliability (by employing lesser redundancy) [15]. The catch is that by provisioning a lesser-resourced inexpensive network, we are compromising with capacity and will thus have lower throughput for user applications. It is also worth noting that current networks are optimized to comply with very high standards in terms of high-aiming service level agreements (SLA) that aim for very high availability (e.g., 99.999%). By considering failure as an option [23]—i.e., by allowing some failures, and not trying to eliminate them completely at a very high cost—future networks can become significantly more cost efficient.

### 3.1.2 Coverage vs. Consumed Power

In networking, there is often a direct relationship between coverage and consumed power: typically, higher-powered transmissions have a large coverage range. Approximate networking can improve the energy efficiency of systems by incorporating intermittency as a degree of freedom for controlling the consumed power. Since nodes do not need to communicate at all times, researchers have proposed putting to sleep parts of the infrastructure—such as the base transceiver station (BTS) of cellular systems [24]—to save on energy costs. The research challenge that arises from this approach is to ensure that the infrastructure can be activated when communication is needed and no message is lost or delayed inordinately.

### 3.1.3 Performance vs. Coverage/Reliability

In certain cases, it may be appropriate to tradeoff coverage for performance, while the opposite may be true in other situations. In wireless networks, there is a tradeoff between the throughput and the coverage (and the reliability) of a transmission—i.e., for higher-rate transmissions, the coverage area is typically smaller, and the chances of bit errors higher.

### 3.1.4 Managing the tradeoffs in networking

While we have described the main tradeoffs involved in approximate networking and have discussed how they may be visualized, the all-important question still remains to be addressed: *How can we effectively manage these approximate networking tradeoffs?* This is very much an open issue and some open important questions regarding tradeoffs are as follows:

1. *How do we quantify when our approximation is working and when it is not?*

2. *How do we measure success in managing the service quality/accessibility tradeoff?*
3. *How do we measure the cost of approximation in terms of performance degradation?*
4. *How to dynamically control the approximation tradeoffs according to the network condition?*
5. *How to also incorporate social optimality into a user's approximation decision? (We note here that selfish use of approximate networking can result undesirably in one user's performance improving at the cost of all others.)*
6. *How to design proper incentives for the service provider and the user so that both act harmoniously for provisioning a customer-centric contextually-appropriate service?*
7. *How to outsource some of things that we do on the current Internet to less-costly and less energy-intensive offline methods (while ensuring that we get enough QoS that is necessary for our applications [25])?*

## 3.2 Resource Pooling & Bottom-Up Networking

*“Innovative bottom-up methods will solve problems that now seem intractable—from energy to poverty to disease.”—Vinod Khosla*

Broadly speaking, resource pooling involves aggregating a collection of networked resources such that they behave collectively as a single unified resource pool and developing mechanisms for shifting load between the various parts of the unified resource pool. The main benefits of resource pooling include greater reliability and increased robustness against failure; better ability to handle surges in load on individual resources; and, increased utilization [26]. Resource pooling is well suited in scarcity-afflicted approximate networking settings, where maintaining dedicated IT infrastructure and staff is especially cost prohibitive for small-scale entrepreneurs, business owners, and non-profits. Resource pooling can also be especially influential in a world with limits since resource pooling naturally allows some slack in dealing with scarcity and failures.

### 3.2.1 Encouraging Versatility, Recombination, and Reuse

The Latin word *versatilis* connotes turning, or having the capable of turning to varied subjects or tasks. In networks burdened with limits, it will be important to reuse existing resources versatilely for various new settings that may arise. It is possible that a future Internet may require internetworking of partially-connected networks using totally different locally-optimal protocol stacks [4]. A good example of a versatile approximate networking technique is the use of software defined radio (SDR). SDRs by their versatile nature are radio chameleons that can use software programmability to run completely different protocols at different times (e.g. CDMA and Wi-Fi). It will also be important for networks operating under limits to maximize their efficiency by avoiding the waste of resources. In this regard, we can deploy *dynamic spectrum access* [27] to provide secondary users (SUs) access to a primary network when it is not being used by the licensed primary users (PUs). We can also deploy *scavenger transport protocols* to provide less-than-best-effort (LBE) service [28]. LBE service can be used to facilitate background applications in accessing the unused capacity of backhaul links without impacting the performance of priority applications.

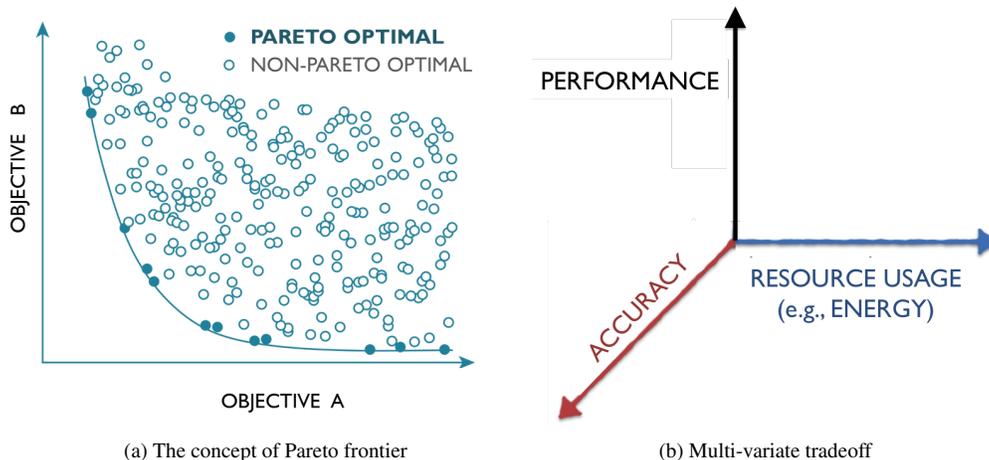


Figure 1: Concepts related to context-appropriate approximate networking tradeoffs

### 3.2.2 Community/Crowdsourced/DIY networks

Being a bottom-up cost-effective approach for building networks, community networking is especially promising for approximate networking under limits. Community networks can be used to implement efficient usage of resources through better resource sharing. In recent times, it has even become possible to develop community cellular networks using low-cost software defined radios (SDRs) and open-source software such as OpenBTS [29]. Such community-driven projects can be used to provide approximate networking services where traditional ideal networking solutions are not feasible. Community networks can allow inclusive services in the future of limits by providing non-priority users (e.g., under-privileged users in developing regions) LBE access to networking services while also ensuring appropriate quality of service (QoS) for priority users (such as the users who are contributing their own networking infrastructure) [30].

### 3.2.3 Multiplicity, Redundancy, and Slack

The principles of multiplicity, redundancy, and using slack may look out of place in a paper on approximate networking. But these ideas are somewhat counterintuitively quite important under limits since any approximate networking solution for such environments that does not have redundancy and slack inbuilt will be debilitatingly fragile. It has been shown in literature that keeping a margin or keeping some slack is a key to frugal innovation [20] and thriving in scarcity-afflicted environments [21]. Approximate networks can also exploit the power of multiplicity by supporting heterogeneous technologies and by resource pooling a diverse collection of paths and thereby unlock the inherent redundancy of the Internet. In particular, approximate networking solutions can leverage the inherent diversity and multiplicity of networks to reap the benefits of increased reliability, efficiency, and fault tolerance [31].

## 3.3 Failure Cognizant Network Design

*“Hoping for the best, prepared for the worst, and unsurprised by anything in between.”—Maya Angelou.*

### 3.3.1 Design for Failure/Collapse

Approximate networking solutions must be designed assuming an inevitable presence of failures/weaknesses/deficiencies. The applications must then be designed to withstand and cope with

some failures while still proving “good enough” services. This is necessary since in the LIMITS setting, we will plausibly deal with many impairments such as long signal propagation delays, high bit error rates (BER), frequent disruptions and unstable or intermittently available links; high congestion; very low data rates; and variable bandwidth. It will be helpful to plan for such a state by assuming severe resource deficiencies such as 25% less power; 25% less connectivity; 10x more volatility; 10x more failures; 10x less non-renewable materials; 10x greater variation in societal structure [15]. Failures should be anticipated and even intentionally utilized where appropriate—e.g., it may be helpful to intentionally cause errors or minor failures; the “random early detection” (RED) congestion control algorithm intentionally drops some packets when the average queue buffer lengths are more than a threshold (e.g. 50%) to signal implicitly to the sender the rising congestion.

### 3.3.2 Robust Design For Avoiding Failure/Collapse

The principle about designing for failure should not be construed to mean that approximate networking should not try to avoid failure. To the contrary, it is very important for approximate networking solutions to be robust<sup>7</sup> and to fail gracefully when subsystems fail. Approximate networking solutions should aim to avoid disruption due to failures by adopting robust tradeoffs that make the solution failure proof or resilient. Towards this end, it has been pointed out in literature that the solutions should “keep the margin” [20] and should have “spare bandwidth” [21] when working in scarcity-afflicted environments.

Approximate networking solutions will also do well to incorporate insights about robust networking gleaned from robust highly-evolved biological and technological systems (such as the Internet) that gracefully degrade when afflicted with failures. Alderson & Doyle [32] have argued that complexity arises in such systems in order to provide robustness to uncertainty in their environments—however, this complexity can also be a source of fragility, leading to a “robust yet fragile” tradeoff in system design. The need to scale out at all levels of the architecture—at both the level of distributed systems and at the macro system level—is also emphasized in the paper by Ragavan [17] for creating “benign systems”, which are

<sup>7</sup>We define some property of a system to be robust if it is invariant with respect to some set of perturbations. Fragility is defined as the opposite of robustness.

computer systems that are less likely to produce harmful impacts to the ecosystem and society.

### 3.3.3 Decentralization

We've earlier discussed how approximate networking should adopt the principles of appropriate technology such as building systems that are simple, locally reproducible, composed of local materials and resources. More generally, in a world burdened by limits, we would like to have resilient infrastructures, and one way of building resilience is to adopt decentralized architectures. Decentralized architectures are more resilient since they can more easily absorb change and disturbance. In a previous work [33], the authors recommend meeting critical human needs such as food, water, energy, communications using alternative decentralized infrastructures (ADIs), which comprise coordinated distributed collections of small-scale systems and services (in preference to large centralized interdependent critical infrastructures that are used in today's settings of abundance and stability). The trend of edge computing, another instance of decentralized computing in which computing is not performed in a centralized cloud but at the edge close to the user, can also be useful in the LIMITS settings [34].

### 3.3.4 Design for Intermittency

In the future world taxed by limits, it will not be possible, nor desirable, to provide networking service all the time. Approximate networking should also be designed while accounting for intermittent availability of resources. This intermittent access may be enforced or volitional. As an example of intentional use of intermittency (also called "duty cycling"), we note that the technique of volitional intermittency can be leveraged in a context-appropriate fashion to provide satisfactory QoS while also saving on energy costs. Approximate networking can draw insights from the rich literature on *delay-tolerant networking* (DTN) [35] and *opportunistic networking* [36] that are also focused on disrupted and intermittently accessible networks.

Traditional cellular networks are designed mostly for performance goals and are not optimized for energy. It has been shown that the "always-on" service approach of traditional cellular networks results in an energy consumption profile that is agnostic of load (i.e., the energy consumption is wastefully high even for light loads) [37]. Such cellular networks can benefit greatly by purposefully employing the approximate networking techniques of intermittency and duty cycling. Such networks can also leverage new advancements in cellular infrastructures—such as the control-data separation architecture for cellular radio access networks [37]—for implementing approximate networking more flexibly.

## 3.4 Scarcity Inspired Network Design

*"What really makes it an invention is that someone decides not to change the solution to a known problem, but to change the question."*—Dean Kamen

Approximate networking can benefit from the insights presented in [20], in which the following guiding principles were provided for frugal innovation in in complex challenging scarcity-afflicted settings: (1) *Seek opportunity in adversity*; (2) *Do more with less*; (3) *Think and act flexibly*; (4) *Keep it simple*; and (5) *Include the margin*.

### 3.4.1 Protocols and Services Optimized for Scarcity

In the future world of scarcity, the majority of the people may be encumbered by poor network connectivity, and prohibitively slow/ unstable services, that are ill suited to the design of conventional protocols and services [38]. This motivates the devel-

opment of optimized protocols and services that can work well in such poor-connectivity scenarios. We motivate this by discussing transport-layer protocols (while noting that scarcity-aware protocols are needed at all layers).

When faced with a severely congested links (which are not uncommon in challenged environments), the congestion control service of conventional Internet transport protocols such as TCP starts to break down as the system tends towards *sub-packet regimes*, where a typical per-flow throughput becomes less than 1 packet per round-trip time (RTT). In small packet regimes, the performance of TCP degrades resulting in severe unfairness, high packet loss rates, and stuttering flows due to repetitive timeouts. For such sub-packet regimes, innovative active queue management (AQM) solutions can be deployed to reduce timeouts and thereby improve fairness and performance predictability [39].

### 3.4.2 Simple Approximate Networking Solutions

Simplicity when coupled with convenience and accessibility can result in wide adoption of approximate networking as it has been shown time and again that users are willing to trade off fidelity of user experience to gain on accessibility and convenience. Simplicity has always been considered a virtuous design trait—e.g., this has been codified in the engineering principles of KISS ("Keep it Simple, Stupid") and the "Occam's Razor" (which recommends adopting the simplest design solution for protocols and not to multiply complexity beyond what is necessary. Previous work has shown that simple protocols with severe constraints can still enable "rich" applications. For example, in situations where mobile users cannot access data services (e.g., due to services not being offered in that location or due to unaffordability): the users can access services through short messaging service (SMS) and voice services. In scenarios where the network is congested, users can even exploit asynchronous voice messages in contrast to live voice calls [40].

### 3.4.3 Sustainable Approximate Networking

While ICT admittedly has many benefits, the unthoughtful use of technology can lead to unintended harmful side effects (e.g., when the society becomes overly reliant on technology, it becomes too reliant on it, and fails to function when technology is disrupted). As we chart out the approximate networking ecosystem, it will be a good time to base approximate networking on the strong architectural foundations of "benign computing" [17], which is focused on minimizing the harmful side effects of technology. In this regard, we can focus on making our approximation networking solutions scale out, fail well, have open design at every level of its structure [17].

## 4. LEVERAGING LIMITS: DOING MORE WITH LESS

*"The impediment to action advances action. What stands in the way becomes the way."*—Marcus Aurelius.

Limits in networking are usually considered as artifacts that constrain performance. But this does not necessarily be the case. We know that many expressive mediums (such as poetry/ paintings) deal with strict constraints, and the beauty and elegance of such media is in how these constraints are managed. The limits in networking can also be viewed analogously—the challenge is to develop approximations that can deal with the limits in liberating ways.

Throwing more resources at a problem is costly and often can lead to sloppy solutions. In fact, it is well known that sometimes removing features can actually improve performance (cf. *Braess' paradox* [41]). Braess' paradox can manifest itself in transportation

networks when opening up new roads can counterintuitively deteriorate traffic conditions while closing down roads can sometimes improve traffic conditions. In the context of networking, it has been shown that counterintuitively the overall efficiency can improve by using a worse service [42].

Doing more with less is not only desirable but will become imperative in the undeveloping future (as the economies of today's advanced countries become stagnate and face growing resource constraints). Fortunately, even with seriously deficient infrastructure, approximate solutions can be remarkably useful. Constrained environments can lead towards lean simple ingenious ideas that can bypass the liability introduced by the constraint; in certain situations, by looking at the problem in the light of restricted resources, a better lean solution may be envisioned that can advance the state of the art more generally (i.e., even for situations sans the constraint). For example, the abundance of resources can mask inefficient design; while additional resources can be used to allay performance bottlenecks, under tight constraints, the possibility of improving implementation's efficiency becomes attractive since it can provide improved performance even with the bottlenecked scarce resources [43].

The resource constraints can also unfetter a "Jugaad"<sup>8</sup> hacker mentality that can lead to novel technical solutions [20]. To illustrate how Jugaad-based thinking can generally advance the state of the art, we highlight how IEEE 802.11 (originally a wireless local area networking standard) was discovered as a technically and economically viable long-distance communication technology by researchers who were driven by the desire to use the low-cost off-the-shelf 802.11 network interface cards (NICs) in constrained settings.

Approximate networking can also lead to "disruptive innovations" [44]. According to Clayton Christensen's "disruptive innovation" theory [44], disruptive innovations typically start as cost-efficient lower-end technologies (let's call them approximate technologies) that do not necessarily meet all the needs of their mainstream customers. For example, consider that Wikipedia initially started as an approximate Britannica Encarta, Skype as an approximate telephone service, and WhatsApp as an SMS service. The users adopt them in droves due to their cost efficiency and these systems often evolve enough eventually to displace high-end "ideal" reigning technologies. The recent uptake of various "over the top" messaging and calling services (such as WhatsApp and Skype) demonstrate the disruptive potential of these approximate networking applications.

## 5. CONCLUSIONS

*"One cannot alter a condition with the same mindset that created it in the first place."*—Albert Einstein

The deep-rooted reliance on infrastructure—which itself depends on many exogenously sourced depletable resources (such as energy and materials)—makes the modern society vulnerable to a disruptive collapse when resources become less available. To cope up with such a likely eventuality—in which the world will be burdened by fundamental limits that will globally affect developing and developed countries alike—we have proposed the idea of "approximate networking". Approximate networking is based on the idea that coping with such a world burdened with limits will require us to adopt context-specific tradeoffs to provide "good enough" ser-

<sup>8</sup>Jugaad is a Hindi/ Urdu work used in the Indian subcontinent for a street-smart hack that can do the job. Jugaad, which can be translated as bricolage, typically involved some ingenious gaming of the system.

vice. In this paper, we have provided some basic building blocks for approximate networking solutions for networks in the LIMITS environment. Determining what these context-appropriate tradeoffs in different LIMITS settings will look like is an important open issue and needs more attention.

## 6. ACKNOWLEDGMENT

The work was funded by the H2020 RIFE project (Grant agreement No.644663).

## 7. REFERENCES

- [1] R. Heinberg, *The party's over: oil, war and the fate of industrial societies*. Clairview books, 2005.
- [2] B. Raghavan, "Networking for undeveloping regions," <http://contraposition.org/blog/2013/04/12/networking-for-undeveloping-regions/>, 2013.
- [3] J. Diamond, *Collapse: How societies choose to fail or succeed*. Penguin, 2005.
- [4] B. Raghavan and J. Ma, "Networking in the long emergency," in *Proceedings of the 2nd ACM SIGCOMM workshop on Green networking*, pp. 37–42, ACM, 2011.
- [5] R. L. Hirsch, R. Bezdek, and R. Wendling, "Peaking of world oil production and its mitigation," *Driving Climate Change: Cutting Carbon from Transportation*, p. 9, 2010.
- [6] J. Mineraud, L. Wang, and J. K. Sasitharan Balasubramaniam, "Hybrid renewable energy routing for isp networks," in *INFOCOM, 2016 Proceedings IEEE*, pp. 1–9, April 2016.
- [7] B. Tomlinson, M. Silberman, D. Patterson, Y. Pan, and E. Bleviss, "Collapse informatics: augmenting the sustainability & ict4d discourse in hci," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 655–664, ACM, 2012.
- [8] W. C. Stirling, *Satisficing Games and Decision Making: with applications to engineering and computer science*. Cambridge University Press, 2003.
- [9] J. Han and M. Orshansky, "Approximate computing: An emerging paradigm for energy-efficient design," in *Test Symposium (ETS), 2013 18th IEEE European*, pp. 1–6, IEEE, 2013.
- [10] V. K. Chippa, D. Mohapatra, K. Roy, S. T. Chakradhar, and A. Raghunathan, "Scalable effort hardware design," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 22, no. 9, pp. 2004–2016, 2014.
- [11] B. Hazeltine and C. Bull, *Appropriate Technology; Tools, Choices, and Implications*. Academic Press, Inc., 1998.
- [12] 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," *Computer*, vol. 38, no. 6, pp. 25–38, 2005.
- [13] B. Tomlinson, J. Norton, E. P. Baumer, M. Pufal, and B. Raghavan, "Self-obviating systems and their application to sustainability," *iConference 2015 Proceedings*, 2015.
- [14] E. Gelenbe and Y. Caseau, "The impact of information technology on energy consumption and carbon emissions," *Ubiquity*, vol. 2015, pp. 1:1–1:15, June 2015.
- [15] B. Raghavan, D. Irwin, J. Albrecht, J. Ma, and A. Streed, "An intermittent energy internet architecture," in *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*, p. 5, ACM, 2012.

- [16] R. Koch, *The 80/20 principle: the secret to achieving more with less*. Crown Business, 2011.
- [17] B. Raghavan, "Abstraction, indirection, and sevarid's law: Towards benign computing," *First Monday*, vol. 20, no. 8, 2015.
- [18] J. Chen, "Computing within limits and ictd," *First Monday*, vol. 20, no. 8, 2015.
- [19] T. Anderson, S. Shenker, I. Stoica, and D. Wetherall, "Design guidelines for robust internet protocols," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 1, pp. 125–130, 2003.
- [20] N. Radjou, J. Prabhu, and S. Ahuja, *Jugaad innovation: Think frugal, be flexible, generate breakthrough growth*. John Wiley & Sons, 2012.
- [21] S. Mullainathan and E. Shafir, *Scarcity: Why having too little means so much*. Macmillan, 2013.
- [22] P. Van Mieghem and L. Vandenberghe, "Trade-off curves for qos routing.," in *INFOCOM*, Citeseer, 2006.
- [23] A. Rice, S. Akoush, and A. Hopper, "Failure is an option," *The Rise and Rise of the Declarative Datacentre*, p. 36, 2008.
- [24] K. Heimerl, S. Hasan, K. Ali, T. Parikh, and E. Brewer, "An experiment in reducing cellular base station power draw with virtual coverage," in *Proceedings of the 4th Annual Symposium on Computing for Development*, p. 6, ACM, 2013.
- [25] J. Greer, *The Long Descent*. New Society Publishers, 2008.
- [26] D. Wischik, M. Handley, and M. B. Braun, "The resource pooling principle," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 5, pp. 47–52, 2008.
- [27] Q. Zhao and B. M. Sadler, "A survey of dynamic spectrum access," *Signal Processing Magazine, IEEE*, vol. 24, no. 3, pp. 79–89, 2007.
- [28] D. Ros and M. Welzl, "Less-than-best-effort service: A survey of end-to-end approaches," *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 2, pp. 898–908, 2013.
- [29] K. Heimerl, S. Hasan, K. Ali, E. Brewer, and T. Parikh, "Local, sustainable, small-scale cellular networks," in *Proceedings of the Sixth International Conference on Information and Communication Technologies and Development: Full Papers-Volume 1*, pp. 2–12, ACM, 2013.
- [30] A. Sathiseelan and J. Crowcroft, "LCD-Net: lowest cost denominator networking," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 2, pp. 52–57, 2013.
- [31] J. Qadir, A. Ali, K.-L. A. Yau, A. Sathiseelan, and J. Crowcroft, "Exploiting the power of multiplicity: a holistic survey of network-layer multipath," *Communications Surveys & Tutorials, IEEE*, vol. 17, no. 4, pp. 2176–2213, 2015.
- [32] D. L. Alderson and J. C. Doyle, "Contrasting views of complexity and their implications for network-centric infrastructures," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 40, no. 4, pp. 839–852, 2010.
- [33] B. Tomlinson, B. Nardi, D. J. Patterson, A. Raturi, D. Richardson, J.-D. Saphores, and D. Stokols, "Toward alternative decentralized infrastructures," in *Proceedings of the 2015 Annual Symposium on Computing for Development*, pp. 33–40, ACM, 2015.
- [34] A. Sathiseelan, L. Wang, A. Aucinas, G. Tyson, and J. Crowcroft, "Scandex: Service centric networking for challenged decentralised networks," in *Proceedings of the 2015 Workshop on Do-it-yourself Networking: an Interdisciplinary Approach*, pp. 15–20, ACM, 2015.
- [35] K. Fall, "A delay-tolerant network architecture for challenged internets," in *ACM SIGCOMM 2003*, pp. 27–34, ACM, 2003.
- [36] C.-M. Huang, K.-c. Lan, and C.-Z. Tsai, "A survey of opportunistic networks," in *Advanced Information Networking and Applications-Workshops, 2008. AINAW 2008. 22nd International Conference on*, pp. 1672–1677, IEEE, 2008.
- [37] A. Mohamed, O. Onireti, M. Imran, A. Imran, and R. Tafazolli, "Control-data separation architecture for cellular radio access networks: A survey and outlook," *IEEE Communications Surveys and Tutorials*, 2015.
- [38] J. Chen, "Computing within limits and ICTD," *First Workshop on Computing within Limits*, 2015.
- [39] J. Chen, L. Subramanian, J. Iyengar, and B. Ford, "TAQ: enhancing fairness and performance predictability in small packet regimes," in *Proceedings of the Ninth European Conference on Computer Systems*, p. 7, ACM, 2014.
- [40] K. Heimerl, R. Honicky, E. Brewer, and T. Parikh, "Message phone: A user study and analysis of asynchronous messaging in rural uganda," in *SOSP Workshop on Networked Systems for Developing Regions (NSDR)*, pp. 15–18, 2009.
- [41] R. Steinberg and W. I. Zangwill, "The prevalence of braess' paradox," *Transportation Science*, vol. 17, no. 3, pp. 301–318, 1983.
- [42] R. Mittal, J. Sherry, S. Ratnasamy, and S. Shenker, "How to improve your network performance by asking your provider for worse service," in *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*, p. 25, ACM, 2013.
- [43] G. Varghese, *Network algorithmics*. Chapman & Hall/CRC, 2010.
- [44] C. Christensen, *The innovator's dilemma: when new technologies cause great firms to fail*. Harvard Business Review Press, 2013.