

Steps towards an Ecology for the Internet

Computing through crisis into the next decade of growth

Anil Madhavapeddy*
University of Cambridge

Sam Reynolds*
University of Cambridge

Alec P. Christie
Imperial College London

David A. Coomes
University of Cambridge

Michael W. Dales
University of Cambridge

Patrick Ferris
University of Cambridge

Ryan Gibb
University of Cambridge

Hamed Haddadi
Imperial College London

Sadiq Jaffer
University of Cambridge

Josh Millar
Imperial College London

Cyrus Omar
University of Michigan

William J. Sutherland
University of Cambridge

Jon Crowcroft
University of Cambridge

ABSTRACT

The Internet has grown from a humble set of protocols for end-to-end connectivity into a critical global system with no builtin “immune system”. In the next decade the Internet will likely grow to a trillion nodes and need protection from threats ranging from floods of fake generative data to AI-driven malware. Unfortunately, growing centralisation has led to the breakdown of mutualism across the network, with surveillance capitalism now the dominant business model. We take lessons from biological systems towards evolving a more resilient Internet that can integrate adaptation mechanisms into its fabric. We also contribute ideas for how the Internet might incorporate digital immune systems, including how software stacks might mutate to encourage more architectural diversity. We strongly advocate for the Internet to “re-decentralise” towards incentivising more mutualistic forms of communication.

CCS CONCEPTS

• **Security and privacy** → **Malware and its mitigation; Artificial immune systems.**

KEYWORDS

ecology, protocols, evolution, botnets, malware, open source, containment, code models, artificial intelligence, specifications

ACM Reference Format:

Anil Madhavapeddy, Sam Reynolds, Alec P. Christie, David A. Coomes, Michael W. Dales, Patrick Ferris, Ryan Gibb, Hamed Haddadi, Sadiq Jaffer, Josh Millar, Cyrus Omar, William J. Sutherland, and Jon Crowcroft. 2025. Steps towards an Ecology for the Internet: Computing through crisis into the next decade of growth. In *The sixth decennial Aarhus conference: Computing X Crisis (AAR 2025), August 18–22, 2025, Aarhus N, Denmark*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3744169.3744180>

*Both authors contributed equally to the paper

1 INTRODUCTION

Ecology is the study of relationships between living things and their environment across all scales, from blue whales in the ocean, to bacteria living on the surface of a leaf. It encompasses all life processes, interactions and adaptations; community structure and composition; competition and foodwebs; ecosystem structure and energy flow; landscapes and life histories; mutualism, parasitism and infectious disease; the abundance and distribution of organisms; and how ecosystems function [97]. Meanwhile, the rather more abiotic Internet is the global system of computer networks that standardises protocols to communicate among each other. It lacks any central control structure and is a network of networks that connect private, public, academic, business and country networks, linked by a broad array of physical connection technologies [98].

Bateson’s classic work “Ecology of Mind” notes that “*the history of evolutionary theory is inevitably a metalogue between man and nature, in which the creation and interaction of ideas must necessarily exemplify evolutionary process*” [9]. In this paper, we explore the idea of an ecology for the Internet, and how the Internet in its current phase of growth can draw parallels from ecological theory to become more resilient, sustainable, and trustworthy into the coming decades as it approaches a trillion connected nodes [83]. To help guide our discussion, we will first clarify our teleology. We are using the term *ecology* directly, not merely as a weak metaphor for a collection of stakeholders in a socio-techno-economic system. We examine those dimensions of ecology that have no intentional purpose [35] and serve to preserve and advance life through selection pressures against competing strategies [7]. The Internet, now entering its fifth decade of existence [64], needs to continue to be based on system designs that provide communication, whatever the adversary does; whatever and whoever the adversary is, even if the adversary is sometimes ourselves.

1.1 The Internet as an Ecology

The Internet started from humble beginnings [17] and has now grown into a behemoth of interconnected nodes with zettabytes of traffic flowing across it daily [16]. As growth continues to multiply, it is facing new existential threats; from the rise of nation-state actors, to the increasing sophistication of cyber-criminals, and the increasing complexity and layering of its protocols resulting in

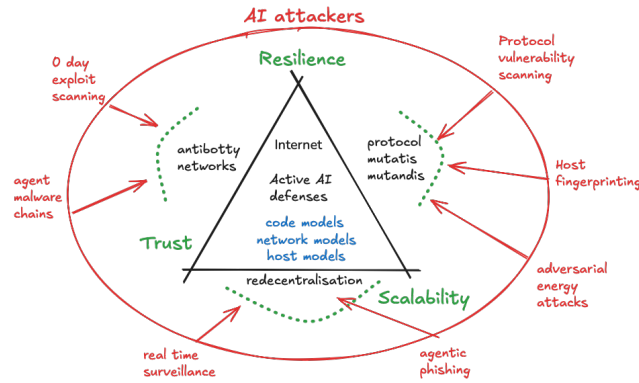


Figure 1: The Internet needs to urgently evolve active defenses against new threats from AI-powered attacks. We discuss three such defences in this paper.

the ossification of its infrastructure [27]. Accelerating all of these threats are emerging AI technologies that can automate attacks in superhuman timescales [24], poison the wells of human information with fake or low-quality generated data [79], and even manipulate the physical world through Internet of Things (IoT) devices [72].

The core purpose of the Internet is to facilitate communications in a decentralized, scalable and extensible manner. It is based on the end-to-end principle [32] that dictates the network itself remain as simple as possible, with complexity pushed to the software on end hosts instead. This has permitted decades of evolutions of communications protocols ranging from the original ARPANET to the modern incarnation [64]. However, the end-to-end principle has also led to the Internet’s current vulnerability to automated attacks as the network itself has no builtin “immune system” to protect itself with, and end hosts cannot upgrade quickly enough.

However, individual Internet hosts can also deploy new software techniques in small clusters without requiring central “permission”. In this sense, they are very similar to the phenomenon of self-organisation that is primarily seen in biological systems. We therefore turn to ecological inspirations, such as plant-pathogen interactions and evolutionary theory, to help us to build a more resilient Internet with more active defence systems.

Table 1 contains an illustrative mapping between biological and Internet systems, beginning firstly from the smallest components and then scaling up to entire populations. Throughout the rest of this paper, we will explore how these biological concepts can be applied to the Internet. Central to our discussion will be we can learn from the natural world to embed new concepts into the Internet architecture that will enable evolution towards a more resilient system despite the number of nodes growing into the trillions. We will also sketch out how some these ideas could be implemented in practice, and what challenges we may face in doing so. Our primary goal is to inspire further research and discussion on this topic, rather than to provide a comprehensive solution.

2 RESILIENCE AGAINST ATTACK

The Internet protocol culture is one of open standards and collaborative processes. Rather than being dictated by central authorities,

Internet standards emerge through “rough consensus and running code”, meaning ideas are adopted when working implementations demonstrate real-world value. Design principles like Postel’s Law (*“be liberal in what you accept, conservative in what you send”*) have encouraged resilience and compatibility [32], while the end-to-end principle keeps the core network simple and pushes complexity to the edges [51]. In theory, this should lead to a diverse ecosystem of implementations by which newer software revisions can extend the protocols in a backwards compatible manner to older endpoints. However, in practice, only a few pieces of software dominate most endpoints. Consequently, if a vulnerability is discovered in these dominant implementations, millions of hosts become immediately susceptible to compromise. And because of the low latency interconnectedness of the vast majority of Internet hosts, this compromise can and has happened within minutes [84].

We can view this widespread software monoculture as being similar to an evolutionary bottleneck: working towards the currently-optimal Internet structure took place under a certain set of tight and idealistic environmental constraints of the time (e.g. hardware and technology availability, an early push for openness and simplicity), which left the network as the fittest version of itself at the time, best adapted to the environment it faced. However, the resultant lack of diversity is now a threat as the pressures and environmental conditions change. The Internet now currently lacks the flexibility to respond to challenges, and these are weaknesses consistent across the entire population due to the software monoculture. The most obviously visible threat is the rise of botnets, which can commandeer vast numbers of machines to launch further attacks [4]. To counteract this, a more resilient model would promote active diversity in the software stack [41]. This model would not only fight against vulnerabilities but also provide shielding to hosts that cannot deploy such countermeasures (consider, for instance, a pacemaker that provides a critical service but cannot be easily updated). A more resilient network would ensure that such devices are protected from potential threats via cooperation and local incentives. Ribiero et al argue that a successful approach to evolving the Internet in such a manner requires backwards compatibility, an evolutionary strategy from the current Internet, and room for new architectures to proliferate [74]. This is a good set of principles to start from; given the rapid growth of botnet- and AI-assisted attacks, we need to build more active translation layers that can adapt (ideally unpredictably) in real-time.

Biological ecosystems also have self-regulation mechanisms [44] where the population of a species is controlled by the availability of resources and nutrient cycles. We know that diverse types of microbes within microbial ecosystems can specialise in distinct functions, contributing substantially to the overall stability of the ecosystem [37]. This also allows for instances of cooperation whereby microbes form alliances¹ for resource acquisition [55], analogous to how Internet routing is a set of cooperative peering agreements [59]. This cooperative imperative also extends across inter- and intra-species relations [12]. In contrast, Internet-connected hosts tend to require individual active management in order to stay connected

¹Lichens are a model example of alliances between fungi and algae or cyanobacteria [2].

Biological Concept	Internet Concept	Section
DNA	Protocol specifications (RFCs), file formats	§3.1
Genes	Software modules, protocol functions, reusable code	§2
Proteins	Operating system processes, code regions	§3.2
Virus	Malware, botnets, Internet worms	§2.1
Chromosome	Software interfaces, formal specifications	§2
Organelle	CPUs, RAM, NICs, SSDs	§3.2
Cell	End nodes (IoT devices, servers, middleboxes)	§3.1
Tissues	Clusters of hosts (data centers, edge clusters)	§3
Organs	Domain-specific networks (e.g., CDN nodes)	§3
Organisms	Entire Internet services (e.g., Google, Facebook)	§4
Populations	Device ecosystems, software monocultures	§4
Communities	Federated networks (e.g., ActivityPub, Bluesky)	§4.1
Ecosystems	Overall Internet resilience	§4.1

Table 1: Just as biological processes exhibit self-organization, we explore in this paper how these principles might apply to the network of networks that is the Internet. The concepts are in rough order of increasing scale.

and secure. Within a network, the collection of hosts is not controlled by the availability of resources, and malware-laden machines can lie dormant and be controlled by botnets for years.

In our first thought experiment, we will discuss an approach to self-regulating Internet networks that can adapt to changes in their environment and take active measures to control the network.

2.1 Active Antibotty Defences

Consider the example of detecting and stopping malware within an urban community. In such a deployment, there are relatively small clusters of hosts that have multiple network connections to each other; for instance, a mobile phone host might gain access to all hosts within another home network when the owner visits a friend and connects to their wifi [67]. The current process of protecting hosts is laboriously manual, involving installation of malware detectors on every endpoint and ensuring the antivirus systems are themselves up-to-date and not themselves going to become routes of exploit. Biological systems operate their defence mechanisms much more autonomously though, via immune systems which fight back routine disease incursions without affecting the overall biological host adversely in routine infections [65]. If we consider malware in the same light as a biological virus, then what is the equivalent of antibodies to automatically detect and eliminate malware in a friend’s home who may not be actively maintaining software updates on all their devices?

One approach would be for every host to act as an antibody² and actively scan the network for signs of malware on nearby friendly hosts. This immediately has a temporal advantage over botnets that must scan whole netblocks [5] to find vulnerable hosts, since each host only checks its local population. When a local host does find evidence of malware, it could take active measures to either isolate the infected host from the network, or directly exploit it before a global botnet does to deliberately patch it. Since the local host is acting in the interests of the immediate community and advertising its intent, it is acting as a vigilante [18] to protect the community

network from harm, and is not incentivised to exploit the host for personal gain.

Such an antibotty³ community network becomes stronger as the community hosts increase their diversity of operating systems and application services (§3.1) since the chances of one of them spotting a vulnerability in their peers increases. There is also the possibility of building higher level services to repair the damage caused by isolating a host; nearby hosts might set up application-level proxies to reroute traffic and even filter traffic until the vulnerable host can be patched. In an extreme case where no patch exists, the offending host could be “killed” by terminating its network connection from its peers — this might happen for an old IoT device which has not received software updates due to obsolescence or repeated takeovers [67]. The antibotty network approach outraces global botnets because local hosts only need to scan orders of magnitude fewer devices in their vicinity, and could therefore be a step towards building a self-regulating Internet ecology. It also turns a weakness into an advantage by using the sophisticated command-and-control mechanisms that have been developed by global botnets for malicious use and deploying them within the community [85].

There is a reason that antibotties are currently a thought experiment though, as great care needs to be taken to ensure that the antibotties do not themselves become a threat. Since every host acts as a vigilante, a structure is needed to ensure that the antibotties do not cause unintended harms. An individual node cannot easily distinguish between mission-critical devices (e.g. a pacemaker) and a non-critical device (e.g. a smart lightbulb in a hallway). Biological systems deal with this via regimented, systematic and ordered processes of development to form a hierarchy of capabilities and protections. This is why, for example, an organism’s heart follows an exact blueprint, or legs do not appear where antennae should be on a fruit fly.⁴ As part of this process, immune systems develop which have multiple ways of responding to external threats (e.g. from mucus to skin), through which different sorts of white blood cells that combat a range of antigens. An important part of this immune system is the blood-brain barrier, which is a semi-permeable

²Or more precisely, act as a T-cell or B-cell; the cells of the immune system which seek out antigens and eliminate them with antibodies or other methods.

³“Antibotty” is our portmanteau of “antibody” and “botnet”

⁴Assuming healthy development and the absence of homeotic transformations.

membrane that is highly selective about what can pass across to the vital brain from the rest of the body, and antibodies typically cannot because they would cause harm to delicate functions [1].

A similar design for the Internet version of immune systems would be to have a hierarchy of controllers to coordinate clusters of antibotties. The design needs to be extended into the network protocols by which vigilante hosts communicate, and the network structure needs to remain resilient to attacks that might exploit the vigilante mechanisms themselves. Like biological systems, antibody networks should establish multiple layers of scoped protections against threats, for example via opportunistic proxies and neighbourly packet inspection on local broadcasts. Although Internet Service Providers have not historically been keen on active countermeasures [25], this attitude may need to change to adapt to the rising tide of AI-powered attacks.

The communication mechanisms behind antibody alerts amongst each other also need some care to not leak private metadata. With existing vigilante mechanisms [18] the alerts also advertise the mechanism behind the attack, which could be exploited by a new adversary to join the network and learn new tricks. One solution would be to seal the alerts and only allow trusted nodes to decode them, but a more radical ecologically inspired mechanism would be to introduce more mutations throughout the software stack in response to the new vulnerability. The biological antibodies produced by species, are by design, incredibly diverse via a process called “somatic hypermutation” [75] whereby mutations are introduced within B-cells to create a diverse range of antibodies. They have error-prone DNA repair and copying proteins by design to create the highest diversity of antibodies possible. We consider how to emulate this in digital antibotties next (§3.1).

3 SCALING THE INTERNET TOWARDS A TRILLION NODES

As the Internet inexorably grows, it needs to do so in a manner that does not make it more brittle to global disasters — either from anthropogenic ones such as malware as discussed earlier, or natural occurrences [31]. The Internet has been quite successful at providing diversity in the physical layer, and transmission systems can use anything that can convey sufficient signal to get an IP packet almost anywhere in the global via RF, optical, electronic, or even mechanical routes. From the IP layer up, we are largely in the domain of software systems which are much easier to reconfigure and program. The key constraint is that we still have to interoperate via protocol specifications that we collectively agree on. As long as network endpoints can understand each other, there is no reason not to have a wide range of such protocols — after all, humans have a wide range of languages, and the natural ecosystem has a much wider range of ways to convey signals between members of same and different species [82, 36, 92] and then even within and between the cells of those individual species..

Mutualistic relationships built over these signals abound between creatures in the wild; in the African Savanna creatures who can smell, hear or see better than each other such as warthogs, impala and giraffes [45] cooperate to give warning to each other of possible predators [21]. In a natural ecosystem “indirect benefits play the main role in explaining cooperation within species” [95], giving us

encouraging signs that we need not advocate for purely immediate metrics of cooperation. Increasing host endpoint diversity would not reduce cooperation as long as there remains consensus on the protocols that allow them to communicate, and indeed in some areas such as wireless networks cooperation is necessary for effective use of the frequency spectrum [40].

How far could we take this diversity within Internet nodes? There exists a staggering diversity of species in nature; mostly millions of varieties of insects [86] and fungal species [68] which in turn influence the development of plants and other aboveground species [28]. These naturally biodiverse systems have evolved to have extreme resilience to changes in the environment, whether from global changes caused by climate changes, pollution, or coincident anthropogenic conversion [90]. And conversely, as ecosystem biodiversity reduces in the anthropocene, the resilience of these natural systems is also declining rapidly [60]. Monocultures in nature not only lead to ecosystem-wide brittleness (such as wildfires having outsize impact on the broader population [56]), but they also close off whole strands of evolutionary exploration. The same holds in electronic systems where a single vulnerability can lead to global compromise [53], thus leading to our next thought experiment.

3.1 Protocol Mutatis Mutandis

A dangerous monoculture of software implementations for end nodes has arisen on the Internet. A few operating systems, namely Windows, macOS and Linux, dominate the landscape, and certain aspects, particularly network stacks, tend to evolve slowly. This ossification occurs for several reasons: the difficulty of changing the userspace APIs that applications use to access the filesystem and network stack [49], the deployment of middleboxes in the network which adopt a fixed and often overly strict interpretation of protocols [69], and (perhaps a controversial position to take) the huge success of the open-source model which encourages code reuse [26].

This all implies that the Internet’s natural state is heading towards one of ossification and centralisation from a software perspective, despite the diversity in physical link mechanisms. What we need is a mechanism that encourages “mutations” to develop in the software stacks of Internet hosts, and a way to measure the effects of these mutations on the network as a whole. This is a protocol implementation evolution that is not driven by a central authority, but by the individual hosts themselves, while remaining respectful of the communication specifications themselves to remain understandable by their peer hosts (which might themselves also be mutating).

The open-source process of code reuse and sharing between different software is similar to Horizontal Gene Transfer [6] whereby, non-sexually, bacteria and archaea and eukaryotes can release DNA (genes) into the environment and be taken up by other cells which can incorporate them directly into their genome (codebase). It has contributed to antibiotic resistance in bacteria, the evolution of eukaryotic cells (all animal cells are eukaryotic) and their endosymbiosis with mitochondria. It allows genetic information to pass between species that would be unable to reproduce sexually as they are too far removed from each other evolutionarily. Doyle observed

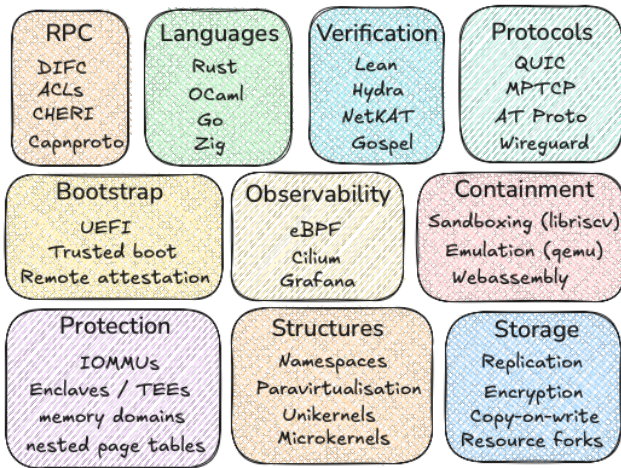


Figure 2: The modern software stack has a number of points where automated patch mutations can be inserted, contained, evaluated and stored without affecting the whole system.

that the Internet IP layer possesses a similar “deconstraining constraint” for electronic communications; IP has a specification that still permits a wide range of interpretations over it [22].

A key aspect of such mutations in software is that they would not just be driven by their genes (that is, the software that was originally written for that host), but could also be guided by the environment within which the hosts are deployed. The Internet, after all, spans everything from remote terrestrial sensors to satellites to embedded pacemakers to vast datacenters. The conventional view of evolutionary adaptations is that they arise from the natural selection of random DNA mutations across generations [30], which would be far too slow for the Internet. However, there has been a decade-long debate in evolutionary biology about a different approach dubbed “extended evolutionary synthesis” [39] that argues that development driven by dynamic conditions also drives how an organism behaves [38]. They note that “evolution must proceed where development leads”, which is a useful perspective for our desire to broaden the Internet’s software diversity.

One emerging mechanism available to kickstart evolution in modern software systems is the rise of AI-driven code models, which are increasingly proficient at program synthesis [46]. While these do not yet generate code as well as a human author [57], with sufficient guardrails, it is possible to create patches for existing code. By developing a system that can patch endpoints incrementally and measure the effects across the entire network, we could perform directed evolution [91] within the software stack more rapidly. Any AI-driven code generation needs adequate context in order to satisfy the constraints of such code mutation. The Internet community has long maintained [33] a rich corpus of plaintext “RFC” specifications that can be interpreted via Large Language Models into formal specifications [81], which can guide LLMs towards generating only permissible protocol messages in code mutations.

3.2 Containment of Mutations

Much like the blood-brain barrier regulates the transfer of chemicals between the circulatory and nervous systems in a biological system, operating systems must also regulate the impacts of software mutations. There are now numerous hardware and software containment layers where code mutations can be inserted safely and with minimal overhead. Figure 2 shows some possible hooks:

- Hardware protections, such as IOMMUs, secure enclaves, nested paging, capabilities and intra-process protections all allow for code mutations to be guarded.
- New operating system structures, such as unikernels and process namespaces, as well as observability mechanisms such as eBPF allow for dynamic checks to be inserted across the software layers.
- Lightweight sandboxing mechanisms, such as Webassembly, are widely deployed in web browsers and cloud environments, allowing for experimental code to be guarded and distributed easily.
- Storage systems now extend traditional POSIX storage interfaces [76] to include support for copy-on-write snapshots, resource forks, encryption and incremental replication, allowing for reliable storage in the face of mutations.
- Newer protocols such as QUIC, Wireguard and the ATProto [34] are all implemented in userspace, allowing for much easier modifications.
- Formal verification frameworks such as Hydra [73] and NetKAT [3] can be used to verify the effects of changes on the network before they are deployed, providing a safety net to routing health.
- Systems languages such as Rust, Go, OCaml and Zig continue to rise in popularity, providing layers of type safety to code changes proposed, as well as formal verification tools such as F* and Lean that provide stronger guarantees for the outputs of code models.

There remain significant challenges to overcome; since malware has free reign to exploit weaknesses, there is a risk that adversarial attacks [101] could allow malware to guide code mutations down undesirable paths that open up new exploits. The same is exactly true in biology as well; viruses hijack cell machinery to make the proteins they need to build more viruses and mimic (or crack) the keys other proteins use to enter cells and hijack them. Research in cell biology has shown that “protein language models” can evolve human antibodies by suggesting mutations that are evolutionarily plausible [29]. The same approach could generate code mutations that are protocol compatible for antibody hosts.

A safer and simpler approach would be to use capability-based programming and hardware extensions [94]. In this model, a piece of code uses a piece of data as both the key and the lock to access a mutual resource. This would allow for both fine-grained tracking of the sources of mutations, and also to craft more sophisticated security policies for individual layers in the software stack [61].

There is also a higher level question about how extensive such patches should be. The existence of sexual reproduction is a puzzle to evolutionary biologists [63], as with each generation females lose half of their genetic material (albeit it is duplicated) and in many cases have no parental care in exchange! The most likely

reason for this is that sexual reproduction results in rapid new combinations of material. These combinations may result in produced diverse lifeforms that are not only exciting new options but are also are different from each other and the parent, so a pathogen or predator that has adapted to the parent faces new genotypes. This is especially important where the pathogens have shorter generation times than the host such as a long-lived oak tree exposed to environmental stresses [93]. So when considering the extent of mutation introduced as part of our software mutatis, we might also swap out entire subsystems with alternative implementations as form of “sexual component mixing”. There are already examples of this within components such as TCP *ex machina* [99], and across operating system components via unikernels [50] that are making their way into conventional operating systems like Linux [71].

4 MUTUALISM AND TRUST

Trust in biological networks can be established because there is constant meaningful and consequential feedback between all parties involved, ranging from the smallest cells to the largest mammals. Many of these relationships are symbiotic in structure to both each other and their context [15, 20]. At a macro level, bird breeding populations, for example, are affected by their local environment network [48]. At a micro level, social interactions among viruses occur whenever multiple viral genomes infect the same cells, hosts, or populations of hosts [43].

One of the most remarkable networks on the planet is the underground mycorrhizal fungal system, which are the dominant form of nutrient exchange between land plants and the soil [100]. The earliest land plants were mycorrhizal⁵ and the root networks that we normally associate with plants only evolved 400 million years after the first fungal networks, as an evolutionary optimisation to facilitate the symbiotic relationship [87].

Mycorrhizal networks exist because the fungi take sugars from plants “in exchange” for moisture and nutrients gathered from the soil by the fungal strands. If trees stopped sharing, it would not be long before the fungi did too. These fungal networks have evolved further to longer range facilitate “nutrient trading networks” that respond to resource inequality by withholding or supplying them to distant nodes [96]. Individual fungal nodes have specialised, with subsets of polymorphic nuclei acting cooperatively, or as traders, or entirely selfishly [58]. The mycorrhizal network itself acts as a key driver for aboveground plant population and biology [89], which in turn provides habitats for many above-ground species. A key point here is that the soil is packed with pathogenic fungi, so plants and fungi have developed intricate handshakes which set the scene for this extraordinary collaboration that happens trillions of times a day [11], but the protocol has yet not been exploited by competing species to “steal” the nutrients!

The Internet looks remarkably similar in many ways, with individual endpoints specialising to certain tasks, and protocols existing for peer-to-peer cooperation. There are longer-range consensus protocols to symbiotically maintain the health of the larger network, such as BGP and DNS that depend on peer communication via a hierarchy. And the Internet substrate itself provides a “habitat”

for the billions of humans that use it every day to communicate among each other. The Internet has never mandated a specific identity model beyond providing addressing and naming mechanisms, instead leaving this to hosts to negotiate. This has led to a healthy proliferation of identity architectures across layers of the protocol stack, from simple username and password combinations, to public key infrastructure, and biometric passkeys. While the early Internet, with its smaller population, exhibited strong mutualism and cooperation, we are now seeing a rise in more parasitic models that break down trust across the network.

The first threat is the rise of advertising driven business models that lead to a centralised attention economy [52]. Trackers deployed across the Internet by a few large providers have led to extremely high resolution user fingerprinting and data collection mechanisms [13]. Google, for example, uses a vast array of third party trackers as “secure signal providers” who share secure signals with advertising bidders competing from clicks from humans.⁶

The second threat is that even such extreme tracking mechanisms are themselves under pressure, as they are not scalable to billions of users who need to prove their “humanity” for websites trying to distinguish between real people and fake clicks [19]. AI has recently made it vastly easier to impersonate humans, and to create fake data that can be used to manipulate these advertising models. This is leading to a crisis of trust of information accessed over the Internet, where it is difficult to determine whether both a given piece of data or the creators are genuine or not. The problem is particularly acute in the context of social communications, where fake news can spread rapidly and cause real-world harm [42]. The move online has changed the economics of news distribution to one of “surveillance capitalism”, contributing to the lowering of diversity of news sources [47] and leaving just a few big players standing who serve a large population.

Mutualism in biological networks — both long-term symbioses and brief exchanges — have transformed entire ecosystems and operate alongside intense competitive interactions for resources [8]. Humans, on the other hand, don’t currently have the ability to easily fight back against the attention-sapping concentration of media arising from both centralised advertising models that are augmented by generative artificial intelligence. The good news for the Internet is that, ultimately, species that “cheat” (that is, undermine the cooperative social contributions of others to further themselves) are often punished by the group and suffer because of it. These cheaters also tend to become less fit over time than other species and cheat out of necessity to maintain their status. Cooperation evolves as species choose cooperative partners, which begets more cooperation [77]. If we consider the current structure to be parasitic in nature, where a few large corporations benefit at the expenses of the larger population of billions, then the Internet is approaching a tipping point where the billions spent on advertising to capture trillions of hours of humans’ attention will be flooded by fakery and will collapse. Ecology tells us that a likely — but not guaranteed [80] — stable fit state after we recover from the parasitic assault is one of mutualism, where the network itself can manage cooperation locally.

⁵The Rhynie Chert, a 407-million-year-old fossilized geothermal site, provides crucial evidence of early land plant-fungal associations, including mycorrhizae.

⁶<https://support.google.com/admanager/answer/10488752?hl=en>

4.1 Redecentralisation

A decade ago in Aarhus, we made an argument for “databoxes” [14], an architecture that shifts private data processing to run locally rather than centrally for sensitive information such as health metrics [103] or travel [23]. Databoxes never took off in the years since, but — as with the unpredictable pace of evolutionary selection — the time may now be right for a redcentralisation of the Internet to emerge! Most Internet trust models have pooled around a monoculture of identity providers who control the roots of trust for almost all of our online communications, with anonymity relegated to specific usecases [54]. But some systems such as DNS and e-mail retain (limited) decentralisation [102], and the fabric of the Internet still allows for plenty of innovation due to the end-to-end principle allowing for new protocols to emerge.

Fake news is of weak quality (or of “lower fitness”) than real news, but its scale and our inability to organise widespread community pushback against it, makes it difficult to punish and allows it to proliferate. A new breed of social network architectures are now emerging to provide that pushback. ActivityPub is one such protocol standardized in the W3C that permits decentralised hosts to communicate with each other for a variety of services such as messaging, microblogging, and photo and video sharing. Bluesky is another, with an underlying “ATProto” [34] that has grown to over 30 million users by Feb 2025 and supports having multiple interoperable providers for every part of the system. One consequence of decentralised algorithms in these systems is that they fundamentally affect how users communicate and perceive each other’s communications with greater agency as to where to focus their attention [88].

It has also been getting simpler for users to *deploy* their own infrastructure. Serverless cloud providers can execute custom user code, custom DNS infrastructure allows for users to own their own domain names and identity online [78], and a self-hosted personal clouds⁷ permit the use of open standards and federated protocols for individual groups and families and organisations to communicate with each other. What’s missing is the symbiotic infrastructure to protect and nurture such activities, which we hope can arrive via mutualistic new protocols that we have discussed here such as antibotties (§2.1) and protocol mutatis mutandis (§3.1). Mutualism in biological networks is also not necessarily the only stable endpoint, with reversion to autonomy also being regularly observed [80], so we should not be surprised if the Internet also oscillates between centralised and decentralised models over the coming decades.

5 CAVEATS

There are several aspects of ecology which we have not discussed in this paper with relation to the Internet’s growing pains. Most obviously, mutualism only works for some biological processes but not for others — it is ideal for nutrient exchange, but not when competing for mates! One aspect of the natural world that we have not included in this discussion, but merits detailed future attention, is *expiry*. In Venki Ramakrishnan’s work on “why we die” [70] he makes it abundantly clear that there are as many reasons why individuals and systems (populations) arise and fall. In nature,

pathogens cause the death of individual cells and bodies, but the fittest survive. In the future we intend to consider how the Internet should incorporate this; we have dealt with this for e-mail spam, for example, but the result has been an increasingly centralised monoculture [102].

Ecology also teaches us that the local spread of disease can be contained by the evolution of disease response, but global rapid spread of diseases (from Covid 19 to ash tree dieback [66]) can be catastrophically rapid. Local defence systems developed over millennia to cope with local diseases do not have the information they need to cope with newcomers, and so some restriction of the flow of information across Internet nodes within political boundaries might actually be critical to ensure the long-term health of the clusters within the larger ecosystem.

It is not yet clear whether the evolution of Internet code is driven by biotic factors (the “Red Queen” model) or abiotic factors (the “Court Jester” model), a topic which is itself an area of debate in the biological sciences [10]. On the one hand, our code mutation proposals can be considered biotic factors and can introduce short-term ecosystem changes. On the other hand, the longer term incentives from the dominant business models of the era may abiotically dominate any such localised strategies. More systemic changes may therefore be required to make room for evolutionary approaches.

6 CONCLUSIONS

The Internet connects around ten billion people and systems. Within your biome alone, there are more than ten trillion microbial cells. And there are 10 million different species of creatures. This scale and diversity survives a wide range of environments and environmental change and, while messy, deserves attention for the wide range of mechanisms that serve that survival, as an inspiration for future technologies. The lack of gene-pool diversity in the Internet and its associated risks has been commented on for over twenty years.⁸ At around the same time, the potential of network systems to evolve was observed by Doyle: “TCP/IP is an example of how architectures that are well-designed for extreme robustness can create evolvability as a side benefit” [22]. Today, we still retain a fundamental optimism towards the long term health and ability of the Internet to resist capture into a singular monoculture.

However, we are not alone in recognising the need for a shakeup now before the ossification sets in too deeply to correct. Farrell and Berjon also recently observed that the “*drive to centralize, control and extract*” has driven the Internet towards an ossified pathology of command and control⁹ and made a powerful case for “rewilding” the Internet. We should collectively derive inspiration from the principles of governing the commons from Ostrom [62], where one main conclusion already sees that rule-based protocols will not scale, and suggest the need to move towards nested ecologically-driven approaches and incentives. We therefore believe it is now the right time to embrace these messy but extraordinarily robust and sustainable mechanisms from nature into artificial systems such as the Internet, so that they too can survive the current crisis and grow and thrive into the coming century.

⁷For example, <https://github.com/RyanGibb/eilean-nix>

⁸<https://web.archive.org/web/20030425082402/http://www.icir.org/mjh/nyt-orig.html>

⁹<https://www.noemamag.com/we-need-to-rewild-the-internet/>

ACKNOWLEDGEMENTS

We gratefully acknowledge our many colleagues in the Cambridge Conservation Initiative and the University of Cambridge Energy and Environment Group who helped debate and shape the arguments in this paper. We thank AI@CAM and the UKRI CRCRM TFS (MR/Y024354) for funding portions of this work, as well as unrestricted donations from Tarides, John Bernstein and Jane Street. Some authors received financial support from Imperial College London through an Imperial College Research Fellowship grant, as well as a Henslow Fellowship funded by the Cambridge Philosophical Society. The funders had no role in the design, analysis, decision to publish, or preparation of the manuscript.

REFERENCES

- [1] N Joan Abbott et al. “Structure and function of the blood–brain barrier”. In: *Neurobiology of disease* 37.1 (2010), pp. 13–25.
- [2] Jessica L Allen and James C Lendemer. “A call to reconceptualize lichen symbioses”. In: *Trends in Ecology & Evolution* 37.7 (2022), pp. 582–589.
- [3] Carolyn Jane Anderson et al. “NetKAT: semantic foundations for networks”. In: *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL ’14. San Diego, California, USA: Association for Computing Machinery, 2014, pp. 113–126. doi: 10.1145/2535838.2535862.
- [4] Dennis Andriess, Christian Rossow, and Herbert Bos. “Reliable Recon in Adversarial Peer-to-Peer Botnets”. In: *Proceedings of the 2015 Internet Measurement Conference*. IMC ’15. ACM, Oct. 2015, pp. 129–140. doi: 10.1145/2815675.2815682.
- [5] Manos Antonakakis et al. “Understanding the mirai botnet”. In: *26th USENIX security symposium (USENIX Security 17)*. 2017, pp. 1093–1110.
- [6] Brian J. Arnold, I-Ting Huang, and William P. Hanage. “Horizontal gene transfer and adaptive evolution in bacteria”. In: *Nature Reviews Microbiology* 20.4 (Nov. 2021), pp. 206–218. doi: 10.1038/s41579-021-00650-4.
- [7] Yoseph Bar-Cohen. *Biomimetics: biologically inspired technologies*. CRC press, 2005.
- [8] Jordi Bascompte. “Mutualistic networks”. In: *Frontiers in Ecology and the Environment* 7.8 (Sept. 2008), pp. 429–436. doi: 10.1890/080026.
- [9] Gregory Bateson. *Steps to an ecology of mind*. en. Chicago, IL: University of Chicago Press, May 2000.
- [10] Michael J. Benton. “The Red Queen and the Court Jester: Species Diversity and the Role of Biotic and Abiotic Factors Through Time”. In: *Science* 323.5915 (Feb. 2009), pp. 728–732. doi: 10.1126/science.1157719.
- [11] Paola Bonfante and Andrea Genre. “Mechanisms underlying beneficial plant–fungus interactions in mycorrhizal symbiosis”. In: *Nature Communications* 1.1 (July 2010). doi: 10.1038/ncomms1046.
- [12] Judith L. Bronstein and Hari Sridhar. “Connecting and integrating cooperation within and between species”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 379.1909 (July 2024). doi: 10.1098/rstb.2023.0203.
- [13] Tomasz Bujlow et al. “A Survey on Web Tracking: Mechanisms, Implications, and Defenses”. In: *Proceedings of the IEEE* 105.8 (Aug. 2017), pp. 1476–1510. doi: 10.1109/jproc.2016.2637878.
- [14] Amir Chaudhry et al. “Personal Data: Thinking Inside the Box”. In: *Aarhus Series on Human Centered Computing* 1.1 (Oct. 2015), pp. 4–4. doi: 10.7146/aahcc.v1i1.21312.
- [15] Guillaume Chomicki et al. “Editorial: Symbiotic Relationships as Shapers of Biodiversity”. In: *Frontiers in Ecology and Evolution* 10 (Feb. 2022). doi: 10.3389/fevo.2022.850572.
- [16] Cisco. *2024 Global Networking Trends Report*. 2024.
- [17] Kerry Coffman and Andrew Odlyzko. “The size and growth rate of the Internet”. In: (1998).
- [18] Manuel Costa et al. “Vigilante: end-to-end containment of internet worms”. In: *SIGOPS Oper. Syst. Rev.* 39.5 (Oct. 2005), pp. 133–147. doi: 10.1145/1095809.1095824.
- [19] Vacha Dave, Saikat Guha, and Yin Zhang. “Measuring and fingerprinting click-spam in ad networks”. In: *Proceedings of the ACM SIGCOMM 2012 conference*. 2012, pp. 175–186.
- [20] Emanuela Del Dottore and Barbara Mazzolai. “Perspectives on Computation in Plants”. In: *Artificial Life* 29.3 (2023), pp. 336–350. doi: 10.1162/artl_a_00396.
- [21] Nathan Diplock et al. “Large mammal declines and the incipient loss of mammal–bird mutualisms in an African savanna ecosystem”. In: *PLOS ONE* 13.8 (Aug. 2018). Ed. by Willem F. de Boer, e0202536. doi: 10.1371/journal.pone.0202536.
- [22] John C. Doyle and Marie Csete. “Architecture, constraints, and behavior”. In: *Proceedings of the National Academy of Sciences* 108.supplement_3 (2011), pp. 15624–15630. doi: 10.1073/pnas.1103557108.
- [23] Chris Elsmore et al. “Confidential carbon commuting: exploring a privacy-sensitive architecture for incentivising ‘greener’ commuting”. In: *Proceedings of the First Workshop on Measurement, Privacy, and Mobility*. MPM ’12. ACM, Apr. 2012, pp. 1–6. doi: 10.1145/2181196.2181201.
- [24] Richard Fang et al. *LLM Agents can Autonomously Exploit One-day Vulnerabilities*. 2024. doi: 10.48550/ARXIV.2404.08144.
- [25] Nick Feamster, Lixin Gao, and Jennifer Rexford. “How to lease the internet in your spare time”. In: *SIGCOMM Comput. Commun. Rev.* 37.1 (Jan. 2007), pp. 61–64. doi: 10.1145/1198255.1198265.
- [26] Daniel Feitosa et al. “CODE reuse in practice: Benefiting or harming technical debt”. In: *Journal of Systems and Software* 167 (Sept. 2020), p. 110618. doi: 10.1016/j.jss.2020.110618.
- [27] Mark Handley. “Why the Internet only just works”. In: *BT Technology Journal* 24.3 (2006), pp. 119–129.
- [28] Marcel G. A. van der Heijden et al. “Mycorrhizal fungal diversity determines plant biodiversity, ecosystem variability and productivity”. In: *Nature* 396.6706 (Nov. 1998), pp. 69–72. doi: 10.1038/23932.
- [29] Brian L. Hie et al. “Efficient evolution of human antibodies from general protein language models”. In: *Nature Biotechnology* 42.2 (Apr. 2023), pp. 275–283. doi: 10.1038/s41587-023-01763-2.

- [30] Eva Jablonka. “A new vision for how evolution works is long overdue”. In: *Nature* 637.8046 (Jan. 2025), pp. 539–541. doi: 10.1038/d41586-025-00054-x.
- [31] Sangeetha Abdu Jyothi. “Solar superstorms: planning for an internet apocalypse”. In: *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. SIGCOMM ’21. Virtual Event, USA: ACM, 2021, pp. 692–704. doi: 10.1145/3452296.3472916.
- [32] J. Kempf. *The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture*. Ed. by R. Austein. Mar. 2004. doi: 10.17487/rfc3724.
- [33] Prashant Khare et al. “The Web We Weave: Untangling the Social Graph of the IETF”. In: *Proceedings of the International AAAI Conference on Web and Social Media* 16 (May 2022), pp. 500–511. doi: 10.1609/icwsm.v16i1.19310.
- [34] Martin Kleppmann et al. “Bluesky and the AT Protocol: Usable Decentralized Social Media”. In: *Proceedings of the ACM Conext-2024 Workshop on the Decentralization of the Internet*. CoNEXT ’24. ACM, Dec. 2024, pp. 1–7. doi: 10.1145/3694809.3700740.
- [35] Alexandra R. Knight and Catherine Allan. “Intentional Ecology: Integrating environmental expertise through a focus on values, care and advocacy”. In: *Humanities and Social Sciences Communications* 8.1 (Nov. 2021). doi: 10.1057/s41599-021-00960-1.
- [36] Kosei Kobayashi et al. *FinchGPT: a Transformer based language model for birdsong analysis*. 2025. doi: 10.48550/ARXIV.2502.00344.
- [37] Christian Kost et al. “Metabolic exchanges are ubiquitous in natural microbial communities”. In: *Nature Microbiology* 8.12 (Nov. 2023), pp. 2244–2252. doi: 10.1038/s41564-023-01511-x.
- [38] Kevin N Lala et al. *Evolution evolving*. en. Princeton, NJ: Princeton University Press, Oct. 2024.
- [39] Kevin Laland et al. “Does evolutionary theory need a rethink?” In: *Nature* 514.7521 (Oct. 2014), pp. 161–164. doi: 10.1038/514161a.
- [40] J.N. Laneman, D.N.C. Tse, and G.W. Wornell. “Cooperative diversity in wireless networks: Efficient protocols and outage behavior”. In: *IEEE Transactions on Information Theory* 50.12 (2004), pp. 3062–3080. doi: 10.1109/TIT.2004.838089.
- [41] Per Larsen et al. “SoK: Automated Software Diversity”. In: *2014 IEEE Symposium on Security and Privacy*. IEEE, May 2014, pp. 276–291. doi: 10.1109/sp.2014.25.
- [42] David MJ Lazer et al. “The science of fake news”. In: *Science* 359.6380 (2018), pp. 1094–1096.
- [43] Asher Leeks et al. “Open questions in the social lives of viruses”. In: *Journal of Evolutionary Biology* 36.11 (Nov. 2023), pp. 1551–1567. doi: 10.1111/jeb.14203.
- [44] Timothy M. Lenton and Marcel van Oijen. “Gaia as a complex adaptive system”. In: *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 357.1421 (May 2002). Ed. by R. V. Solé and S. A. Levin, pp. 683–695. doi: 10.1098/rstb.2001.1014.
- [45] Walter Leuthold. “Relations with Animals of Other Species”. In: *African Ungulates*. Springer Berlin Heidelberg, 1977, pp. 71–83. doi: 10.1007/978-3-642-81073-2_8.
- [46] Yujia Li et al. “Competition-level code generation with AlphaCode”. In: *Science* 378.6624 (Dec. 2022), pp. 1092–1097. doi: 10.1126/science.abq1158.
- [47] Timothy Libert and Reuben Binns. “Good News for People Who Love Bad News: Centralization, Privacy, and Transparency on US News Sites”. In: *Proceedings of the 10th ACM Conference on Web Science*. WebSci ’19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 155–164. doi: 10.1145/3292522.3326019.
- [48] Yu-Heng Lin et al. “Environmental quality mediates the ecological dominance of cooperatively breeding birds”. In: *Ecology Letters* 26.7 (May 2023), pp. 1145–1156. doi: 10.1111/ele.14226.
- [49] Anil Madhavapeddy and Dave Scott. “Unikernels: Rise of the Virtual Library Operating System”. In: *ACM Queue* 11.11 (Nov. 2013), pp. 30–44. doi: 10.1145/2557963.2566628.
- [50] Anil Madhavapeddy et al. “Unikernels: library operating systems for the cloud”. In: *Proceedings of the eighteenth international conference on Architectural support for programming languages and operating systems*. Houston Texas USA: ACM, Mar. 2013, pp. 461–472. doi: 10.1145/2451116.2451167.
- [51] Stephen McQuistin et al. “Characterising the IETF through the lens of RFC deployment”. In: *Proceedings of the 21st ACM Internet Measurement Conference*. IMC ’21. ACM, Nov. 2021. doi: 10.1145/3487552.3487821.
- [52] Filippo Menczer and Thomas Hills. “The attention economy”. In: *Scientific American* 323.6 (2020), pp. 54–61.
- [53] David Moore et al. “Inside the slammer worm”. In: *IEEE Security & Privacy* 1.4 (2003), pp. 33–39.
- [54] Richard Mortier et al. “Using Dust Clouds to Enhance Anonymous Communication”. In: *Security Protocols XVIII*. Ed. by Bruce Christianson and James Malcolm. Vol. 7061. Berlin, Heidelberg: Springer, Mar. 2014, pp. 54–59. doi: 10.1007/978-3-662-45921-8_10.
- [55] Fariha Mostafa et al. “Microbial markets: socio-economic perspective in studying microbial communities”. In: *microLife* 5 (2024). doi: 10.1093/femsm/uaq016.
- [56] “Spreading like wildfire”. In: *Nature Climate Change* 7.11 (Nov. 2017), pp. 755–755. doi: 10.1038/nclimate3432.
- [57] Claudia Negri-Ribalta et al. “A systematic literature review on the impact of AI models on the security of code generation”. In: *Frontiers in Big Data* 7 (May 2024). doi: 10.3389/fdata.2024.1386720.
- [58] Ronald Noë and E. Toby Kiers. “Mycorrhizal Markets, Firms, and Co-ops”. In: *Trends in Ecology and Evolution* 33.10 (Oct. 2018), pp. 777–789. doi: 10.1016/j.tree.2018.07.007.
- [59] William B Norton. *The Internet peering playbook: connecting to the core of the Internet*. DrPeering Press, 2011.
- [60] Tom H. Oliver et al. “Declining resilience of ecosystem functions under biodiversity loss”. In: *Nature Communications* 6.1 (Dec. 2015). doi: 10.1038/ncomms10122.
- [61] Cyrus Omar, Patrick Ferris, and Anil Madhavapeddy. “Modularizing Reasoning about AI Capabilities via Abstract Dijkstra Monads”. In: *the 12th ACM SIGPLAN Workshop on Higher-Order Programming with Effects*. Sept. 2024.

- [62] Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Sept. 2015. doi: 10.1017/cbo9781316423936.
- [63] Sarah P. Otto. “The Evolutionary Enigma of Sex”. In: *The American Naturalist* 174.S1 (July 2009), S1–S14. doi: 10.1086/599084.
- [64] Camille Paloque-Bergès and Valérie Schafer. “Arpanet (1969–2019)”. In: *Internet Histories* 3.1 (2019), pp. 1–14.
- [65] Jacqueline Parkin and Bryony Cohen. “An overview of the immune system”. In: *The Lancet* 357.9270 (2001), pp. 1777–1789.
- [66] Marco Pautasso et al. “European ash (*Fraxinus excelsior*) dieback—A conservation biology challenge”. In: *Biological conservation* 158 (2013), pp. 37–49.
- [67] Stanislaw Piasecki, Lachlan Urquhart, and Professor Derek McAuley. “Defence against the dark artefacts: Smart home cybercrimes and cybersecurity standards”. In: *Computer Law and Security Review* 42 (Sept. 2021), p. 105542. doi: 10.1016/j.clsr.2021.105542.
- [68] Andy Purvis and Andy Hector. “Getting the measure of biodiversity”. In: *Nature* 405.6783 (May 2000), pp. 212–219. doi: 10.1038/35012221.
- [69] Costin Raiciu et al. “How hard can it be? Designing and implementing a deployable multipath TCP”. In: *9th USENIX symposium on Networked Systems Design and Implementation*. 2012, pp. 399–412.
- [70] Venki Ramakrishnan. *Why We Die: The New Science of Ageing and Longevity*. Hachette UK, 2024.
- [71] Ali Raza et al. “Unikernel Linux (UKL)”. In: *Proceedings of the Eighteenth European Conference on Computer Systems*. EuroSys ’23. Rome, Italy: Association for Computing Machinery, 2023, pp. 590–605. doi: 10.1145/3552326.3587458.
- [72] Jingjing Ren et al. “Information Exposure From Consumer IoT Devices: A Multidimensional, Network-Informed Measurement Approach”. In: *Proceedings of the Internet Measurement Conference*. Association for Computing Machinery, 2019, pp. 267–279. doi: 10.1145/3355369.3355577.
- [73] Sundararajan Renganathan et al. “Hydra: Effective Runtime Network Verification”. In: *Proceedings of the ACM SIGCOMM 2023 Conference*. ACM SIGCOMM ’23. New York, NY, USA: Association for Computing Machinery, 2023, pp. 182–194. doi: 10.1145/3603269.3604856.
- [74] Guilherme Ribeiro et al. “Internet Architecture Evolution: Found in Translation”. In: *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*. 2024, pp. 300–307.
- [75] Jana Ridani et al. “Somatic Hypermutation”. In: *Molecular Biology of B Cells*. Elsevier, 2024, pp. 235–256. doi: 10.1016/b978-0-323-95895-0.00012-x.
- [76] Tom Ridge et al. “SibylFS: formal specification and oracle-based testing for POSIX and real-world file systems”. In: *Proceedings of the 25th Symposium on Operating Systems Principles*. Monterey California: ACM, Oct. 2015, pp. 38–53. doi: 10.1145/2815400.2815411.
- [77] Christina Riehl and Megan E. Frederickson. “Cheating and punishment in cooperative animal societies”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 371.1687 (Feb. 2016). doi: 10.1098/rstb.2015.0090.
- [78] Charalampos Rotsos et al. “Lost in the Edge: Finding Your Way with DNSSEC Signposts”. In: *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*. Washington, D.C.: USENIX Association, Aug. 2013.
- [79] Mark A. Runco. “The discovery and innovation of AI does not qualify as creativity”. In: *Journal of Cognitive Psychology* (Dec. 2024), pp. 1–10. doi: 10.1080/20445911.2024.2436362.
- [80] J SACHS and E SIMMS. “Pathways to mutualism breakdown”. In: *Trends in Ecology and Evolution* 21.10 (Oct. 2006), pp. 585–592. doi: 10.1016/j.tree.2006.06.018.
- [81] Prakhar Sharma and Vinod Yegneswaran. “PROSPER: Extracting Protocol Specifications Using Large Language Models”. In: *Proceedings of the 22nd ACM Workshop on Hot Topics in Networks*. ACM, 2023. doi: 10.1145/3626111.3628205.
- [82] Pratyusha Sharma et al. “Contextual and combinatorial structure in sperm whale vocalisations”. In: *Nature Communications* 15.1 (May 2024). doi: 10.1038/s41467-024-47221-8.
- [83] Philip Sparks. *The economics of a trillion connected devices*. 2017.
- [84] Stuart Staniford, Vern Paxson, and Nicholas Weaver. “How to own the internet in your spare time”. In: *11th USENIX Security Symposium*. USENIX, 2002.
- [85] Brett Stone-Gross et al. “Your botnet is my botnet: analysis of a botnet takeover”. In: *Proceedings of the 16th ACM conference on Computer and Communications Security*. ACM, Nov. 2009, pp. 635–647. doi: 10.1145/1653662.1653738.
- [86] Nigel E. Stork. “How Many Species of Insects and Other Terrestrial Arthropods Are There on Earth?” In: *Annual Review of Entomology* 63.1 (Jan. 2018), pp. 31–45. doi: 10.1146/annurev-ento-020117-043348.
- [87] Christine Strullu-Derrien et al. “The origin and evolution of mycorrhizal symbioses: from palaeomycology to phylogenomics”. In: *New Phytologist* 220.4 (Mar. 2018), pp. 1012–1030. doi: 10.1111/nph.15076.
- [88] Aditya Surve, Aneesh Shamraj, and Swapneel Mehta. *How Decentralization Affects User Agency on Social Platforms*. ICWSM, June 2024. doi: 10.36190/2024.74.
- [89] Leho Tedersoo, Mohammad Bahram, and Martin Zobel. “How mycorrhizal associations drive plant population and community biology”. In: *Science* 367.6480 (Feb. 2020). doi: 10.1126/science.aba1223.
- [90] Rose Thorogood et al. “Understanding and applying biological resilience, from genes to ecosystems”. In: *npj Biodiversity* 2.1 (Aug. 2023). doi: 10.1038/s44185-023-00022-6.
- [91] Rongzhen Tian et al. “Establishing a synthetic orthogonal replication system enables accelerated evolution in *E. coli*”. In: *Science* 383.6681 (Jan. 2024), pp. 421–426. doi: 10.1126/science.adk1281.
- [92] Theron S. Wang et al. “Phonetic and Lexical Discovery of Canine Vocalization”. In: *Findings of the Association for Computational Linguistics: EMNLP 2024*. Ed. by Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen. Miami, Florida, USA: Association for Computational Linguistics, Nov. 2024, pp. 13972–13983. doi: 10.18653/v1/2024.findings-emnlp.816.
- [93] PM Wargo. “Consequences of environmental stress on oak: predisposition to pathogens”. In: *Annales des sciences forestières*. Vol. 53. 2-3. EDP Sciences. 1996, pp. 359–368.

- [94] Robert NM Watson et al. “CHERI: A hybrid capability-system architecture for scalable software compartmentalization”. In: *2015 IEEE Symposium on Security and Privacy*. IEEE. 2015, pp. 20–37.
- [95] Stuart A. West et al. “Ten recent insights for our understanding of cooperation”. In: *Nature Ecology and Evolution* 5.4 (Jan. 2021), pp. 419–430. doi: 10.1038/s41559-020-01384-x.
- [96] Matthew D. Whiteside et al. “Mycorrhizal Fungi Respond to Resource Inequality by Moving Phosphorus from Rich to Poor Patches across Networks”. In: *Current Biology* 29.12 (June 2019), 2043–2050.e8. doi: 10.1016/j.cub.2019.04.061.
- [97] Wikipedia. *Ecology*. <http://en.wikipedia.org/w/index.php?title=Ecology&oldid=1275647999>. 2025.
- [98] Wikipedia. *Internet*. <http://en.wikipedia.org/w/index.php?title=Internet&oldid=1275606545>. 2025.
- [99] Keith Winstein and Hari Balakrishnan. “TCP ex machina: computer-generated congestion control”. In: *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*. SIGCOMM ’13. Hong Kong, China: Association for Computing Machinery, 2013, pp. 123–134. doi: 10.1145/2486001.2486020.
- [100] Peter Wohlleben. *The hidden life of trees: What they feel, how they communicate—Discoveries from a secret world*. Vol. 1. Greystone Books, 2016.
- [101] Kaichun Yao et al. “CARL: Unsupervised Code-Based Adversarial Attacks for Programming Language Models via Reinforcement Learning”. In: *ACM Trans. Softw. Eng. Methodol.* 34.1 (Dec. 2024). doi: 10.1145/3688839.
- [102] Luciano Zembruzki et al. “Examining the Centralization of Email Industry: A Landscape Analysis for IPv4 and IPv6”. In: *2023 IEEE Symposium on Computers and Communications*. 2023. doi: 10.1109/ISCC58397.2023.10217834.
- [103] Yuchen Zhao et al. “Privacy-preserving activity and health monitoring on databox”. In: *Proceedings of the Third ACM International Workshop on Edge Systems, Analytics and Networking*. 2020, pp. 49–54.