

The Evolution of Distributed Computing Systems: From Fundamentals to New Frontiers

Dominic Lindsay¹, Sukhpal Singh Gill², Daria Smirnova¹, and Peter Garraghan¹

¹School of Computing and Communication, Lancaster University, UK

²School of Electronic Engineering and Computer Science, Queen Mary University of London, UK

d.lindsay4@lancaster.ac.uk, s.s.gill@qmul.ac.uk, d.smirnova@lancaster.ac.uk, p.garraghan@lancaster.ac.uk

Abstract: Distributed systems have been an active field of research for over 60 years, and has played a crucial role in Computer Science, enabling the invention of the Internet that underpins all facets of modern life. Through technological advancements and their changing role in society, distributed systems have undergone a perpetual evolution, with each change resulting in the formation of a new paradigm. Each new distributed system paradigm - of which modern prominence include Cloud computing, Fog Computing, and the Internet of Things (IoT) – allows for new forms of commercial and artistic value, yet also ushers in new research challenges that must be addressed in order to realize and enhance their operation. However, it is necessary to precisely identify what factors drive the formation and growth of a paradigm, and how unique are the research challenges within modern distributed systems in comparison to prior generations of systems. The objective of this work is to study and evaluate the key factors that have influenced and driven the evolution of distributed system paradigms, from early mainframes, inception of the global inter-network, and to present contemporary systems such as Edge computing, Fog Computing and IoT. Our analysis highlights assumptions that have driven distributed system appear to be changing, including (i) an accelerated fragmentation of paradigms driven by commercial interests and physical limitations imposed by the end of Moore's law, (ii) a transition away from generalized architectures and frameworks towards increasing specialization, and (iii) each paradigm architecture results in some form of pivoting between centralization and decentralization coordination. Finally, we discuss present day and future challenges of distributed research pertaining to studying complex phenomena at scale and the role of distributed systems research in the context of climate change.

Keywords: Distributed Computing, Computing Systems, Evolution, Green Computing

1. Introduction

Societal prosperity of the latter half of the 21st century has been underpinned by the Internet, formed by large-scale computing infrastructure composed of distributed systems which have accelerated economic, social and scientific advancement [1]. The complexity and scale of such systems have been driven by increased societal demand and dependence on such computing infrastructure, which in turn has resulted in the formation of new distributed system paradigms. In fact, these paradigms have evolved in response to technological changes and usage, resulting in alterations to the operational characteristics and assumptions of the underlying computing infrastructure. For example, early mainframe systems provided centralised computing and storage interfaced by teletype terminals. Clustering and packet switching alongside advancement in microprocessor technology and GUIs transferred computing from large mainframes operated remotely to home PCs [5][6]. Standardisation of network protocols enabled global networks-of-networks to exchange messages for global applications [1]. Organisations developed frameworks and protocols capable of offloading computation to remote machine pools of computing resources such as processing, storage and memory [2][3], eventually incorporating sensing and actuator objectives with embedded network capabilities [4]. Thus, distributed systems paradigms have evolved to distribute and facilitate service from centralised

clusters, extending infrastructure beyond the boundaries of central networks forming paradigms such as IoT and Fog computing [8][9].

For the past 60 years distributed system paradigms have conceptually evolved to meet challenges introduced by an ever-changing computing infrastructure and society [47]. From mainframes to clusters, clusters to Cloud, and Cloud to distributed and decentralised infrastructures encompassing the IoT to Edge Infrastructure [52]. Yet paradigms still retain the same underlying characteristics and elements that define their operation [40]. Each is defined by persistent research activities and are often driven by the development of new capabilities, such as security [76], hardware accelerators [77], edge computing [23] and power efficiency [60]. Whilst application framework have evolved to meet challenges presented by integrating with wider eco-systems, ranging from distributed clouds to highly specialised application specific infrastructures [48][74][75]. As such distributed paradigms require constantly evolving middleware's, communication protocols, and secure isolation mechanisms [53].

This work focuses on ascertaining the key characteristics and elements of distributed and networked systems, critically appraise the historical driving technologies and social behaviour that drove their paradigm formation, whilst identifying key trends across the paradigms including system architecture fragmentation, centralisation and decentralisation pivoting, and delays in paradigm conceptualisation to creation by tracing the impact of networked systems on society. From these findings we discuss how future distributed systems will support decentralisation of computation services through composition of decentralised computation platforms specialised to meet workload specific performance goals, forming exponentially larger systems capable of holistic operational requirements including capability and energy availability. Finally we summarise how a dynamic centralised/decentralised distributed paradigms may form and will shape the direction of future computer science research as well as their potential impact within greater society.

The rest of the article is structured as follows: Section 2 presents the background of distributed systems. Section 3 the evolution of the distributed system paradigms. Section 4 analyses trends and observations across all paradigms. Section 5 discusses future challenges facing distributed systems, and Section 6 presents our conclusions.

2. Background

Distributed systems describe a class of computing system in which hardware and software components are connected by means of a network, and coordinate their actions via message passing in order to meet a shared objective [11][12]. Whilst paradigms exhibit differing operational behaviour and leverage various technologies, these systems are defined by their underlying core characteristics and elements that facilitate their operation.

2.1 Characteristics

Transparent Concurrency: Distributed Systems are inherently concurrent, with any participating resource accessible via any number of local or remote processes. The capacity and availability of such a system can be increased by adding resources that require mechanisms for accounting and identification. Such a system is vulnerable to volatile inter-actor behaviours and must be resilient to node failure as well as lost and delayed messages [16]. The management and access of objects, hardware or data in a distributed networked environment is also of particular importance due to potential for physical resource contention [2][6][7][13].

Lack of Shared Clock Computing: Systems maintain their own independent time, interpreted from a variety of sources, and as such Operating Systems (OSs) are susceptible from clock skew and drift.

Furthermore, detecting when a message was sent or received is important for ensuring correct system behaviour. Therefore, events are tracked by means of conceptual Logical and Vector clocks; by sequencing messages, processes distributed across a network are able to ensure total event ordering [10][14][15].

Dependable and Secure Operation: Components of a distributed system are autonomous, and service requests are dependent on correct transaction of operation between sub-systems. Failure of any subsystem may affect the result of service requests and may manifest in ways that are difficult to effectively mitigate. Fault tolerance and dependability are key characteristics towards ensuring the survivability of distributed systems and allow services to recover from faults and whilst maintaining correct service [16].

2.2 Elements

Physical System Architecture: Physical system architecture identifies physical devices that exchange messages in a distributed system and what medium they communicate over. Early distributed systems such as mainframes were physically connected to clients. Later packet switching enabled long-haul multi-hop communication. Cellular networks incorporate mobile computing systems, whilst modern systems host services at specialised hardware between services providers and consumers. Initial designs of distributed systems aimed to provide service across local or campus wide networks of tens to hundreds of machines, and were focused on the development of operating systems and remote storage [1] [2]. Early efforts were designed to explore potential challenges and demonstrate their feasibility [9] and to enhance their functional and non-functional properties (performance, security, dependability, etc).

Entities: A logical perspective of a distributed system describes several process exchanging messages in order to achieve a common goal [17] [18]. Contemporary systems extend this definition by considering logical and aggregate entities, such as Objects and Components, used for abstracting resource and functionality [19]. Here systems are exposed as well-defined interfaces capable of describing natural decomposition of functional software requirements, and enabled exploring the loose-coupling between interchangeable components for domain specific problems found in distributed computing [20]. More recent systems leverage web services and micro-services, that consider their deployment to physical hardware as well as constraints including locality, utilization and stakeholders' policies [35]. Grid and Cloud computing enable distributed computing by abstracting processing, memory and disk space aggregation [21] whereas Fog and Edge computing emphasize integrating mobile and embedded devices [22][28].

Communication Models: Several communication models support distributed systems [24] [25] [26] including (i) *Inter-process Communication*: Enabling two different processes to communicate with each other by means of operating system primitives such as pipes, streams, and datagrams in a client - server architecture; (ii) *Remote Invocation*: Mechanisms and concepts enabling a process in one address space to affect execution of operations, procedures and methods in another address space; and (iii) *Indirect Communication*: Mechanisms enabling message exchanges between one to many processes via an intermediary. In contrast with previous communication models, senders and receiving processes are decoupled, and responsible for facilitating message exchange is passed to the intermediary [37] [38].

Consensus and Consistency: Distributed systems make decisions amongst groups of cooperating processes each possessing possibly inconsistent states. Consensus algorithms are a mechanism in which a majority subset of nodes or 'quorum' can fulfil a client request negotiate a truth and fulfil a client request. Replication and partitioning are common techniques used to improve system scalability, reliability and availability [16] when exposed to volatile environments. Consistency is a challenge to both replicated, partitioned storage and consensus algorithms [10][16].

Table 1. Timeline of Distributed System Paradigms Formation and Key Technological Drivers

Year	Driver	Technology & Paradigm	Model Elements			
			Physical	Conceptual	Entities	Communication
1960 - 1970	Clustering and packet switching (1967-1977)	Inter-process Communication (IPC)	Mainframe and telnet clients.	Client terminal connections share mainframe resources.	Clients (teletype terminals) & servers. Hosts (servers), switches, routers and mainframes	Datagram transport (ATM, X.25)
		Client-server	Local networks interconnected over packet switching infrastructure primarily for research activity.	Networks, provide specific services to private networks, accessible clients across geographic and organisational boundaries.		
		Supercomputer				
		ARPANET and early Internet				
1970-1980	GUI (WIMP), x86 architecture Internet protocols (1974-1984)	ARPANET	Local networks interconnected over packet switching infrastructure.	Private networks provide services across geographic and organisational boundaries.	Hosts (servers), switches, routers and mainframes. Mainframes provide specialised co-processors enabling parallel request processing from clients at scale.	IP addressable hosts are able to communicate by means of datagrams.
		Unix	Home teletype computers and home video games. Early GUI based home computer based systems. Increased memory.	Domains translated to IP addresses for identifying networked hosts. DNS system created.		
		Initial conception of TCP/IP and UDP protocols				
		Distributed Operating Systems				
1980-1990	POSIX.1 Remote Procedure Call (RPC) HTTP and HTML (1985-1990)	Home Computer (Apple LISA, ZX Spectrum, etc)	Mainframe terminals replaced with 8086 microprocessor architectures.	Hosts interconnected by IP addresses and switches; DNS provides address translation; networks remain centralised.	Clusters of microprocessor machines displace monolithic mainframes.	TCP/IP becomes standard internet protocol of internet.
		Standardized TCP/IP and Initial internet	BBS boards begin to appear hosted and & operated by consumers.	Networks ARPANET, NSFNET, DECNET made obsolete by WAN infrastructure via TCP/IP.		
		Generalised OS (drivers)	Move from centralised mainframes to decentralised computers outside of research.			
1990-2000	Middleware Peer to peer protocols	HTTP (TBL)	WWW leads to form geographic internet, services now provided by home servers and clustered machines.	DNS, WWW, and TCP/IP enable decentralised internet. Peer to Peer architecture enables highly decentralised file sharing, parallel processing, and online gaming applications.	Most services now provisioned via off-the-shelf -machines organised into clusters. Servers provide resources described by Uniform Resource Locators.	Remote objects and procedures, enabled development of early middleware HTTP over TCP/IP popularise internet. P2P protocols, group communication
		HTML				
		WWW	Home systems connected by dial up modems.			
		P2P computing				
		Mobile Computing				
2000-2010	High speed broadband x86 Virtualization Hypervisors	Web Services	Educational organizations form Grids for scientific goals.	Grids computing provides orchestration across organizational boundaries.	Most services provisioned via off-the-shelf-machines organised into clusters described by Uniform Resource Locators. Grids and Cloud provide resource pooling (CPU, memory, storage).	Cluster middleware REST, WSDL, XML, JSON, MQTT, XMPP (application layer group comm) Xen and KVM hypervisor.
		Grid computing	VM para-virtualization application mobility.	VMs enables resource isolation between applications on shared hardware.		
		Community Computing	Services and resource consolidation to datacenter.			
		Virtualized Commodity Clusters	Rise of smart phone adoption and mobile computing.	Web services allow further service abstraction from physical hardware.		
		Cloud computing				
2010-2020	Software Defined Networks Containerization	IoT	Fog nodes	Specialization of computing tasks and hardware (GPU, NPU, smart phones, sensors) Remote resources (Storage, processing).	Containers become increasingly prominent Cloudlets	P4, Openflow Open SVN.
		Edge Computing	Smart objects and edge infrastructure			
		Fog Computing	Edge datacenters			

Consistency in distributed systems can be defined as strong consistency, where any update to a partition of a data set is immediately reflect in any subsequent accesses, or weak consistency in which updates may experience delay before they are propagated through the system and are reflected in subsequent access's.

3. The Evolution of Distributed Systems

Distributed systems have continued to evolve in response to various scientific, technological and societal factors. This has given rise to new forms of computer systems, as well as adaptation of paradigms from Client-Server through to IoT and Fog Computing [26]. However, the core characteristics and model elements discussed in Section 2 have remained relatively constant, with the precipitating paradigm augmenting (or re-engineering) technology from prior paradigms. Table 1 provides a detailed a timeline of key distributed paradigm formation, technologies that enabled their realisation, and a description of their respective elements. The formation of distributed systems does not occur in a vacuum, and is influenced by factors spanning other computer science disciplines (e.g. HCI, security), societal exposure, education, and business strategy [24] [25].

Due to the sheer volume of potential influences, we have focused our discussions pertaining to major technological advances and impact upon distributed system elements.

The Mainframe (1960-1967): Mainframes machines of the early 1960's provided time sharing service to local clients that interacted with teletype terminals [29]. Such system conceptualised the client-server architecture, prevalent in present day distributed systems design [30]. The client process connects and requests server processes, enabling a single time-sharing system to multiplex resources amongst clients [31]. Mainframes remained prohibitively expensive and were the focus of supercomputing engineers that lead to the innovation of early disk-based storage and transistor memory [32].

Cluster Networks (1967-1974): The late 1960s and early 1970's saw the development of packet switching, and clusters of off-the-shelf computing components were identified as a cheaper alternative to more powerful yet more expensive supercomputer and mainframes [61]. New programming environments and resource abstractions were developed abstracting resource across local networks of machines [1][2]. This time period also saw the creation of ARPANET and early networks that enabled global message exchange [3], allowing for services hostable on remote machines across geographic bounds decoupled from a fixed programming model. Cerf & Karn [3][4] defined the TCP/IP protocol that facilitated datagram and stream orientated communication over a packet switched autonomous network of networks [39].

Internet & Home PCs (1974-1985): During this era, the Internet was created. Whilst early NCP-based ARPANET systems were characterised by powerful timesharing systems serving multiple clients over networks, new technologies such as TCP/IP had begun to transform the Internet into a network of several backbones, linking local networks to the wider Internet [3]. Thus, the number of hosts connected to the network began to grow rapidly, and centralised naming systems such as HOSTS.TXT could not scale sufficiently [5]. Domain Name Systems (DNSs) were formalised in 1985 and were able to transform hosts domain names to IP addresses; the Unix BIND system was the first public implementation of the DNS. Computers such as Xerox Star and Apple LISA utilizing early WIMP based GUIs demonstrated the feasibility of computing within the home, providing applications such as video games and web browsing to consumers.

World Wide Web (1985-1996): During the late 1980s and early 1990s, the creation of HyperText Transport Protocol (HTTP) and HyperText Markup Language (HTML) [6] resulted in the first web browsers, website, and web-server¹. Standardisation of TCP/IP provided infrastructure for interconnected network of networks known as the World Wide Web (WWW). This enables explosive growth of the number of hosts connected to the Internet, and was the public's first large societal exposure to Information Technology [3][6]. Mechanisms such as Remote Procedure Calls (RPCs) were invented, allowing for the first time applications interfaced with procedure, functions and method across address spaces and networks [7].

P2P, Grids & Web Services (1994-2000): Peer to Peer (P2P) applications such as Napster and Seti@Home demonstrated it was feasible for a global networks of decentralised cooperating processes to perform large-scale processing and storage. P2P enabled a division of workload amongst different peers/computing nodes whereby other peers could communicate with each other directly from the application layer [8]² without the requirement of central coordinator. The creation of Web Services enables further abstraction of the system interface from implementation in the Web [40]. Rather than facilitate direct communication between clients and servers, Web Services mediated communication via a brokerage service [33]. Scientific communities identified that creating federations for large pools of computing resources from commodity hardware could achieve capability comparable to that of large supercomputing systems [41]. Beowulf enables resource sharing amongst process by means of software libraries and middle-wares, conceptualising clustered infrastructure as a single system [42]. Grid computing enabled open access to computing resources and storage by means of open-protocols and middleware. This time period also saw the creation of effective x86 virtualization [43], which became a driving force for subsequent paradigms.

Cloud, Mobile & IoT (2000-2010): A convergence of cluster technology, virtualization, and middleware resulted in the formation of the Cloud computing that enabled creating service models for provision application and computing resource as a service [34]. Driven primarily by large technology organization who constructed large-scale datacenter facilities, computation and storage began a transition from the client-side to the provider side more similar to that of mainframes in the 1960s and 1970s [35] [36]. Mobile computing enabled access to remote resources from resource constrained devices with limited network access [43] [66] IoT also began to emerge from the mobile computing and sensor network communities providing common objects with sensing, actuating and networking capabilities, contributing towards building a globally connected network of 'things' [44].

Fog and Edge Computing (2010-present): Whilst data produced by IoT and Mobile computing platforms continued to increase rapidly, collecting and processing the data in real-time was, and still remains an unsolved issue [27]. This resulted in forming Edge computing whereby computing infrastructure such as power efficient processors, and workload specific accelerators are placed between consumer devices and datacenter providers [66]. Fog computing provides mechanisms that allow for provisioning applications upon edge devices [45][46], capable of coordinating and executing dynamic workflows across decentralised computing systems. The composition of Fog and Edge computing paradigms further extended the Cloud computing model away from centralised stakeholders to decentralized multi-stakeholder systems [45] capable of providing ultra-low service response times, increased aggregate bandwidths and geo-aware provisioning [23][27]. Such a system may comprise of one-off federations or clusters, realised to meet single application workflows or act as intermediate service brokers, and provide common abstractions such as

1 The first webpage --- <http://info.cern.ch/hypertext/WWW/TheProject.html>

2 History of Distributed Systems --- <https://medium.com/microservices-learning/the-evolution-of-distributed-systems-fec4d35beffd>

utility and elastic computing across heterogeneous, decentralised networks of specialised embedded devices, contrasting with centralised networks found in clouds [22].

4. Trends & Observations

By appraising the evolution of the past six decades of distributed system paradigms shown in Table 1, it is apparent that a variety of technological advancements within computer science have driven the formation of new distributed paradigms. It is thus now possible to observe longer-term trends and characteristics of particular interest within distributed systems research.

4.1 Diversification of Paradigms

There appears to be an increased diversification of distributed system paradigms as the research area has matured as shown in Figure 1. This is predominantly driven by two factors: First, it is observable that the acceleration of paradigm formation was precipitated by the invention of the WWW in 1999. This is intuitive as this event enabled distributed systems to transition away from specialized research focused activities into greater society, with each sector requiring specific requirements from entertainment to commercial use. The second reason is that the maturity of fundamental technologies (TCP/IP, HTTP, Unix) created a platform that heavily emphasised abstraction to interconnect heterogeneous platforms in an effective manner, hence future paradigms were able to build upon these concepts. Figure 2 also demonstrates how distributed system paradigms transitioned from a potentially ‘niche’ research with a development singular track within the computer science community towards an area spanning a wide variety of paradigms coinciding with the time the Internet and WWW gained traction.

4.2 Architecture Pivoting from Centralization to Decentralization

The creation of a new technology appears to drive the next distributed system paradigm, and respectively alter its respective degree of centralization as shown in Figure 1. The creation of a new paradigm results in researchers revisiting fundamental mechanisms (schedulers, fault tolerance, monitoring) to ensure that they are capable of effectively operating within the new set of system assumptions. This is exemplified when considering responsibilities frequently carried out by scientists; a principle purpose of peer-review within the research community is ascertaining whether proposed approaches exhibit suitable differences from previous

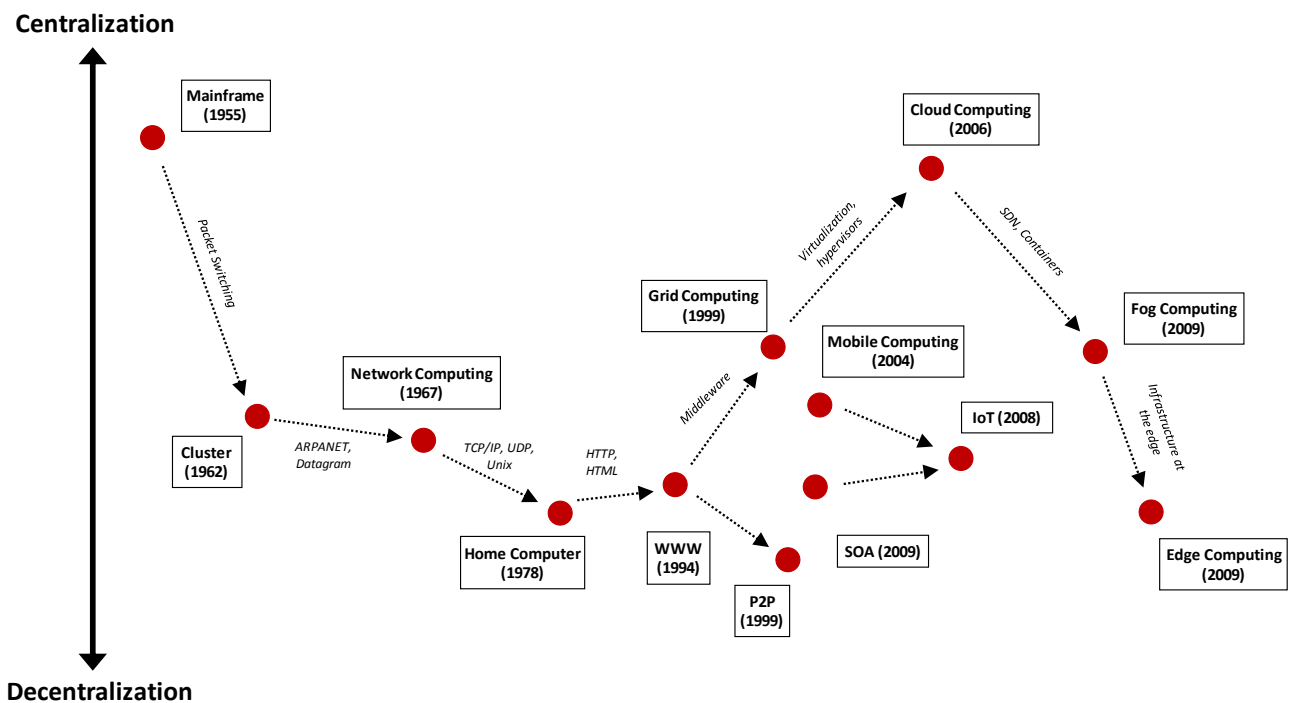


Figure 1. Depiction of distributed system paradigm evolution.

paradigms to determine their novelty (or whether it is a ‘reinvention of the wheel’). This is apparent when considering a number of papers created that attempt to clearly distinguish between paradigms that leverage shared technologies [21]. We observe that the majority of paradigms predominantly are decentralized in nature, with the exception of Cloud computing which follows many similarities with the centralized mainframe in terms of the coordination of computational resources within a datacenter facility which users access via web APIs.

4.3 Time Between System Conception & Creation




Figure 2. Time gap between paradigm conception and creation.

The delay between the description of a potential paradigm and actual successful implementation in recent years appears to be shorter in contrast to previous decades as shown in Figure 2. It is worth noting that ascertaining the precise publication fully credited in accurately describing the full realization of a paradigm due to a single individual or group is not necessarily feasible. Thus, we have attempted to seek papers which first define the appropriate terminology and paradigm description that were later adopted. As shown Figure 2, the formative years of distributed systems between 1960 - 1996 saw an average delay of 13 years and after the adoption of the WWW saw an average 8.8-years delay. It is observable that most paradigm are conceived and created sometime within 3-10 years, with the exception between 1960 – 1990 which is likely due to insufficient technologies when first envisioned, Later paradigms again appear to be relatively short in duration to create, and is likely a by-product of increased maturity of the research area, combined with its pervasiveness within society and growth of research activity within each respective paradigm (i.e. there are a sizable proportion of distributed researchers whom focus on a particular paradigm).

5. Future of Large-Scale Computing Infrastructure

5.1 Accelerated Paradigm Specialization

It is observable that specific distributed system paradigms have a particular affinity for tackling different objectives; whilst Cloud computing is capable of handling generalized application workload, paradigms such as edge computing and fog computing have been envisioned to be particularly effective for sensor actuation and increasingly important latency requirements. A growing number of microprocessors are being designed to accelerate specific tasks (such as graphics and machine learning using GPUs and NPUs, respectively). In tandem, the end of Moore's law indicates that by 2025 chip density will reach a scale where heat dissipation and quantum uncertainty make transistors unreliable [54]. When combining all of these factors together, it is apparent computing systems are in the process of undergoing massive diversification. This diversification is not solely limited to hardware but can also be observed in software.

For example, the last decade has seen resource management undergo a transition from centralized monolithic scheduling to decentralized model architecture [47, 48, 66]. Centralized schedulers maintain a global view of cluster state and are therefore able to make high quality placement decisions at the cost of latency [3][4][49][50]. However, decentralized schedulers maintain only partial state about the cluster, and so they are able to make low latency decision at the cost placement quality [51]. As a result, we envision that further diversification and fragmentation of the distributed paradigm will continue to accelerate and affect all of its respective elements. For example, it is not hard to envision that the system that enables an infrastructure autonomous vehicle operation being substantially different to that of remote sensor networks and smart phones; we are already seeing such diversification with making custom OSs and applications for these scenarios. In the case of cluster resource management there have been an increased research activity in hybrid schedulers, capable of multiplexing centralized and decentralized architectures [52,53], and we expect that future distributed systems must be capable of architectural adaptivity in response to changes to operation.

5.2 Generalization against Specialization

Related to paradigm specialization discussed in Section 5.1, the distributed systems research area appears to be at a particularly interesting cross-roads; ensure that system paradigms are designed to be generalizable to handle a wide variety of operational conditions and scenarios (at the cost of performance and efficiency), or alternatively focus on creating more specialized and bespoke distributed system more suitable to a particular task at the expense of generalization and portability. While the wide-spread adoption of the x86 architecture, middleware, and virtualization have reinforced that historically the community has championed generalizable and portability, continued diversification of paradigms and technological limitations have begun discussion whether the axis is pivoting in the other direction [55]. This is further reinforced by increasing customization of microprocessors, OSs, and power management techniques for particular use case scenarios. For example, the increased uptake of deep learning has resulted in further increase into research into GPU and NPUs inside and outside of the datacentre, as well as creation of cluster resource schedulers specifically for deep learning [65].

5.3 Complexity at Scale & the Role of Academic Research

An area of potential future research challenge moving forward is how to understand these future distributed systems at scale. For many years Computer scientists have leveraged well-structured system abstractions in order to reduce the complexity to understand component interactions and assumptions. However increasingly there have been difficulties in handling unseen emergent behaviour within massive-scale distributed systems

[62] that require rethinking well-established assumptions for system mechanisms [63]. Moreover, with the rapid uptake of new technologies such as deep learning and reinforcement learning to conduct decision making of system operation [64], whilst introduction of temporal applications and mobile compute will likely lead to increased complexity of distributed system operation at scale. In relation to the academic research community, where there is a substantive reliance on simulation or small to medium-scale distributed systems, it will continue to become increasingly difficult to evaluate effectiveness of their approaches when exposed to emergent behaviour within systems at scale. Whilst production systems from industry can greatly support understanding of distributed systems at scale, it does not provide an avenue to conduct experiments within a controller environment to test hypothesis effectively.

5.4 The Green Agenda

Growing end-use demand for applications and subsequent data generation in the regions of Exabyte will usher in the first system at Exascale by 2020, and eventually Zettascale by 2035 [56]. Whilst an achievement in itself, it also brings a variety of associated challenges. One challenge which is particularly problematic is the enormous power requirements that will be necessary for operation. ICT presently consumes more than 10% of the global electricity annually [57]. The creation of ever larger systems through efficiency improvements is in fact detrimental due to the Rebound Effect [58] that causes even greater demand and consumption. At a time where climate change and a 1.5°C increase in global temperatures by 2100 due to Greenhouse Gas Emissions [59], we foresee energy and GHG emissions being increasingly important for future distributed system paradigms. This is not solely increasing energy-efficiency as we see today, but more fundamental concerns related to systems assuming operate constant stable power sources, integration with renewable energy sources, and alternative methods for reducing energy consumption but also computation itself. An area of particular interest is that of holistic coordination of energy management (asynchronous computing, voltage scaling, Wake-on-LAN, cooling, etc.) [60] towards studying and treating systems as living eco-systems, as opposed to individual components in isolation.

5.5 Shifting from Centralised Systems to Decentralised Edge

The evolution of centralised systems towards decentralised system transformed many industries and organisations which have resulted in significant contributions towards economic growth worldwide [80]. With the emergence of Big Data, centralised cloud systems have played an important role to process both structured and unstructured data in an efficient manner [67]. With the rapid adoption of IoT technology, these systems are able to process large amount data using various machine learning algorithms. It is difficult to process real-time jobs on centralised cloud systems due to increases in latency and response time, and incurs various complexities: New distributed applications (cryptocurrencies, the machine economy etc.) require computing models which are not compatible with existing centralised cloud systems [66]. As the adoption of edge computing is increasing, decentralised edge systems have been positioned to be particularly effective process user workloads immediately on powerful edge devices without the reliance upon large cloud datacenters [68], thus reducing round trip communication times at the cost of reduced computational performance. The evolution from centralised cloud systems to decentralised edge is growing among various industries while executing IoT based decentralised applications [69]. It is likely that given sufficient time and technological innovations, this pendulum will swing in the opposite direction, whereby computationally powerful decentralised systems will in turn form centralised architectures (and possibly an assortment of centralised systems coordinated via federation).

5.6 Distributed Green Computing

Rapid growth in large scale distributed application servicing paradigms ranging from Big Data and Machine Learning to the Internet of Things, are increasingly responsible for the world's energy consumption and as such a major contributor to environmental pollution [79]. One such example includes distributed Machine Learning systems [73] – comprising of clusters of GPUs dedicated to Deep Learning applications; require effective energy management aware scheduling policies [70]. As such new orchestration mechanism capable of capturing GPU, CPU, and memory energy characteristics [71] informing new scheduling algorithms prioritising energy consumption in contrast with traditional performance and fairness scheduling objectives [60] [77] [78]. Such scheduler should holistically consider energy consumption and account for out of band costs including impact of workload consolidation on cooling systems [60] [78]. Furthermore, exergy and energy source can be utilised to further inform datacentre operators about the carbon impact of their infrastructure. Whilst, hybrid energy grids utilizing green intermittent decentralised energy sources including solar and wind can provide clean energy whilst brown energy source can be utilized at peak time, minimized reliance of fossil fuels energy sources, and achieve new sustainable computing standards [72].

6. Conclusions

In this paper, we have discussed and evaluated the evolution of the distributed paradigm over the past six decades by focussing on the development and decentralised pivoting of networked computing systems. We have identified core elements of distributed systems by describing their physical infrastructure, logical entities and communication models. We examine how cross cutting factors such conceptual and physical models influence centralisation and decentralisation across various paradigms. We observe long term trends in distributed systems research, by identifying influential links between system paradigms, and technological breakthroughs. Of particular interest, we have observed that distributed system paradigms have undergone a long history of decentralisation up until the inception of the World Wide Web. In the following years, pervasive computing paradigms --- such as the Internet of Things --- brought about by advancements and specialisation of microprocessor architecture, operating systems designs, and networking infrastructure further diversified both infrastructure and conceptual systems. Furthermore, it is apparent that the diversification of distributed systems paradigms that begun at conception of the World Wide Web is likely to further accelerate due to increased emphasis on decentralisation and prioritization of specialized hardware and software for particular problems within domains such as machine learning and robotics. This is somewhat removed from the past few decades which has emphasized generality and portability of distributed system operation and as such will be the focus of research efforts over the coming years. Moreover, there are potentially difficult challenges on the horizon related to the upfront cost of operating large systems testbeds out of reach for most academic laboratories, and the impact of climate change and how it shapes future system design.

Acknowledgements

This work is supported by the UK Engineering and Physical Sciences Research Council (EP/P031617/1).

References

- [1] M. Armbrust et al., “Above the Clouds: A Berkeley View of Cloud Computing,” EECS Dep. Univ. California, Berkeley, no. JANUARY, pp. 1–25, 2009.
- [2] A. Botta, W. De Donato, V. Persico, and A. Pescap, “Integration of Cloud Computing and Internet of Things : A Survey”, *Future Generation Computer Systems*, Vol. 56, pp. 684-700, 2016.

- [3] M. I. Xinghuo Yu, Fellow IEEE, and Yusheng Xue, "Smart Grids: A Cyber-Physical Systems Perspective," *Proc. IEEE* | Vol. 104, no. 5, pp. 1058–1070, 2016.
- [4] Cisco Systems, "Fog Computing and the Internet of Things: Extend the Cloud to Where the Things Are," www.cisco.com, p. 6, 2016.
- [5] Leslie Lamport, "Time, clocks, and the ordering of events in a distributed system," *Commun. ACM*, vol. 21, no. 7, pp. 558–565, 1978.
- [6] K. W. Chow Yuan-Chieh, "Models for dynamic load balancing in a heterogeneous multiple processor system," *IEEE Trans. Comput.*, vol. C, no. 5, pp. 354–361, 1979.
- [7] A. D. Birrell and B. J. A. Y. Nelson, "Implementing Remote Procedure Calls," vol. 2, no. 1, pp. 39–59, 1984.
- [8] T. G. Walker Bruce, Popek Gerald, English Robert, Kline Charles, "The LOCUS Distributed Operating System," pp. 49–70, 1983.
- [9] A. D. Birrell, R. Levin, M. D. Schroeder, and R. M. Needham, "Grapevine: an exercise in distributed computing," *Commun. ACM*, vol. 25, no. 4, pp. 260–274, 1982.
- [10] L. Lamport, R. Shostak, and M. Pease, "The Byzantine Generals Problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [11] P. H. Enslow, "What is a Distributed Data Processing System?," vol. 11, no. 1, pp. 13–21, 1978.
- [12] L. Gerard, "Distributed Systems - Towards a Formal Approach," *IFIP Congr.*, 1977.
- [13] D. Thain, T. Tannenbaum, and M. Livny, "Distributed computing in practice: The Condor experience," *Concurr. Comput. Pract. Exp.*, vol. 17, no. 2–4, pp. 323–356, 2005.
- [14] C. Figde, "Logical Time in Distributed Computing systems," *Computer (Long. Beach. Calif.)*, pp. 28–33, 1991.
- [15] M. Friedemann, "Virtual Time and Global States of Distributed Systems," *SIAM J. Comput.*, vol. 28, no. 5, pp. 1829–1847, 1999.
- [16] L. C. Algirdas Avizienis, Laprie Jean-Claude, Randell Brian, "Basic Concepts and Taxonomy of Dependable and Secure Computing," *IEEE Trans. Dependable Secur. Comput.*, vol. 1, no. 1, pp. 11–33, 2004.
- [17] V. S. Sunderam, G. A. Geist, J. Dongarra, and R. Manchek, "The PVM concurrent computing system: Evolution, experiences, and trends," *Parallel Comput.*, vol. 20, no. 4, pp. 531–545, 1994.
- [18] W. Gropp, "An Introduction to MPI Parallel Programming with the Message Passing Interface," pp. 1–48, 1998.
- [19] P. K. Gummadi, S. D. Gdbble, and U. Washington, "A Measurement Study of Napster and Gnutella as Examples of Peer-to-Peer File Sharing Systems," *Comput. Commun. Rev.*, no. January, p. 2002, 2002.
- [20] D. P. Anderson, J. Cobb, E. Korpela, M. Lebofsky, and D. Werthimer, "Seti@home An Experiment in Public-Resource Computing," *Commun. ACM*, vol. 45, no. 11, pp. 56–61, 2002.
- [21] I. Foster, Y. Zhao, I. Raicu, and S. Lu, "Cloud Computing and Grid Computing 360-degree compared," *Grid Comput. Environ. Work. GCE 2008*, pp. 1–10, 2008.
- [22] P. Mell and T. Grance, "The NIST Definition of Cloud Computing Recommendations of the National Institute of Standards and Technology," *Nist Spec. Publ.*, vol. 145, p. 7, 2011.
- [23] R. K. Naha et al., "Fog Computing: Survey of Trends, Architectures, Requirements, and Research Directions," vol. X, pp. 1–31, 2018.
- [24] R. Baheti and H. Gill, "Cyber-physical Systems," *Impact Control Technol.*, no. 1, pp. 161–166, 2011.
- [25] S. Karnouskos, "Cyber-physical systems in the SmartGrid," *2011 9th IEEE Int. Conf. Ind. Informatics*, vol. 1 VN-re, 2011.
- [26] D. Evans, "The Internet of Things - How the Next Evolution of the Internet is Changing Everything," *CISCO white Pap.*, no. April, pp. 1–11, 2011.
- [27] S. S. Gill, P. Garraghan, and R. Buyya. "ROUTER: Fog enabled cloud based intelligent resource management approach for smart home IoT devices." *Journal of Systems and Software* 154 (2019): 125-138.

- [28] S. Singh and I. Chana. "A survey on resource scheduling in cloud computing: Issues and challenges." *Journal of grid computing* 14, no. 2 (2016): 217-264.
- [29] M. J. Flynn, "Very High-speed Computing Systems," vol. 54, no. 12, pp. 1901–1909, 1966.
- [30] S. Singh, I. Chana and M. Singh. "The journey of QoS-aware autonomic cloud computing." *IT Professional* 19, no. 2 (2017): 42-49.
- [31] J. K. Casavant Thomas, "A Taxonomy of Scheduling in General-Purpose Distributed Computing Systems," vol. 14, no. 2, 1988.
- [32] K. Compton and S. Hauck, "Reconfigurable Computing : A Survey of Systems and Software," vol. 34, no. 2, pp. 171–210, 2002.
- [33] J. Yu and R. Buyya, "A Taxonomy of Workflow Management Systems for Grid Computing," pp. 1–31.
- [34] S. Singh and I. Chana, "QoS-Aware Autonomic Resource Management in Cloud Computing: A Systematic Review," vol. 48, no. 3, 2015.
- [35] A. Celesti, "Open Issues in Scheduling Microservices in the Cloud the types of devices that might," pp. 81–88, 2016.
- [36] B. M. Leiner et al., "Internet Society (ISOC) All About the Internet : A Brief History of the Internet Internet Society (ISOC) All About the Internet : A Brief History of the Internet," pp. 1–18, 2000.
- [37] Cerf VG; RE Icahn, "A Protocol for Packet Network Intercommunication," *ACM SIGCOMM Comput. Commun. Rev.* 71 Vol. 35, Number 2, April 2005, vol. 35, no. 2, pp. 71–82, 1974.
- [38] D. K. Mockapetris Paul, "Development of the Domain Name System," *SIGCOMM '88 Symp. Commun. Archit. Protoc.*, 1988.
- [39] D. Lindsay, S. S. Gill, and P. Garraghan. "PRISM: an experiment framework for straggler analytics in containerized clusters." In *Proceedings of the 5th International Workshop on Container Technologies and Container Clouds*, pp. 13-18. 2019.
- [40] C. Peltz, "Web services orchestration and choreography," *IEEE Internet Comput.*, 36 (10), 46–52, 2003.
- [41] I. Foster, C. Kesselman, and S. Tuecke, "The Anatomy of the Grid," *Hand Clin.*, vol. 17, no. 4, pp. 525–532, 2001.
- [42] T. Sterling, D. J. Becker, D. Savarase, J. E. Dorband, U. A. Ranawake, and C. V Packer, "BEOWULF: A parallel workstation for scientific computation," *Proceedings of the 24th International Conference on Parallel Processing*. pp. 2–5, 1995.
- [43] S. S. Gill, X. Ouyang, and P. Garraghan. "Tails in the cloud: a survey and taxonomy of straggler management within large-scale cloud data centres." *The Journal of Supercomputing* (2020): 1-40
- [44] A. Whitmore, A. Agarwal, and L. Da Xu, "The Internet of Things — A survey of topics and trends," no. March 2014, pp. 261–274, 2015.
- [45] A. Brogi, S. Forti, C. Guerrero, and I. Lera, "How to Place Your Apps in the Fog - State of the Art and Open Challenges," 2019.
- [46] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge Computing: Vision and Challenges," *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, 2016.
- [47] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. D. Joseph, R. H. Katz, S. Shenker, and I. Stoica, "Mesos: A platform for fine-grained resource sharing in the data center.," in *NSDI*, 2011, vol. 11, pp. 22–22.
- [48] V. Vavilapallih, A. Murthyh, C. Douglasm, M. Konarh, R. Evansy, T. Gravesy, J. Lowey, S. Sethh, B. Sahah, C. Curinom, O. O'Malleyh, S. Agarwali, H. Shahh, S. Radiah, B. Reed, and E. Baldeschwieler, "Apache Hadoop YARN," in *SoCC* , 2013, pp. 1–16.
- [49] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, "Large-scale cluster management at google with Borg," in *Proceedings of the Tenth European Conference on Computer Systems, EuroSys '15*, (New York, NY, USA), ACM, 2015, pp. 18:1–18:17.
- [50] I. Gog, M. Schwarzkopf, A. Gleave, R. M. N. Watson, and S. Hand, "Firmament: Fast, centralized cluster scheduling at scale," in *Proc. 12th USENIX Symp. Oper. Syst. Design Implement.*, 2016, pp. 99–115.

- [51] K. Ousterhout, P. Wendell, M. Zaharia, I. Stoica, "Sparrow: distributed, low latency scheduling", Proceedings of the 24th ACM Symposium on Operating Systems Principles, 2013, pp. 69-84.
- [52] P. Delgado, F. Dinu, A.-M. Kermarrec, and W. Zwaenepoel, "Hawk: Hybrid datacenter scheduling," in USENIX ATC, 2015, pp. 499–510.
- [53] K. Karanasos, S. Rao, C. Curino, C. Douglas, K. Chaliparambil, G. M. Fumarola, S. Heddaya, R. Ramakrishnan, and S. Sakalanaga, "Mercury: Hybrid centralized and distributed scheduling in large shared clusters," in USENIX ATC, 2015, pp. 485–497.
- [54] M. Waldrop "The Chips are Down for Moore's Law", Nature, 2016.
- [55] G. Blair "Complex Distributed Systems: The Need for Fresh Perspectives", IEEE ICDCS, 1410-1421, 2018.
- [56] X. Liao, "Moving from Exascale to Zettascale Computing: Challenges and Techniques", Froniers of Information Technology & Electronic Engineering, pp. 1236-1244, 2018.
- [57] W. V. Heddeghem, et al. "Trends in Worldwide ICT Electricity Consumption from 2007 to 2012", Computer Communications, 2014.
- [58] C. Gossart, "Rebound Effects and ICT: A Review of the Literature", ICT Innovations for Sustainability, pp.435-448, 2014.
- [59] IPCC, "Global Warming of 1.5 °C", Intergovernmental Panel on Climate Change, 2018.
- [60] X. Li, et al "Holistic virtual machine scheduling in cloud datacenters towards minimizing total energy", IEEE Transactions on Parallel and Distributed Systems, pp. 1317-1331, 2018.
- [61] G. M. Amdahl, "Validity of the single processor approach to achieving large scale computing capabilities," AFIPS spring Jt. Comput. Conf., pp. 1–4, 1967.
- [62] S. S. Gill and A. Shaghghi. "Security-Aware Autonomic Allocation of Cloud Resources: A Model, Research Trends, and Future Directions." Journal of Organizational and End User Computing (JOEUC) 32, no. 3 (2020): 15-22.
- [63] P. Garraghan, et al "Emergent Failures: Rethinking Cloud Reliability at Scale", IEEE Cloud Computing, vol. 5, pp. 12-21, 2018.
- [64] J. Gao, "Machine Learning Applications for Data Center Optimization", Google White Paper, 2014.
- [65] W. Xiao, et al, "Gandiva, Introspective Cluster Scheduling for Deep Learning" OSDI, 2018.
- [66] S. S. Gill et al. "Transformative Effects of IoT, Blockchain and Artificial Intelligence on Cloud Computing: Evolution, Vision, Trends and Open Challenges." Internet of Things (2019): vol. 8, 100118.
- [67] A. J. Ferrer, J. Manuel Marquès, and J. Jorba. "Towards the decentralised cloud: Survey on approaches and challenges for mobile, ad hoc, and edge computing." ACM Computing Surveys 51, no. 6 (2019): 1-36.
- [68] M. A. Khan, F. Algarni, and M. T. Quasim. "Decentralised Internet of Things." In Decentralised Internet of Things, pp. 3-20. Springer, Cham, 2020.
- [69] I. Psaras. "Decentralised edge-computing and iot through distributed trust." In Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services, pp. 505-507. 2018.
- [70] S. S. Gill, P. Garraghan, V. Stankovski, G. Casale, R. K. Thulasiram, S. K. Ghosh, K. Ramamohanarao, and R. Buyya. "Holistic resource management for sustainable and reliable cloud computing: An innovative solution to global challenge." Journal of Systems and Software 155 (2019): 104-129.
- [71] R. Yang, C. Hu, X. Sun, P. Garraghan, T. Wo, Z. Wen, H. Peng, J. Xu, and C. Li. "Performance-aware speculative resource oversubscription for large-scale clusters." IEEE Transactions on Parallel and Distributed Systems 31, no. 7 (2020): 1499-1517.
- [72] S. S. Gill, S. Tuli, A. N. Toosi, F. Cuadrado, P. Garraghan, R. Bahsoon, H. Lutfiyya et al. "ThermoSim: Deep learning based framework for modeling and simulation of thermal-aware resource management for cloud computing environments." Journal of Systems and Software 164 (2020): 110596.
- [73] W. Xiao, R. Bhardwaj, R. Ramjee, M. Sivathanu, N. Kwatra, Z. Han, P. Patel, X. Peng, H. Zhao, Q. Zhang, F. Yang, L. Zhou. 2018. Gandiva: introspective cluster scheduling for deep learning. In Proceedings of the 13th

- USENIX conference on Operating Systems Design and Implementation (OSDI'18). USENIX Association, USA, 595–610.
- [74] B. Burns, B. Grant, D. Oppenheimer, E. Brewer, and J. Wilkes, “Borg, omega, and kubernetes,” *Commun. ACM*, vol. 59, no. 5, pp. 50–57, 2016.
 - [75] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker, and I. Stoica, “Discretized streams: Fault-tolerant streaming computation at scale,” *SOSP 2013 - Proc. 24th ACM Symp. Oper. Syst. Princ.*, no. 1, pp. 423–438, 2013.
 - [76] S. Arnavtsov et al., “SCONE: Secure linux containers with Intel SGX,” *Proc. 12th USENIX Symp. Oper. Syst. Des. Implementation, OSDI 2016*, pp. 689–703, 2016.
 - [77] I. R. Z. Michael Kaufmann, IBM Research Zurich, Karlsruhe Institute of Technology; Kornilios Kourtis, “The HCl Scheduler: Going all-in on Heterogeneity,” 9th {USENIX} Work. Hot Top. Cloud Comput. (HotCloud 17), pp. 1–7, 2017.
 - [78] K. Ma, X. Li, W. Chen, C. Zhang, and X. Wang, “GreenGPU: A holistic approach to energy efficiency in GPU-CPU heterogeneous architectures,” *Proc. Int. Conf. Parallel Process.*, pp. 48–57, 2012.
 - [79] A. Alqahtani, E. Solaiman, P. Patel, S. Dustdar, R. Ranjan (2019). Service level agreement specification for end-to-end IoT application ecosystems. *Software: Practice and Experience*, 49, 12, pp. 1689-1711
 - [80] A. Chandra, J. Weissman, and B. Heintz. "Decentralized edge clouds." *IEEE Internet Computing* 17, no. 5 (2013):