BenchCouncil Transactions

TBench

Volume 2, Issue 3

2022

on Benchmarks, Standards and Evaluations

Research Article

High fusion computers: The IoTs, edges, data centers, and humans-in-the-loop as a computer Wanling Gao, Lei Wang, Mingyu Chen, Jin Xiong, ... Jianfeng Zhan

Ounderstanding hot interconnects with an extensive benchmark survey Yuke Li, Hao Qi, Gang Lu, Feng Jin, ... Xiaoyi Lu

Review Article

 A review of Blockchain Technology applications for financial services
 Mohd Javaid, Abid Haleem, Ravi Pratap Singh, Rajiv Suman, Shahbaz Khan

ISSN: 2772-4859

Copyright © 2023 International Open Benchmark Council (BenchCouncil); sponsored by the Institute of Computing Technology, Chinese Academy of Sciences. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd.

BenchCouncil Transactions on Benchmarks, Standards and Evaluations (TBench) is an open-access multi-disciplinary journal dedicated to benchmarks, standards, evaluations, optimizations, and data sets. This journal is a peer-reviewed, subsidized open access journal where The International Open Benchmark Council pays the OA fee. Authors do not have to pay any open access publication fee. However, at least one of BenchCouncil International register the authors must Symposium on Benchmarking, Measuring and Optimizing (Bench) (https://www.benchcouncil.org/bench/) and present their work. It seeks a fast-track publication with an average turnaround time of one month.

Contents

High fusion computers: The IoTs, edges, data centers, and humans-in-the-loop as a computer
W. Gao, L. Wang, M. Chen, J. Xiong, C. Luo, W. Zhang, Y. Huang, W. Li, G. Kang, C. Zheng, B. Xie, S. Dai, Q. He, H. Ye, Y. Bao and J. Zhan
Understanding hot interconnects with an extensive benchmark survey
A review of Blockchain Technology applications for financial services <i>M. Javaid, A. Haleem, R.P. Singh, R. Suman and S. Khan</i>
TBench Editorial Board 47
TBench Call For Paper 48
Bench 2022 Call For Paper

Contents lists available at ScienceDirect

BenchCouncil Transactions on Benchmarks, Standards and Evaluations

journal homepage: www.keaipublishing.com/en/journals/benchcouncil-transactions-onbenchmarks-standards-and-evaluations/

Research Article

KeAi

High fusion computers: The IoTs, edges, data centers, and humans-in-the-loop as a computer

Wanling Gao ^{a,b}, Lei Wang ^{a,b}, Mingyu Chen ^{a,b}, Jin Xiong ^{a,b}, Chunjie Luo ^{a,b}, Wenli Zhang ^{a,b}, Yunyou Huang ^c, Weiping Li ^d, Guoxin Kang ^{a,b}, Chen Zheng ^e, Biwei Xie ^{a,b}, Shaopeng Dai ^{a,b}, Qian He ^f, Hainan Ye ^f, Yungang Bao ^{a,b}, Jianfeng Zhan ^{a,b,*}

^a State Key Lab of Processors, and Research Center for Advanced Computer Systems, Institute of Computing Technology, Chinese Academy of Sciences, No. 6 Kexueyuan South Road, Haidian District, 100190, Beijing, China

^b University of Chinese Academy of Sciences, No. 19 (A) Yuquan Road, Shijingshan District, 100049, Beijing, China

^c Guangxi Normal University, No. 15, Yucai Road, Qixing District, Guilin, 541010, Guangxi, China

^d Civil Aviation Flight University of China, No. 46, Section 4, Nanchang Road, Guanghan, 618307, Sichuan, China

e Institute of Software, Chinese Academy of Sciences, No. 4, South 4th Street, Zhongguancun, 100190, Beijing, China

^f Beijing Institute of Open Source Chip, No. 7 Caoqiao, Northwest Corner of Haidian Bridge, Haidian District, 100080, Beijing, China

ARTICLE INFO

Keywords: Emerging and future applications Safety-critical Mission-critical IoT Edge Data center Humans-in-the-loop Open-source computer systems High Fusion Computers HFC

ABSTRACT

Emerging and future applications rely heavily upon systems consisting of Internet of Things (IoT), edges, data centers, and humans-in-the-loop. Significantly different from warehouse-scale computers that serve independent concurrent user requests, this new class of computer systems directly interacts with the physical world, considering humans an essential part and performing safety-critical and mission-critical operations; their computations have intertwined dependencies between not only adjacent execution loops but also actions or decisions triggered by IoTs, edge, datacenters, or humans-in-the-loop; the systems must first satisfy the accuracy metric in predicting, interpreting, or taking action before meeting the performance goal under different cases.

This article argues we need a paradigm shift to reconstruct the IoTs, edges, data centers, and humans-inthe-loop as a computer rather than a distributed system. We coin a new term, high fusion computers (HFCs), to describe this class of systems. The fusion in the term has two implications: fusing IoTs, edges, data centers, and humans-in-the-loop as a computer, fusing the physical and digital worlds through HFC systems. HFC is a pivotal case of the open-source computer systems initiative. We laid out the challenges, plan, and call for uniting our community's wisdom and actions to address the HFC challenges. Everything, including the source code, will be publicly available from the project homepage: https://www.computercouncil.org/HFC/.

1. Introduction

The past decades have witnessed solid achievements and ambitious plans on planetary-scale infrastructures. Typical examples include but are not only limited to Grid computing [2], planet-scale data centers hosting internet services or called warehouse-scale computers [3], virtual supercomputers in the cloud [4], interplanetary Internet [5], networked systems of embedded computers [6], planet-scale computing networks [7,8] or planetary computer [9]. However, emerging and future applications raise daunting challenges beyond the reach of the state-of-the-practice systems.

Emerging and future applications rely heavily on systems consisting of IoTs, edges, data centers, and humans-in-the-loop [10]. These networked systems of embedded computers [6] or IoTs, collaborating with data centers, edges, and humans-in-the-loop, can radically change

* Correspondence to: Institute of Computing Technology, Chinese Academy of Sciences, No. 6 Kexueyuan South Road, Haidian District, 100190, Beijing, China *E-mail addresses:* gaowanling@ict.ac.cn (W. Gao), wanglei_2011@ict.ac.cn (L. Wang), cmy@ict.ac.cn (M. Chen), xiongjin@ict.ac.cn (J. Xiong),

luochunjie@ict.ac.cn (C. Luo), zhangwl@ict.ac.cn (W. Zhang), huangyunyou@gxnu.edu.cn (Y. Huang), weiping2012c@gmail.com (W. Li), kangguoxin@ict.ac.cn (G. Kang), zhengchenjason@hotmail.com (C. Zheng), xiebiwei@ict.ac.cn (B. Xie), daishaopeng@ict.ac.cn (S. Dai), heqian@bosc.ac.cn (Q. He), yehainan@bosc.ac.cn (H. Ye), baoyg@ict.ac.cn (Y. Bao), zhanjianfeng@ict.ac.cn (J. Zhan).

¹ According to [1] a safety-critical system is "a system whose failure may result in injury, loss of life or serious environmental damage, e.g., a control system for a chemical manufacturing plant". https://ifs.host.cs.st-andrews.ac.uk/Books/SE9/Web/Dependability/CritSys.html

² According to [1] a mission-critical system is "a system whose failure may fail some goal-directed activity, e.g., a navigational system for a spacecraft"; a business-critical system is "a system whose failure may result in very high costs for the business using that system, e.g., customer accounting system in a bank". https://ifs.host.cs.st-andrews.ac.uk/Books/SE9/Web/Dependability/CritSys.html

https://doi.org/10.1016/j.tbench.2022.100075

Received 20 September 2022; Received in revised form 26 October 2022; Accepted 26 October 2022

Available online 28 October 2022

2772-4859/© 2022 The Authors. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).





the way people interact with the physical world and perform safetycritical,¹ or mission-critical,² tasks. This new class of systems appears in many forms and continues to expand: "implemented as a kind of digital nervous system to enable instrumentation of all sorts of spaces, ranging from in situ environmental monitoring to surveillance of battlespace conditions" [6], e.g., climate change monitoring and defense systems; embodied as integrated instrumentation, operation, maintenance, and regulation facilities of critical physical infrastructure, e.g., industrial digital twin, energy infrastructure management and civil aviation regulation; "employed in personal monitoring strategies (both defense-related and civilian), synthesizing information from sensors on and within a person with information from laboratory tests and other sources" [6], e.g., medical emergency applications; augmented as an extension of the real-world life for entertainment, education, and social activities, e.g., Metaverse; instrumented as a kind of digital sensing and autonomic control systems to perform safety-critical or mission-critical tasks, e.g., automotive driving and interplanetary explorations.

Significantly different from warehouse-scale computers that nonstop serve independent concurrent user requests [3], the new class of computer systems has three unique requirements. First, they directly interact with the physical world - considering humans an essential part: human-in-the-loop, performing safety-critical and mission-critical operations, and a significant fraction of actions may have an irreversible effect. The role of humans and their interactions with the other system components cannot be ignored in the final impact on the physical world. Second, unlike Internet services that process independent concurrent requests across data centers, their computations have intertwined dependencies between not only adjacent execution loops (which we call internal dependencies) but also actions or decisions triggered by IoTs, edge, data centers, or humans-in-the-loop (which we call external dependencies), and traverse different paths through and around IoTs, edges, and data centers. Third, under this highly entangled state, the systems must first satisfy the quality metric before meeting the performance goal under different conditions, like worst-case, average-case, and best-case. The quality metric measures the accuracy of an application, task, or algorithm in predicting, interpreting, or taking action.

Even considering a simple IoT application — 95% queries reporting the status and 5% user queries accessing the database, serving 10billion devices needs about one thousand to one hundred thousand nodes under the threshold of 50 ms response time. Further, considering the three unique system requirements mentioned above and the sea change in computing, data access, and networking patterns, this new class of computer systems demand resources that are several orders of magnitude beyond the reach of the state-of-the-practice systems, which raises daunting challenges.

This article argues that we need a paradigm shift to rebuild the IoTs, edges, data centers, and humans-in-the-loop as a computer rather than a distributed system. We coin a new term, high fusion computers (HFCs), to describe this new class of computer systems. According to the Oxford English Dictionary, fusion means "the process or result of joining two or more things together to form a single entity". Fusion in HFCs has two-fold meanings: fusing IoTs, edges, data centers, and humans-in-the-loop as a computer, fusing the physical world and digital world. Fig. 1 shows a concept viewpoint of HFCs. Our intuition in rebuilding the HFC systems is simple: We explicitly value the role of humans-in-the-loop in different contexts and consider them as essential components of the system; we aggressively embrace co-design from vertical and horizontal dimensions, and will co-explore the design space from the algorithms, runtime systems, resource management, storage, memory, networking, and chip systems from a vertical dimension; Meanwhile, we will consider the close collaboration among IoTs, Edges, data centers, and humans-in-the-loop from a horizontal dimension, and ponder how to facilitate the interactions of humans-in-the-loop with other hardware and software systems.

To dismantle the complexity of building the systems and improve the efficiency, we use the funclet abstraction, architecture, and methodology, inspired by the philosophy of building large systems out of smaller functions [11,12]. The funclet abstract represents "the common proprieties of basic building blocks: each funclet is a well-defined, independently deployable, testable, evolvable, and reusable functionality with modest complexity; funclets interoperate with each other through well-defined interconnections" [12]. Four kinds of funclets form the four-layer funclet architecture: chiplet, HWlet, envlet, and service layers, respectively [12].

We take HFC as a pivotal example of the open-source computer system plan. We abstract reusable functions (funclets) across system stacks among IoTs, edges, and data centers. Based on funclets, we rebuild the IoTs, edges, data centers, and humans-in-the-loop as a computer in a structural manner, with full-fledged functions of autonomic resource discovery, management, programming, workload scheduling, and coordinated collaboration between software, hardware and human components. Our plans are three-fold. First, we value the importance of benchmarks and funclet-based standards in evaluating and building the systems. Second, we emphasize the methodology and tool to facilitate the workload-driven exploration of the system and architecture design space. Third, we will provide the first open-source implementation of the funclet architecture of HFC systems.

We organize the rest of this paper as follows. Section 2 explains the motivation. Section 3 illustrates the HFC challenges. Section 4 explores the HFC software and hardware design space. Section 5 describes our plan. Section 6 summarizes the related work. Section 7 concludes.

2. Motivation

In this section, we first analyze seven typical emerging and future applications' unique requirements, then explain why we need to build an HFC system.

2.1. The requirements of emerging and future applications

This subsection analyzes seven emerging and future applications. Table 1 characterizes those applications. I detail two applications specifically as follows.

2.1.1. Medical emergency management

According to the data from the World Bank, there are more than 723 million people over the age of 65 in the world in 2020, accounting for 9.321% of the world's total population [14]. What is worse, despite the slowdown in world population growth, the proportion of people over the age of 65 is growing rapidly, which will account for 16% of the total population by 2050 [15,16]. Due to the decline of physical function and pathological changes, the elderly will experience many unplanned emergencies in terms of companionship, nursing, medical treatment, etc., bringing massive pressure to the emergency medical care of the entire lifecycle in the future [17]. In addition, the surge of patients will rapidly overwhelm the overcrowded medical facilities when a disaster occurs [18].

Many computing technologies (e.g., IoT, AI, cloud computing) are introduced to support the health system to overcome current and future dilemmas of the elderly emergency medical care [19–22]. However, as shown in Fig. 2, Many elderly emergency medical care issues of the entire lifecycle remain unsolved, posing enormous challenges to the computer systems sustaining emergency medical care applications.

• Task types: emergent and mainly safe-critical. The elderly are more likely to experience medical emergencies (e.g., stroke, myocardial infarction, falls, etc.) than younger adults, and these events often cause more significant harm to the elderly. When an emergency medical event occurs, it is required that the medical system can handle the safe-critical task in real-time and that



Fig. 1. A Concept Viewpoint of High Fusion Computers (HFCs).

Table 1

Summarization of emerging or future applications.

Summarization of e	inerging of future	applications.							
Application	Critical or typical tasks	Task type	Effect	Metrics	IoT Devices	Dependencies*	Data Management	Access Patterns	
Medical	Emergency detection	Safety-critical	Reversible	Worst-case	Camera, blood pressure monitor, spirometer,	Observe, fuse, recommend, train	Image, video,	Real-time write;	
emergency	Rescue planning	Mission-critical	Reversible	Average-case	gyroscopes, CT scanner,	(Internal & external	relational data, XML	random read	
	Rapid diagnosis	Safety-critical	Reversible	Worst-case	mass spectrometer, etc.	dependencies)			
	Trajectory planning	Safety-critical			Camera, LiDAR, Radar,	Observe, fuse, act,	Image, video, LAS	Real-time write;	
Autonomous driving	Surrounding object detection	Safety-critical	Irreversible	Worst-case	ultrasonic, GNSS, GPS, etc.	& external dependencies)	text, XML, float matrix, csv	Periodic burst read	
	Autonomic control	Mission-critical							
Smart defense systems	Battlespace surveillance	Safety-critical	Irreversible	Worst-case	Seismic, acoustic, magnetic, and imaging sensors or terminal control units, etc.	Observe, orient, decide, act (Internal & external dependencies)	SEG-Y files, MP3, WAV, image	Real-time write; Real-time read	
Digital Twin	Smart manufacturing	Mission-critical	Reversible	Average-case	Cameras, sensors, analog-to-digital	Observe, model, decide, control	Image, binary, relational data	Periodic write; Periodic read	
0	Oil well drilling	Safety-critical	Irreversible	Worst-case	converter, digital-to-analog converter, etc.	(Internal & external dependencies)			
	Airport security	Safety-critical		Worst-case	Cameras, Radars, VHF, Flight-Data Acquisition Unit (FDAU), etc.	Observe, fuse, decide, and alert (Internal & external dependencies)			
Civil aviation	Air navigation	Mission-critical					Image, binary, relational data	Real-time write; Random read	
safety regulation	Anomaly detection	Safety-critical	- inteversible						
	Scenario generating	Mission-critical	Reversible	Average-case	Head-mounted display	Observe, recognize,	Dynamic multimedia, relational data		
Metaverse	Avatar maintaining	Mission-critical	Reversible	Average-case	(HMD), Handheld devices (HHDs), etc.	external dependencies)		Real-time read	
	Decentralized finance	Mission-critical	Irreversible	Worst-case					
Interplanetary	Knowledge discovery	Mission-critical	Reversible	Best-case	Catallita Chago proba	Observe, recognize,	Image, Hierarchical data format(HDF), Network Common Data Form (NetCDF)	Real-time write; Batch transfer; Random read	
explorations	Collision avoidance	Safety-critical	Irreversible	Worst-case	Robots, etc.	(Internal & external dependencies)			
	Space navigation	Mission-critical	Irreversible	Worst-case					

*Internal dependency indicates the dependencies between adjacent execution loops.

*External dependency indicates the dependency between actions or decisions triggered by IoTs, edge, data centers, or humans-in-the-loop.

the medical system can reasonably allocate medical resources for rapid rescue. It is worth noticing that medical experts play a decisive role in the system. Medical emergency management systems consider medical professionals a reliable external component in the control loop, which we call reliable-human-in-the-loop. In this scenario, the system may make recommendations, but the medical expert takes the responsibility, and the decision made by the system is Reversible. • **Metrics:** In the worst-case and average-case, the quality of computation results is vital. Though the medical experts will take the final responsibility, the systems are valuable only by providing high-quality and interpretable computation results. To provide real-time and safe-critical services, the worst-case performance is important besides the average performance, including the latency and throughput. Since the workloads are often spiky, the systems must gracefully handle overloading. Due to the specificity of medical care, security and privacy are always the first issues that medical systems have to consider. The future healthcare systems involve a more significant number and variety of devices, a larger population, and more applications, and its complexity brings more significant challenges to security and privacy.

- Various IoT devices generate a massive volume of heterogeneous data. Unlike traditional emergency medical care, modern elderly emergency medical care has been extended to their daily lives: before, during, and after the hospital. The emergency medical systems not only rely on a large number of different sensors (such as cameras, gyroscopes, blood pressure monitors, etc.) and also need to access various professional medical equipment (such as CT scanner, mass spectrometer, spirometer, etc.). These devices generate a large amount of heterogeneous data. The emergency medical care system needs to integrate and process large amounts of heterogeneous data in real-time to provide emergency medical services to the elderly throughout their life cycle.
- Computation patterns, computation dependencies, and interaction patterns:

The computation patterns follow the observe, fuse, recommend, and train patterns. The IoT devices observe the data of the patients at different levels. The system may fuse various observations at the edge or data center. The data centers or edges will train and update an AI model through the widely collected and labeled data. The IoT or edge makes a recommendation like alert and further-taken actions. The medical experts make a final decision.

The computation dependencies are mainly internal dependencies — the dependency between adjacent execution loops of observing, fusing, recommending, and training. For example, a patient's previous diagnosis and treatment recommendations would impact their subsequent recommendations. Besides, external dependencies exist between the decisions triggered by different IoTs, edges, data centers, or humans. Typical examples include new emergency cases reported by the IoTs, the newly trained models, and interventions brought out by the experts.

Over the other applications, the interaction patterns are simpler. Each IoT works within different conditions, and each computation may trigger different algorithms but only involve local data. Recommendations may be made at edges locally, involving collected data with different spatial-temporal scopes. When training the model, the data are widely collected from IoTs or edges and annotated at the data center for further training. Or in another manner, the labeled data at IoT or edges are distributed training using federation learning techniques [23,24]. Then lightweight models are deployed at IoT and edges.

• Data management and access pattern: Patients generate a large number of real-time data from various sensors and professional medical devices. The formats of patient data are diverse — image, video, relational data, XML, etc [25]. Patient data is written into the emergency medical care system in real-time. When an emergency medical event is detected, the emergency medical care system immediately initiates rapid diagnosis and develops rescue planning. Patient data helps medical experts understand how a critical health event unfolds, uncover the geographic characteristics of events, and locate the nearest medical resources. Patient data is accessed only when the patient experiences a medical emergency. Therefore, medical emergencies result in non-periodic random access patterns.

2.1.2. Autonomous driving

Autonomous driving is a promising technology that changes the way people travel. According to the standard of SAE International [26], autonomous driving is classified into six levels—"no driving automation



Fig. 2. The healthcare lifecycle of the elderly.

(Level 0), driver assistance (Level 1), partial driving automation (Level 2), conditional driving automation (Level 3), high driving automation (Level 4), and full driving automation (Level 5)" [26]. With its continuous development, self-driving cars will hit the roads and enter into a highly-automated era in the future [27].

- Task types: highly-automated and mainly safety-critical. The future autonomous driving would be highly-automated, even fully-automated, and consider no human in the control loop. The corresponding system needs to perceive and collect multi-source and multi-dimensional data in real-time and respond within several milliseconds. Considering the properties of high autonomy, hard real-time, and potentially destructive effects, the decision and action made by the system are irreversible. It will be a system failure in hard real-time when missing a deadline. In auto-driving, missing a deadline will be catastrophic.
- **Metrics:** In the worst-case and average-case, the quality of computation results is vital as no person takes responsibility. The systems must provide high-quality and interpretable computation results. The worst-case performance of autonomous driving is highly significant. For the safety of cars, pedestrians, and surroundings, an autonomous driving system is demanded to manage and coordinate massive self-driving cars synchronously, assure the performance of almost all vehicles and guarantee the worst-case performance make the tail latency as low as possible.

Security and privacy are critical challenges in autonomous driving. Besides the traditional security issue, security also means the car can perceive the environment, make decisions, and take actions correctly.

• Various IoT devices generate a massive volume of heterogeneous data. Autonomous driving depends on a large number of sensors; even a single car may deploy multiple kinds of sensors [28,29] including cameras, ultrasonic radar, millimeter-wave radar, lidar, IMU (Inertial Measurement Unit), etc., to obtain the environment information comprehensively. The input data are multi-source and heterogeneous; thus, the system should be able to fuse multi-sensor data for quick processing.

• Computation patterns, computation dependencies, and interaction patterns:

The computation patterns follow the observe, fuse, act, coordinate, and train patterns. The IoT devices observe and detect the data of the weather, surroundings, road lanes, traffic signs, pedestrians, other vehicles, etc. The system fuses various observations, makes a decision, and acts locally at the edge (within each car), including the control of the steering wheel, brake, speed, acceleration, and engine [30]. Meanwhile, vehicles, roads, and surroundings will synchronize their data with each other through data centers and finally coordinate their behaviors. In the background, different models are trained and updated regularly. The computations have severe internal and external dependencies. Internally, a self-driving car's current status and actions would impact its subsequent computations and actions through the observe, fuse, act, and coordinate loop. Externally, a selfdriving car's status and action would affect the behaviors of the other vehicles.

• Data management and access pattern: The autonomous driving system relies on a large number of sensors to provide comprehensive environmental information, such as traffic signs, pedestrians, and weather. The environment information is continuously written for trajectory planning, surrounding object detection, and autonomous control. The dataset includes LAS binary, float matrix, CSV file, etc [31–33]. Moreover, in the morning and evening rush hour, cars often need real-time path planning to avoid traffic congestion. The autonomous driving system needs to handle periodic burst accesses.

2.2. Why we need to build an HFC system?

2.2.1. The three unique requirements of HFC systems

Significantly different from the warehouse-scale computers that non-stop serve user requests [3], HFC systems have three unique requirements. First, they directly interact with the physical world considering humans as an essential part: human-in-the-loop [10], and perform safety-critical and mission-critical operations. Each action may have an irreversible effect. Some systems treat humans as "an external and unpredictable element in the control loop" [10], which we call unreliable-human-in-the-loop. For example, many security systems rely on a "human in the loop" to perform security-critical functions [34], but humans are incompetent and often fail in their security roles [34]; In contrast, the other systems consider humans, who make the final decision, a reliable component in the control loop, which we call reliable-human-in-the-loop. For example, medical expert plays a decisive role in medical emergency management systems. Meanwhile, more scenarios "bolster a closer tie with the human through human-in-theloop controls that consider human skills, intents, psychological states, emotions, and actions inferred through sensory data" [10], which we call collaborative-human-in-the-loop. Collaborative-human-in-the-loop indicates that humans are complements of the other components of the system. Still, an uncoordinated collaboration between a human being and other system components may result in disaster.

Second, unlike Internet services that process independent concurrent requests on planet-scale data center infrastructures, HFC computations have intertwined dependencies between not only adjacent execution loops but also actions or decisions triggered by IoTs, edge, data centers, or humans. Third, under this extremely entangled state, the systems must first satisfy the accuracy metric in predicting, interpreting, or taking action before meeting the performance goal under different conditions, like worst-case, average-case, and best-case.

We take autonomous driving as an example. The systems directly interact with the world. Each action has an irreversible effect. The current status and action of a self-driving car would impact its subsequent computations and actions; a self-driving car's status and action would affect the behaviors of the other vehicles. Autonomous driving must first ensure the accuracy of the decision (quality) and then guarantee the worst-case performance (tail latency).

2.2.2. HFCs demand resources that are several orders of magnitude beyond the reach of the state-of-the-practice systems

Even only taking the worst-case performance metrics - tail latency as examples, we illustrate that the state-of-the-practice systems are far from satisfying the processing requirements. Even for a much simpler application with simpler computation patterns (our motivating example) compared to those in Table 1 – a simplified smart home application with 95% queries reporting the status and sending heartbeat packets, and 5% queries processing user requests and accessing the Redis database, serving vast concurrent connections is still tricky. Zhang et al. [35] simulate this application using a million-level client load generator (MCC) and evaluate the service capacity of kernel TCP and mTCP v2.1 on an X86 server equipped with Intel Xeon E5645 processor, Centos 7.2, Kernel 3.10.0, and 64 GB memory. Taking 50 ms as the 99th percentile latency threshold, the kernel TCP supports one hundred thousand concurrent connections, while for the user-level mTCP network stack, the number is nine hundred and sixty thousand [35]. Accordingly, to achieve ten billion concurrent connections under the threshold of 50 ms, more than ten thousand nodes, even one hundred thousand nodes, would be needed.

Considering the three unique requirements discussed in Section 2.2.1 and the other factors, HFCs demand resources that are several orders of magnitude beyond the reach of the state-of-the-practice systems. The other factors considered in this rough estimation include scheduling, complex computation and interaction patterns, complex data access patterns, heterogeneous systems and networks, longer communication links, and differentiated processing abilities of IoTs, edges and data centers, which will aggravate the situation exponentially.

We further present several factors in detail. For example, as shown in Fig. 3, a lot of emerging and future applications usually require nearly perfect quality (i.e., predicting and interpreting accuracy in Autonomous driving) and the worst-case tail latency within several milliseconds, which are overlooked in the above motivating example.

From the perspective of task scheduling, different scheduling strategies may significantly impact performance. The previous experiments reveal that different placement and scheduling policies of data and workloads across IoTs, edges, and data centers may substantially affect the overall performance: even with the same infrastructure, the gap may achieve dozens or even hundreds of times considering the total response time [36]. The scheduling strategies are affected by multiple factors, including task complexity, device processing capability, available resources, network condition, pending tasks, etc. However, the state-of-the-practice solutions provide separate management of IoT, edge, and data center, lacking a global perspective, and further hardly to discover the optimal scheduling strategy. Consequently, unified management and efficient collaboration across IoTs, edges, and data centers are required to assure high performance and resource utilization.

Instead of simple status reports in our motivating example, HFC computations are much more complex and intricate, such as object recognition interference, OODA (observe, orient, decide, and act) [37] or even complex Avatar behavior in terms of machine learning or big data processing; An HFC system manages multi-source and heterogeneous data from various IoT devices, manifesting diverse data access patterns in real-time, random, burst, periodic, non-periodic, and batch manners, which also significantly impact the performance; IoT devices have a vast number that substantially outweighs the size of Internet users, with notable discrepant functions.

3. The challenges

This section lays out several HFC challenges.

Organizability and manageability challenges. Unlike a traditional computer system, e.g., a supercomputer or warehouse-scale computer, an HFC system is geographically distributed, consisting of IoTs, edges, data centers, and humans-in-the-loop. Moreover, they are dynamic. For example, in smart defense systems and applications [6],



Fig. 3. This figure compares the worst-case quality and performance requirements between the emerging and future applications in Table 1 and the traditional applications. The quality requirement indicates the required model accuracy of the application or tasks; for example, autonomous driving requires nearly 100% accuracy to assure safety. The worst-case performance requirement illustrates the required tail latency of the application. For example, autonomous driving requires extremely low tail latency within several milliseconds.

sensors or terminal control units can be dispersed by airdrop, inserted by artillery, and/or individually placed by an operation team [6]. In an extreme case, the spatial realm even has no bound. For example, unmanned spacecraft may have no bounded destination in interplanetary exploration applications.

So it is challenging to discover the resources and assemble them into a computer system. The challenges lie in managing these resources efficiently, keeping their survivability in the case of highly possible failures, guaranteeing the immunizability from malicious intrusions and attacks, and improving the programmability and schedulability of massive funclets across a large scale of IoTs, edges, and data centers.

Collaboration challenges between software, hardware and people components. More systems exhibit "a closer tie with the human through human-in-the-loop controls" [34]. However, not all these controls are reliable-human-in-the-loop. For example, some systems treat humans as "an external and unpredictable element in the control loop" [10]; thus, the action may have an irreversible effect and result in disaster. The challenge is (1) how do we handle the dilemma of choosing between the system and humans' decision, especially considering "human skills, intents, psychological states, emotions, and actions inferred through sensory data" [10]? (2) how do we support the collaboration between the system and humans? (3) how do we decide the respective responsibility in the partnerships? On the other hand, when the system makes a solo decision, especially for safety-critical functions or worst-case performance, how do we verify and validate its behavior? What is the human' responsibility behind the system? Let us look at these challenges from the perspective of intelligent defense systems described in Table 1. These challenges are not abstract but vivid and concrete in terms of casualties and losses of people and equipment.

Irreversible effect challenges. Most HFC systems or applications perform safety-critical and mission-critical operations, directly interacting with the physical world. Each action may have an irreversible effect under unreliable-human-in-the-loop and collaborative-human-in-the-loop conditions. Even with humans-in-the-loop, considering the

human's reaction time, it is hard to make a timely decision and action in an emergency. This irreversible effect demands that the systems' behaviors be verified and validated in advance; the systems can trace the impact of its attributing factors or causes or even achieve interpretability.

Most HFC systems need to explicitly state the quality of computation results and performance constraints in the entire process, including design, implementation, verification, and validation, referring to its target applications' correctness and performance constraints. For example, an autonomous driving application requires a worst-case design that assures high accuracy and tail latency within several milliseconds. In this case, we need to holistically verify and validate the systems and algorithms in terms of quality and performance in different cases like best-case, worst-case, and average-case. We have not gained enough experience in this regard.

Ecosystem wall challenge. An entire HFC ecosystem not only consists of the ensemble of the respective IoT, edge, and data center ecosystems but also involves multifarious meanwhile disparate technologies like processor design, operating system, toolchain, middleware, networking, etc. Moreover, the design of IoT, edge, and data centers follows distinct guidelines and targets according to their unique requirements or constraints, thus generating various ecosystems with different scopes and boundaries, which we call ecosystem walls. For example, the processor design of IoT pursues low energy consumption and a small chip area, while the data center regards performance as the first element. The daunting complexities caution us that we cannot reinvent the wheel. Instead, we need to be compatible with the ecosystem while improving the performance, energy efficiency, and other primary metrics to generate a positive change force that overcomes the ecosystem inertia [12].

The effective evaluation challenges. Generally, we need to deploy a system in a real-world environment and run a real-world application scenario to evaluate the performance and provide optimization guidelines. However, the real-world environment and emerging/future application scenarios are inaccessible and costly in assessing and verifying an HFC system. Hence, benchmarks as proxies of emerging/future application scenarios and simulators that emulate real-world systems are necessities for designing, evaluating and optimizing an HFC system.

- · Benchmarking challenges. The real-world emerging or future application scenarios are incredibly complex, involving the interconnection of IoT, edge, data center, and human-in-the-loop, and including intricate and lengthy execution path [38]. Constructing benchmarks for HFC systems faces significant challenges. First, the benchmarks should reflect the three unique requirements discussed in Section 2.2.1, which are not easily embodied in a benchmark manner. Second, the application designers and providers are concerned about the interaction, interconnection, task assignment, and end-to-end performance across IoT, edge, data centers, and human-in-the-loop. Thus, the reality of a benchmark is fundamental. Considering the complexity and confidential issues, the real-world application scenario is not fit to be used directly as a benchmark; hence, a simplified scenario is necessary. However, the real-world scenario contains hundreds or thousands of modules and components; even a tracing or monitoring tool can hardly figure out the execution path and call graph. Simplifying the real-world scenario while maintaining the critical parts is a big challenge. Moreover, considering the humans-in-theloop behaviors while constructing the benchmarks are a complex problem.
- Simulation and validation challenges. At the early stage of system and architecture evaluation, the simulator plays a vital role due to the vast manufacturing investment of time and money and the immaturity of the corresponding ecosystem. For example, the effectiveness of the improved processor design, memory access technologies, etc., is evaluated on a simulator. Considering the cost of building a real-world HFC system, a simulation or validation system supporting the whole environment simulation and technology verification is significant. However, the complexity and diversity of application scenarios pose substantial challenges in building such a simulator.

First, there has no unified interface for different application scenarios or architectures like IoT, edge, and data center. Thus, it is challenging to manage different architectures and support various scenarios.

Second, simulation accuracy is a crucial metric. High accuracy means the simulator can reflect similar running characteristics to the real world and exhibit running differences under different system environments. Considering the difficulties of multiple-level or multiple-scale simulation, including hardware level, e.g., processor chip, cache, memory hierarchy, disk, and software level, e.g., operation system, ensuring the simulation accuracy is necessary but challenging.

4. Exploring the solution space

We adopt a funclet methodology and architecture to facilitate exploring the HFC software and hardware design space. According to [12], the funclet represents "the common proprieties of basic building blocks" across the systems. Each funclet has the following characteristics [12]:"each contains a well-defined and evolvable functionality with modest complexity; each can be reusable in different contexts; each can be independently tested and verified before integrating; each can be independently deployable; each can interoperate with other funclets through a well-defined bus interface or interconnection".

The funclet architecture consists of four layers: chiplet, HWlet, envlet, and servlet. A chiplet is "an integrated circuit (IC) with modest complexity, providing well-defined functionality" [39,40]; it is "designed to be susceptible to integration with other chiplets, connected with a die-to-die interconnect scheme" [39,40]. A servlet is "an independently deployable and evolvable component that serves users with a well-defined and modest-complexity functionality" [12]. Microservices [41] or cloud functions [42] are two forms of servlet. An HWlet is "an independently deployable, replaceable, and accessible hardware component, e.g., CPU, memory, storage" [12]. An envlet is "an independently deployable and evolvable environment component with well-defined functionality that supports the management of servlets" [12].

The funclet architecture uses a three-tuple: {funclet set architecture (FSA), organization, system specifics} [12] methodology to describe the architecture. According to [12], the FSA refers to "the actual programmer-visible function set [43], serving as the boundary between two adjacent layers and among different funclets in the same layer"; The organization includes "the high-level aspects of how funclets in the same layer and adjacent layers collaborate"; The system specifics describe "the design and implementation of a system built from funclets" [12]. The advantages of the funclet architecture are discussed in [12]: it increases technology openness, improves productivity, and lowers cost; relieves the complexity of building systems; improves reusability and reliability.

Fig. 4 illustrates the reference HFC architecture and components. For the chiplet layer, we focus on workload-driven chiplet designs. For the HWlet layer, we pay attention to the message interface-based memory system and data path network processor. For the envlet layer, we concentrate on the HFC operating systems and the performancedeterministic distributed storage systems. Also, we provide several tools. A system design tool suite is provided to help designers explore the HFC design space across chiplet, HWlet, envlet, and servlet. A scenario simulator is built across IoTs, edges, and data centers to accelerate innovative technologies' deployment and verification; A full-stack optimization tool is provided. We construct a series of benchmarks and microservices for various HFC applications for the servlet layer. Several scenario benchmarks are proposed as the proxy of real-world application scenarios, aiming to support the whole-stack evaluation and provide feedback to scenario simulators or even real-world scenarios.

Fig. 5 shows a funclet-based HFC architecture in terms of {funclet set architecture (FSA), organization, system specifics}, which is derived from an open-source computer system initiative [12]. Each layer of chiplet, HWlet, envlet, and servlet specifies a set of *FSAs* for IoTs, edges, and data centers. The *FSAs* are composable and collaborative through the same-layer and adjacent layer. The organization specifies "the high-level aspects of how funclets in the same layer and adjacent layers collaborate" [12], including IoT–IoT, Edge–Edge, Datacenter–Datacenter, IoT–Edge, IoT–Datacenter, Edge–Datacenter, and IoT–Edge–Datacenter collaborations. The *system specifies* shows how to implement an HFC system. The different layers for IoTs, edges and data centers are interconnected through the funclet-based open standards, including interconnection, interface, protocol, and networking across IoTs, edges, data centers, and humans-in-the-loop.

4.1. Benchmarks

As the foundation of system design and optimization, benchmarks are of great significance for developing HFC systems. We aim to propose a series of benchmarks that reflect three unique characteristics and other important factors of HFC computations for designing and evaluating HFC systems.

Scenario benchmarks. From the perspective of the whole-stack system benchmarking, e.g., the interconnection and communication of IoTs, edges, data centers, and humans-in-the-loop, we propose scenario benchmarks as the proxy of the real-world application scenario. The construction follows a scenario-distilling methodology that formalizes a real-world application scenario as a graph model and distills it into a combination of essential tasks and components [38]. This methodology identifies "the critical path and primary modules of a real-world scenario since they consume the most system resources and are the core focuses for system design and optimization" [38], thus reducing complexity.



Fig. 4. The overview of the reference architecture.

IoT Benchmarks. From the perspective of middle-level modular benchmarking, we construct IoT benchmarks to evaluate mobile and embedded devices. The IoT benchmarks include the lightweight IoT workloads and light-weight AI workloads.

CPU Benchmarks. We present CPU benchmarks covering typical workloads from emerging and future application scenarios from the low-level architecture benchmarking perspective. Constructing CPU benchmarks adopts a traceable methodology, managing the traceable processes from problem definition, problem instantiation, solution instantiation, and measurement [44].

4.2. Chiplets

This subsection presents the workload-driven chiplet design methodology to explore the ideal architecture for each class of emerging and future applications. We aim to provide reusable building blocks considering the different PPA (performance, power, area) requirements for IoT, edge, and data center.

For the first step, we perform comprehensive workload characterization on a broad spectrum of tasks within target applications. For each application in Table 1, we analyze the computation patterns drilling down into the critical computations within the execution loop, like OODA in smart defense scenarios across IoT, edge, and data center. Then we analyze the interaction patterns covering the interactions between IoT–IoT, IoT–edge, IoT–datacenter, edge–edge, edge–datacenter, and datacenter–datacenter. Above all, the analysis contains computation, memory access, networking, and other characteristics. According to the results, on a single level of IoT, edge, or data center, we classify their characteristics into several classes for each pattern, like computation. After that, we will obtain several classes, including (IoT, computation), (IoT, memory access), (IoT, networking), (edge, computation), (edge, memory access), (edge, networking), (data center, BenchCouncil Transactions on Benchmarks, Standards and Evaluations 2 (2022) 100075

computation), (data center, memory access), (data center, networking), etc.

For the second step, we attempt to define the ideal chiplet architecture for different IoT, edge, data center layers and different analyzed patterns. Specifically, according to the classifications in Step 1, we define the computation, memory, networking chiplets for IoT, edge, and data center, respectively. Each layer of IoT, edge, or data center will contain multiple chiplets for different patterns of computation, memory access, networking, etc. Additionally, each pattern may contain multiple chiplet designs according to the classifications of workload characteristics.

For the third step, we validate the chiplet architecture design and further performs improvements according to the feedback. We adopt FPGA-based simulation and evaluate the scenario, IoT, and CPU benchmarks to conduct the functionality and performance validation. Further, we explore the upgrades and design optimizations based on the validation results.

The chiplet architecture design contains a loop of workload characterization, chiplet design, and validation until the output designs satisfy the application requirements.

4.3. Hwlets

This subsection presents the HWlets solutions, primarily focusing on two innovative HWlets: a message interface-based memory system and a data path network processor.

4.3.1. The message interface based memory system

As the boundary between internal and external memory is blurred, a computer system may face different memory devices with various latency, bandwidth, granularity, and capacity. Thus, the challenge is providing a universal memory interface and a unified memory system so that the programmers do not need to switch between bytelevel load/store CPU instructions (for internal memory) or block-level read/write I/O operations.

We have proposed a message-interface-based memory access approach to solve various "memory wall" problems [45]. We plan to extend the message interface to include internal and external memory to build a unified memory system. The system assumes a high bandwidth, high concurrency, and low latency network, which we believe will come soon.

Instead of only using a fixed command format and address for a memory request, we propose to use a message that contains rich semantics to express a memory request. The semantic information consists of size, sequence, priority, process id, persistence, etc., or even array, link pointer, and locks. Furthermore, the memory resource provider is no longer a simple dumb device but with different local computation capabilities to service a message request.

The message interface base memory system decouples the data access from the data organization. A client does not need to know the details or memory resource organization like banks and rows of an SDRAM, even the exact location of the data. The message interfacebased memory system also decouples the data access from data transfer. Small data requests can be combined into a large network message. A large data request can be divided into multiple messages.

There are three critical components to implementing a message interface-based memory system. (1) a CPU core generates concurrent memory requests with semantic information. Traditional load/storebased instructions are too simple to express a rich-semantic memory request. Instead, we need kinds of asynchronous and operand-variable instructions. (2) a memory controller with a message interface. The controller should group and assemble memory requests from internal CPU cores into coarse-grained messages to exchange with different memory servers through the network. (3) various message-interfaced memory servers. The memory server manages local storage media and accepts message requests and responses after specific local processing. We will implement the previous two components as chiplets and the message-interfaced memory servers as HWlets.



Fig. 5. An HFC instance of the four-layer funclet architecture based on [12].

4.3.2. The data path network processor

High-speed networks have been an indispensable part of modern computer systems. There are two distinct technical routes for network acceleration: offloading [46–48] and onloading [49]. Offloading means to offload part of the network processing to an external accelerator card, namely smart NICs [46,50], and save CPU resources for application logic. For example, the checksum of TCP packets can be computed on high-end NICs [48]. Although various smart NICs have been proposed, none of them has occupied the dominant position in the market. More recently, smart NICs have a new name — DPU (Data processing unit) [51]. On the other hand, onloading means pushing all network packet processing onto general-purpose CPU cores to utilize the everincreasing computing power of the CPU core. DPDK by Intel [52] is the industry standard for onloading with the help of hardware features built-in Intel CPUs.

Whether offloading or onloading, there are many functions in a network stack that can be accelerated by hardware, for example, checksum, encryption/decryption, table lookup, keyword matching, queuing, ordering, etc. Additionally, several functions can only be efficiently processed by the CPU, such as fragments, lists, buffers, order, and other complex data structures.

Hence, we argue that the critical point is organizing these accelerating resources efficiently. Equipping a general-purpose CPU with accelerators accessed via the I/O bus is not the most efficient solution. Likewise, putting a powerful general-purpose multi-core CPU into a NIC will not change much. The challenge here is how to design efficient control and data path of accelerators inside a CPU to combine both the accelerator units' special functions and the CPU cores' general processing ability.

For general-purpose processors, NIC is always an "external" device. The processor has to initialize a DMA operation to move data from the I/O bus to the local memory or cache to access packet data. Most built-in processors in smart NICs still have such structures. That results in uncontrollable processing latency, so these processors are only suitable for processing control paths. Only hardware logic like FPGA or special function-limited core like P4 engine can process linespeed data paths. They can reach the line speed only because their functions are simple and deterministic enough. Generally, they cannot process complex semantic information in data paths that need complex data structures to store the state of many concurrent transactions.

We propose to design a processor for line-speed processing data paths, which we call the data path network processor (a datapath processor). We will implement the datapath processor as an accelerator. However, the network packet will be the first-class citizen in the datapath processor. The packet stream leaves the register file of the primary CPU and arrives at the datapath processor directly without going through a complex memory hierarchy. The datapath processor has full functionality, including access to the cache and main memory. Thus, the datapath processor can hold and process complex state information necessary for the data plane. The datapath processor also has accelerating units on the local bus; data exchange between the datapath processor core and accelerator can be low latency, highly paralleled, and fine-grain. We have not seen such the structure of the datapath processor before, but fortunately, open-source processors, like RISC-V based, allow us to design a novel processor architecture freely.

4.4. Envlets

This subsection presents the Envlets solutions, primarily focusing on the HFC OS and the HFC distributed storage systems.

4.4.1. The HFC operating system

In the HFC scenarios, computing is ubiquitous, consisting of geographically distributed, heterogeneous hardware devices with different performance and power consumption constraints. Hence, we need a more efficient way to improve the organizability and manageability of the HFC systems. Our OS solutions aim to provide the following features:

(1) The new OS should have a flexible system structure. For different devices and workloads, specific OS capabilities need to be built. OS is no longer limited to a single kernel running directly in the local node but a distributed OS architecture that adapts to hardware resources in different computing nodes. OS needs to be able to reconfigure or rebuild itself in run-time to adapt to various scenarios.

(2) To efficiently manage the distributed heterogeneous hardware, the new OS should rebuild a general and intelligent device driving framework to discover, identify, register, access, and drive the massive hardware resources automatically. In addition, it can establish a soft bus connection for interactions with the immunizability from the malicious intrusions and attacks.

(3) To meet different HFC workloads' performance targets, the new OS needs to build fine-grained resource metering and application profiling features to facilitate efficient scheduling for improving performance and resource utilization.

(4) The HFC ecological boundary is open. As a result, OS faces security challenges, such as end-to-end device authentication access, identification issues in an open environment, and security isolation. Under the premise of "zero trust", we need to embody security enhancement strategies in the native OS kernel.

4.4.2. The performance-deterministic distributed storage systems

The emerging and future applications heavily rely on advanced techniques like big data and AI, performing hybrid and concurrent tasks with different requirements, e.g., latency-critical and throughputcritical, and pursuing the worst, average, or best-case performance. However, these hybrid tasks usually adopt different systems and architectures, and have distinct data access patterns and requirements [53, 54]. In addition, the worst-case tail latency poses great challenges even for Internet services in data centers [55–57], let alone much more complex applications across IoTs, edges, and data centers. Thus, to efficiently serve these applications and tasks, building a single distributed storage system (DSS) that provides *deterministic performance and high throughput* is an urgent demand [58–60]. Note that the deterministic performance means that a DSS should enforce differentiated tail latency SLOs for concurrent latency-critical tasks. Throughput means the total QPS (requests per second) or bandwidth of a DSS.

The design of a DSS needs to consider the characteristics of storage devices. We conclude two development trends of storage devices. On the one hand, the devices will be increasingly faster with microsecondscale or even lower latency. Storage devices have experienced several technological breakthroughs in the past twenty years, such as the development of commercial SSD products, NVM-based SSD products (Intel Optane SSD [61]), and persistent memory products (Intel Optane PM [62]). Compared to HDD and ordinary SSD, NVM-base devices have much lower latency. In addition, emerging fast networks (e.g., 200 Gbps and 400 Gbps Infiniband) have round-trip latency of less than 1 µs [63]. These low latency devices put forward high demand to the storage systems [64]. On the other hand, the devices will contain enhanced computation capacities, such as computational storage drives [65,66], SmarkNICs [67], and programmable switches [68]. Many studies propose to offload several tasks to the devices and have shown the performance advantages, like offloading query processing and data (de-)compression to SmartSSDs [69-72], data replication and file system functions to SmartNIC [73,74], and global memory management, load balance and data cache to programmable switches [75-77]. Hence, the design of a DSS needs to make full use of these in-device computing resources.

Considering the application requirements and device characteristics, building a DSS faces serious challenges. First, due to the distinct states of different machines/threads, latency spikes [55,78], schedulability issues [79], load burst [80-83], and resource contention inside storage devices and the network stack [84-86], it is extremely hard to guarantee deterministic performance and high throughput. Second, many technologies have been proposed to achieve low latency, including using poll instead of interrupts [87], kernel-bypass I/O like user-space communication mechanism [52,88], user-space device drivers [89,90], and user-space file systems [91,92]. However, the previous DSS only adopt a single technology, and it is challenging to integrate and benefit from all these technologies in a single DSS. Third, there are increasing computation capacities inside devices through either FPGA or lowenergy embedded processors. Previous work attempts to offload partial computation to in-device computing logic [93-97]. We argue that offloading partial infrastructure software like DSS [74,98] and SQL engine [69,71] rather than user applications are better [74]. However, due to the complicated functionalities of DSS, it is a tough thing.

Our solutions for a novel DSS include the following innovations. First, a new DSS should exploit the leading technologies for emerging devices with μ s-scale latencies, including polling device events, user-space I/O, and run-to-completion request processing. Existing efforts focus on one individual technology and fail to integrate all of them

into a single system in a systematic way for efficiency. Second, a new DSS should embed scheduling along the whole I/O path from the clients to the servers and storage devices. Schedulable architectures should be used at each layer along the I/O path, including client-side, network, server-side, and storage devices. Moreover, exploiting in-device computing power through software–hardware co-design controls resource utilization at the fine granularity and reduces latency. Third, a new DSS should offload some functionalities to the devices without reducing total throughput and impairing the performance of individual applications. Existing efforts exploit one type of in-device computing power, either computational storage, SmartNIC, or programmable switches. Unlike them, a new DSS should make full use of all these in-device computing powers to achieve better performance and energy efficiency.

4.5. Tools

This section presents our tools: the system design tools, the full-stack optimization tool, and the scenario simulator.

4.5.1. The system design tools

We propose system design tools to help designers explore the HFC design space. The system design tools include the chiplet design tool, the HWlet design tool, the envlet design tool, and the servlet design tool, which correspond to the funclet architecture. Each design tool provides the simulation-validation-development tool suite. For example, for the chiplet design tool, we propose a whole-picture simulation to explore the co-design space of the chiplet across stacks. Unlike the traditional microarchitecture-level simulation, such as the GEM5, our whole-picture simulation is across full system stack levels. including three hierarchical levels: the IR (intermediate representation) level, ISA, and microarchitecture levels [99]. The whole-picture simulation combines the IR, ISA, and microarchitecture level simulations. The design decisions at IR, ISA, and microarchitecture levels are ISAindependent, microarchitecture-independent, or specific to the actual processor's microarchitecture. Combining the design decisions from the IR, ISA, and microarchitecture, the user can explore the co-design space across stacks. Furthermore, we propose cycle-accurate and bit-accurate circuit simulations and verification tools for the validation. The validation tool is based on general x86 computing and heterogeneous FPGA resources and provides an on-demand service for chiplet validation. For the development, we propose AI-based open-sourced EDA tools for integrated circuit design and development.

4.5.2. The full-stack optimization tool

The full-stack optimization tool has the following challenges:

(1) The optimization object is uncertain. Finding the performance bottleneck in an HFC system is non-trivial. Users may feel confused about the optimization objects because of the optimization possibilities on IoTs, edges, or data centers and the complex hierarchies of algorithms, frameworks, software, and hardware.

(2) The optimization space is vast. There are thousands of optimization dimensions of the algorithm, software, and hardware, and the values of the variate vary in an extensive range. As a result, the optimization space is exceptionally huge.

(3) The optimization target is diverse. Different application scenarios have additional user requirements. In addition to the vital importance of accuracy, some applications are sensitive to latency; some require high throughput, and some are concerned with energy consumption. Therefore, different application scenarios have other optimization goals.

The tool covers the optimization from vertical and horizontal dimensions. From the vertical dimension, we will co-explore the optimization space from the algorithm, software, and hardware. For example, for deep learning applications, jointly optimizing the network's architecture and the hardware accelerators is promising in improving performance and reducing energy consumption. We will consider the close collaboration among IoTs, Edges, data centers, and humans-in-the-loop from the horizontal dimension. For example, we will automatically offload whole or partial deep neural network computations from end devices to more powerful devices, such as edges or data centers.

Automatically co-optimization is non-trivial because of the vast optimization space. There are thousands of optimization dimensions of the algorithm, software, and hardware, and the values of the variate vary in an extensive range. As a result, the optimization space is too huge to complete the search. Reinforcement learning has shown powerful capabilities for the problem of searching for optimal policies in a vast space. Evaluating is expensive in co-optimizing the algorithm, software, and hardware across the IoTs, edges, and data center. We will investigate the state-of-the-art learning algorithms and evaluation strategies and develop the corresponding tools for automatic optimization.

4.5.3. The scenario simulator

The scenario simulator is a miniature of the real system, which contains unified interfaces and replaceable components to enable rapid deployment and verification of innovative technologies. The scenario simulator manages the whole environment of a computer system, e.g., processor chip, operating system, memory, network, etc. It covers the complete execution and interaction across IoTs, edges, data centers, and humans-in-the-loop. It can demonstrate the effects visually, e.g., running results, performance, and power consumption, under different technologies, deployments, or parameter settings, for example, the latency performance of an autodrive scenario using other memory devices and network protocols. For the first step, we plan to provide a network simulator that simulates the communication patterns of representative scenarios like big data and artificial intelligence, supporting different networking technologies. The involved components are replaceable, and we can easily use emerging technology to replace the existing one and verify its effectiveness. We will expand the scenario simulator to the whole HFC environment for the next step.

5. Our plan

We aim to define a new paradigm — IoTs, edges, data centers, and humans-in-the-loop as a computer and launch an open-source high fusion computer (HFC) system initiative. The goal is to vastly enhance the system capabilities under specific energy and cost constraints for most emerging and future applications.

We abstract reusable functions (funclets) across system stacks among IoTs, edges, and data centers to guide the HFC system design and evaluation. We first propose to define a series of benchmarks and funclet-based standards and then build the tools to facilitate the workload-driven exploration of the system and architecture design space. Finally, we provide open-source implementations of an HFC system. We will perform system co-design from the vertical and horizontal dimensions throughout the process. Vertically, we comprehensively explore the algorithms, runtime systems, resource management, storage, memory, networking, and chip technologies. Horizontally, we deeply discover the collaboration and interaction among IoT, edges, data centers, and humans-in-the-loop.

We plan to build the first open-source implementation of an HFC system using an iterative and evolving way. Fig. 6 shows the development milestone of our HFC system. We first focus on one or two typical application scenarios and essential funclets across IoT, edge, and data centers to reduce the complexity. Then we expand the focus and update or replace the technologies gradually. A scenario simulator is beneficial to the expansion, update, and replacement. Finally, we will summarize the experience and lessons during this period and dedicate ourselves to the contributions of a useful HFC system and related advanced technologies.

6. Related work

The development of the computer industry witnessed a series of computer systems concepts and implementations, as shown in Fig. 7.

In 1936, Turing proposed to invent the single machine to compute any computable sequence, a concept of a "universal" computing device [100,101]. This kind of "universal" would incur additional performance, cost, or energy overhead called "Turing Tax" – a fundamental question that computer architects aim to reduce [100]. John Gage first proposed the phrase "the network is the computer" in 1984 [102,103]. In 1985, Lewis proposed the concept of the "Internet of things" in a speech to the Congressional Black Caucus Foundation 15th Annual Legislative Weekend [104]. In the mid-1990s, grid computing was proposed to provide computing power, data, and software on-demand, through standardizing the protocols [2]. In 2001, EmNets [6] were proposed and referred to as networked systems of embedded computers. In 2005, cyber-physical systems were proposed to "bridge the cyber-world of computing and communications with the physical world" [105-107]. In 2009, Google proposed the concept of "the data center as a computer", or called warehouse-scale computers (WSCs), to efficiently deliver good levels of Internet service performance [108]. In 2009, the Chinese Academy of Sciences predicted that humancyber-physical ternary computing would be a development trend in the next 50 years [109]. In 2012, the "Industrial Internet of Things (IIoT)" concept, also known as "Industrial Internet", was proposed to integrate the latest technologies, intelligence systems, and devices and apply them to the entire industrial economy [110,111]. In 2016, the director of Storage SRE at Google illustrated how they do planet-scale engineering for a planet-scale infrastructure — keep all its services up and running and reduce the downtime [3]. In 2017, Li et al. pointed out that human-cyber-physical ternary intelligence is the leading technology and the main driving force of the new economy in the next 15-20 years [112]. In 2021, Mike Warren's group worked with the EC2 team. launching a virtual supercomputer in the cloud – 4.096 EC2 instances with 172,692 cores. This run achieved 9.95 PFLOPS (actual performance), ranking at 40th on the June 2021 TOP500 list [4]. In 2021 and 2022, Wang et al. and Xu et al. pointed out "a new era of human-cyber-physical ternary computing with diverse, intelligent applications over trillions of devices" [7,8] and further proposed the concept of "Information Superbahn", to achieve high system goodput and application quality of service [7,8].

7. Conclusion

We call attention to the fact that more and more emerging and future applications rely heavily upon systems consisting of Internet of Things (IoT), edges, data centers, and humans-in-the-loop. We characterized this new class of systems and coined a new term, high fusion computers (HFCs), to describe them. Significantly different from warehouse-scale computers that non-stop serve independent concurrent user requests, HFCs directly interact with the physical world, considering humans an essential part and performing safety-critical and mission-critical operations; their computations have intertwined dependencies between not only adjacent execution loops but also actions or decisions triggered by IoTs, edge, data centers, or humans-in-theloops; the systems must first satisfy the accuracy metric in predicting, interpreting, or taking action before meeting the performance goal under different cases. HFCs raise severe challenges in system evaluation, design, and implementation.

We summarize several HFC challenges: organizability and manageability, collaborations between software, hardware, and people components, irreversible effect, ecosystem wall, and effective evaluation. To tackle the above challenges, we propose reconstructing IoTs, edges, data centers, and humans-in-the-loop as a computer rather than a distributed system; we adopt a funclet methodology of building large systems out of smaller functions and exploring HFC design space in a structural manner. We will provide the first open-source implementation of the funclet architecture of HFC systems. The source code will be publicly available from the project homepage: https://www. computercouncil.org/HFC/.



Fig. 6. Development milestone of HFC system.



Fig. 7. An overview of the related work.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Dr. Jianfeng Zhan contributes the concept of the IoTs, edges, data centers, and humans-in-the-loop as a computer. Dr. Yungang Bao proposes considering networking as a critical part. Dr. Jianfeng Zhan and Dr. Lei Wang co-coin the term: high fusion computers (HFCs). Dr. Jianfeng Zhan contributes the abstract and Section 1. Section 2.1.1 is contributed by Dr. Jianfeng Zhan, Dr. Yunyou Huang, and Mrs. Guoxin Kang. Section 2.1.2 is contributed by Dr. Wanling Gao, Dr. Jianfeng Zhan, and Mrs. Guoxin Kang. Dr. Wanling Gao contributes the analysis of autonomous driving, civil aviation safety regulation, and interplanetary explorations in Table 1. Dr. Jianfeng Zhan contributes the analysis of medical emergency management, autonomous driving, and smart defense systems in Table 1. Dr. Yunyou Huang contributes the analysis of medical emergency management in Table 1. Dr. Chunjie Luo and Mr. Hainan Ye contribute the analysis of Metaverse in Table 1. Dr. Lei Wang contributes the analysis of digital twin in Table 1. Dr. Weiping Li contributes the analysis of civil aviation safety regulation in Table 1. Mrs. Guoxin Kang contributes all the analysis of data management and access patterns in Table 1. Section 2.2.1 is contributed by Dr. Jianfeng Zhan and Dr. Wanling Gao. Section 2.2.2 is contributed by Dr. Wanling Gao, Dr. Jianfeng Zhan, and Dr. Wenli Zhang. Section 3 is contributed by Dr. Jianfeng Zhan and Dr. Wanling Gao. Dr. Yungang Bao contributed the organizability and manageability challenge. Section 4.1 is contributed by Dr. Wanling Gao and Dr. Jianfeng Zhan. Section 4.2 is contributed by Dr. Wanling Gao, Dr. Jianfeng Zhan, and Dr. Biwei Xie. Section 4.3 is contributed by Dr. Mingyu Chen. Section 4.4.1 is contributed by Dr. Chen Zheng and Dr.

Lei Wang. Section 4.4.2 is contributed by Dr. Jin Xiong. Section 4.5.1 is contributed by Dr. Lei Wang. Section 4.5.2 is contributed by Dr. Chunjie Luo and Dr. Biwei Xie. Section 4.5.3 is contributed by Dr. Wanling Gao. Section 5, Section 6, and Section 7 are contributed by Dr. Jianfeng Zhan and Dr. Wanling Gao. Figs. 1, 3, 6, 7 is contributed by Dr. Wanling Gao. Fig. 2 is contributed by Dr. Yunyou Huang. Fig. 4 is contributed by Dr. Lei Wang, Mr. Shaopeng Dai, and Mr. Qian He. Fig. 5 is contributed by Dr. Wanling Gao, Dr. Chunjie Luo, and Mr. Qian He. Dr. Jianfeng Zhan, Dr. Wanling Gao, and Dr. Lei Wang proofread throughout the paper. We are very grateful to Dr. Sa Wang, Dr. Yisong Chang, Dr. Di Zhao, Dr. Mi Zhang, Mrs. Fan Zhang, and Mr. Hongxiao Li for their discussions and contributions.

References

- [1] Ian Sommerville, nineth ed., Addison-Wesley, 2011.
- [2] Ian Foster, Yong Zhao, Ioan Raicu, Shiyong Lu, Cloud computing and grid computing 360-degree compared, in: 2008 Grid Computing Environments Workshop, IEEE, 2008, pp. 1–10.
- Melissa Binde, How Google does planet-scale engineering for planet-scale infrastructure, 2016, https://www.youtube.com/watch?v=H4vMcD7zKM0.
- [4] Jeff Barr, Planetary-scale computing 9.95 PFLOPS & position 40 on the TOP500 list, 2021, https://aws.amazon.com/cn/blogs/aws/planetary-scalecomputing-9-95-pflops-position-41-on-the-top500-list/.
- [5] Scott Burleigh, Vinton Cerf, Robert Durst, Kevin Fall, Adrian Hooke, Keith Scott, Howard Weiss, The InterPlaNetary internet: A communications infrastructure for Mars exploration, Acta Astronaut. 53 (4–10) (2003) 365–373.
- [6] National Research Council, et al., Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers, National Academies Press, 2001.
- [7] Zhiwei Xu, Guojie Li, Ninghui Sun, Information superbahn: Towards new type of cyberinfrastructure, Bull. Chin. Acad. Sci. (Chinese Version) 37 (1) (2022) 46–52.
- [8] Xiaohong Wang, Sa Wang, Hongwei Tang, Xiaohui Peng, Building substrate for national strategy of new infrastructure construction-practice and thought of information superbahn testbed, Bull. Chin. Acad. Sci. (Chinese Version) 36 (9) (2021) 1066–1073.

- [9] Microsoft, A planetary computer for a sustainable future, 2022, https:// planetarycomputer.microsoft.com/.
- [10] David Sousa Nunes, Pei Zhang, Jorge Sá Silva, A survey on human-in-the-loop applications towards an internet of all, IEEE Commun. Surv. Tutor. 17 (2) (2015) 944–965.
- [11] Gordon E. Moore, et al., Cramming More Components onto Integrated Circuits, McGraw-Hill, New York, 1965.
- [12] Jianfeng Zhan, Open-source computer systems initiative: The motivation, essence, challenges, and methodology, BenchCouncil Trans. Benchmarks Standards Eval. 2 (1) (2022) 100038.
- [13] Qi Chen, Airborne lidar data processing and information extraction, Photogramm. Eng. Remote Sens. 73 (2) (2007) 109.
- [14] The World Bank, Population ages 65 and above, total, 2021, https://data. worldbank.org/indicator/SP.POP.65UP.TO. (Last Accessed on 4 February 2022).
- [15] U.N. Desa, World Population Prospects 2019: Highlights, Vol. 11, no. 1, United Nations Department for Economic and Social Affairs, New York (US), 2019, p. 125
- [16] Eduardo Teixeira, Hélder Fonseca, Florêncio Diniz-Sousa, Lucas Veras, Giorjines Boppre, José Oliveira, Diogo Pinto, Alberto Jorge Alves, Ana Barbosa, Romeu Mendes, et al., Wearable devices for physical activity and healthcare monitoring in elderly people: A critical review, Geriatrics 6 (2) (2021) 38.
- [17] Debajyoti Pal, Suree Funilkul, Nipon Charoenkitkarn, Prasert Kanthamanon, Internet-of-Things and smart homes for elderly healthcare: An end user perspective, IEEE Access 6 (2018) 10483–10496.
- [18] JeongGil Ko, Jong Hyun Lim, Yin Chen, Rvăzvan Musvaloiu-E, Andreas Terzis, Gerald M Masson, Tia Gao, Walt Destler, Leo Selavo, Richard P Dutton, MEDiSN: Medical emergency detection in sensor networks, ACM Trans. Embedded Comput. Syst. (TECS) 10 (1) (2010) 1–29.
- [19] Sathiyabhama Balasubramaniam, Rajeswari Kurubarahalli Chinnasamy, IoTbased noninvasive wearable and remote intelligent pervasive healthcare monitoring systems for the elderly people, Intell. Pervasive Comput. Syst. Smarter Healthc. (2019) 141–158.
- [20] Adriana Alexandru, Dora Coardos, Eleonora Tudora, IoT-Based healthcare remote monitoring platform for elderly with fog and cloud computing, in: 2019 22nd International Conference on Control Systems and Computer Science, CSCS, IEEE, 2019, pp. 154–161.
- [21] Sumit Majumder, Emad Aghayi, Moein Noferesti, Hamidreza Memarzadeh-Tehran, Tapas Mondal, Zhibo Pang, M. Jamal Deen, Smart homes for elderly healthcare-recent advances and research challenges, Sensors 17 (11) (2017) 2496.
- [22] Ying Liu, Lin Zhang, Yuan Yang, Longfei Zhou, Lei Ren, Fei Wang, Rong Liu, Zhibo Pang, M Jamal Deen, A novel cloud-based framework for the elderly healthcare services using digital twin, IEEE Access 7 (2019) 49088–49101.
- [23] Yuan Liang, Yange Guo, Yanxia Gong, Chunjie Luo, Jianfeng Zhan, Yunyou Huang, FLBench: A benchmark suite for federated learning, in: BenchCouncil International Federated Intelligent Computing and Block Chain Conferences, Springer, 2020, pp. 166–176.
- [24] Yunyou Huang, Xiuxia Miao, Ruchang Zhang, Li Ma, Wenjing Liu, Fan Zhang, Xianglong Guan, Xiaoshuang Liang, Xiangjiang Lu, Suqing Tang, et al., Training, testing and benchmarking medical AI models using clinical AIBench, BenchCouncil Transactions Benchmarks Standards Eval. (2022) 100037.
- [25] Tom J. Pollard, Alistair E.W. Johnson, Jesse D. Raffa, Leo A. Celi, Roger G. Mark, Omar Badawi, The eICU collaborative research database, a freely available multi-center database for critical care research, Sci. Data 5 (1) (2018) 1–13.
- [26] Sergej S. Shadrin, Anastasiia A. Ivanova, Analytical review of standard Sae J3016 "taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles" with latest updates, Avtomobil'. Doroga. Infrastruktura. 3 (21) (2019) 10.
- [27] NTT DATA, When the car takes over A glimpse into the future of autonomous driving, 2020, https://us.nttdata.com/en/-/media/assets/whitepaper/mfg-autonomous-cars-white-paper.pdf.
- [28] De Jong Yeong, Gustavo Velasco-Hernandez, John Barry, Joseph Walsh, Sensor and sensor fusion technology in autonomous vehicles: A review, Sensors 21 (6) (2021) 2140.
- [29] Jorge Vargas, Suleiman Alsweiss, Onur Toker, Rahul Razdan, Joshua Santos, An overview of autonomous vehicles sensors and their vulnerability to weather conditions, Sensors 21 (16) (2021) 5397.
- [30] Jelena Kocić, Nenad Jovičić, Vujo Drndarević, An end-to-end deep neural network for autonomous driving designed for embedded automotive platforms, Sensors 19 (9) (2019) 2064.
- [31] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Scharwächter, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, Bernt Schiele, The cityscapes dataset, in: CVPR Workshop on the Future of Datasets in Vision, Vol. 2, sn, 2015.
- [32] Will Maddern, Geoffrey Pascoe, Chris Linegar, Paul Newman, 1 year, 1000 km: The Oxford RobotCar dataset, Int. J. Robot. Res. 36 (1) (2017) 3–15.
- [33] Andreas Geiger, Philip Lenz, Christoph Stiller, Raquel Urtasun, Vision meets robotics: The kitti dataset, Int. J. Robot. Res. 32 (11) (2013) 1231–1237.
- [34] Lorrie F. Cranor, A Framework for Reasoning About the Human in the Loop, Advanced Computing Systems Professional and Technical Association, 2008.

- [35] Wenli Zhang, Ke Liu, Yifan Shen, Yazhu Lan, Hui Song, Mingyu Chen, Yuanfei Chen, Labeled network stack: A high-concurrency and low-tail latency cloud server framework for massive IoT devices, J. Comput. Sci. Tech. 35 (1) (2020) 179–193.
- [36] Tianshu Hao, Jianfeng Zhan, Kai Hwang, Wanling Gao, Xu Wen, AI-oriented workload allocation for cloud-edge computing, in: 2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing, CCGrid, IEEE, 2021, pp. 555–564.
- [37] Chet Richards, Boyd's OODA Loop, FHS, Sjøkrigsskolen, 2020.
- [38] Wanling Gao, Fei Tang, Jianfeng Zhan, Xu Wen, Lei Wang, Zheng Cao, Chuanxin Lan, Chunjie Luo, Xiaoli Liu, Zihan Jiang, Aibench scenario: Scenariodistilling ai benchmarking, in: 2021 30th International Conference on Parallel Architectures and Compilation Techniques, PACT, IEEE, 2021, pp. 142–158.
- [39] Tao Li, Jie Hou, Jinli Yan, Rulin Liu, Hui Yang, Zhigang Sun, Chiplet heterogeneous integration technology-status and challenges, Electronics 9 (4) (2020) 670.
- [40] Mark Lapedus, The Good and Bad of Chiplets, 2020, https://semiengineering. com/the-good-and-bad-of-chiplets/.
- [41] Irakli Nadareishvili, Ronnie Mitra, Matt McLarty, Mike Amundsen, Microservice Architecture: Aligning Principles, Practices, and Culture, "O'Reilly Media, Inc.", 2016.
- [42] Eric Jonas, Johann Schleier-Smith, Vikram Sreekanti, Chia-Che Tsai, Anurag Khandelwal, Qifan Pu, Vaishaal Shankar, Joao Carreira, Karl Krauth, Neeraja Yadwadkar, et al., Cloud programming simplified: A berkeley view on serverless computing, 2019, arXiv preprint arXiv:1902.03383.
- [43] John L. Hennessy, David A. Patterson, Computer Architecture: A Quantitative Approach.
- [44] Jianfeng Zhan, A BenchCouncil view on benchmarking emerging and future computing, BenchCouncil Trans. Benchmarks Standards Eval. 2 (2) (2022) 100064.
- [45] Licheng Chen, Mingyu Chen, Yuan Ruan, Yongbing Huang, Zehan Cui, Tianyue Lu, Yungang Bao, MIMS: Towards a message interface based memory system, J. Comput. Sci. Tech. 29 (2) (2014) 255–272.
- [46] Ming Liu, Tianyi Cui, Henry Schuh, Arvind Krishnamurthy, Simon Peter, Karan Gupta, Offloading distributed applications onto smartnics using ipipe, in: Proceedings of the ACM Special Interest Group on Data Communication, 2019, pp. 318–333.
- [47] Yanfang Le, Hyunseok Chang, Sarit Mukherjee, Limin Wang, Aditya Akella, Michael M Swift, TV Lakshman, UNO: Uniflying host and smart NIC offload for flexible packet processing, in: Proceedings of the 2017 Symposium on Cloud Computing, 2017, pp. 506–519.
- [48] YoungGyoun Moon, SeungEon Lee, Muhammad Asim Jamshed, KyoungSoo Park, {AccelTCP}: Accelerating network applications with stateful {TCP} offloading, in: 17th USENIX Symposium on Networked Systems Design and Implementation, NSDI 20, 2020, pp. 77–92.
- [49] Matthew GF Dosanjh, Ryan E Grant, Patrick G Bridges, Ron Brightwell, Reevaluating network onload vs. offload for the many-core era, in: 2015 IEEE International Conference on Cluster Computing, IEEE, 2015, pp. 342–350.
- [50] Yan Yan, Arash Farhadi Beldachi, Reza Nejabati, Dimitra Simeonidou, P4enabled smart nic: Enabling sliceable and service-driven optical data centres, J. Lightwave Technol. 38 (9) (2020) 2688–2694.
- [51] Idan Burstein, Nvidia data center processing unit (DPU) architecture, in: 2021 IEEE Hot Chips 33 Symposium, HCS, IEEE, 2021, pp. 1–20.
- [52] DPDK, Data plane development kit, 2022, https://www.dpdk.org/.
- [53] Fengfeng Pan, Yinliang Yue, Jin Xiong, Daxiang Hao, I/O characterization of big data workloads in data centers, in: Big Data Benchmarks, Performance Optimization, and Emerging Hardware, Cham, 2014, pp. 85–97.
- [54] Steven W.D. Chien, Stefano Markidis, Chaitanya Prasad Sishtla, Luis Santos, Pawel Herman, Sai Narasimhamurthy, Erwin Laure, Characterizing deeplearning I/O workloads in TensorFlow, in: 2018 IEEE/ACM 3rd International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems, PDSW-DISCS, IEEE, 2018, pp. 54–63.
- [55] Jeffrey Dean, Luiz André Barroso, The tail at scale, Commun. ACM 56 (2) (2013) 74–80.
- [56] Aad Van Moorsel, Metrics for the internet age: Quality of experience and quality of business, in: Fifth International Workshop on Performability Modeling of Computer and Communication Systems, Vol. 34, no. 13, Citeseer, 2001, pp. 26–31.
- [57] Tobias Hobfeld, Raimund Schatz, Martin Varela, Christian Timmerer, Challenges of QoE management for cloud applications, IEEE Commun. Mag. 50 (4) (2012) 28–36.
- [58] Andrew Wang, Shivaram Venkataraman, Sara Alspaugh, Randy Katz, Ion Stoica, Cake: Enabling high-level SLOs on shared storage systems, in: Proceedings of the Third ACM Symposium on Cloud Computing, SoCC '12, New York, NY, USA, 2012.
- [59] Ana Klimovic, Heiner Litz, Christos Kozyrakis, ReFlex: Remote flash ≈ local flash, ACM SIGARCH Comput. Archit. News 45 (1) (2017) 345–359.
- [60] Liuying Ma, Zhenqing Liu, Jin Xiong, Dejun Jiang, QWin: Enforcing tail latency SLO at shared storage backend, 2021, arXiv preprint arXiv:2106.09206.

- [61] Intel, Intel optane[™]DC SSD series, 2020, https://www.intel.com/content/ www/us/en/products/details/memory-storage/data-center-ssds/optane-dc-ssdseries.html.
- [62] Intel, Intel optane[™] persistent memory, 2020, https://www.intel.com/ content/www/us/en/products/details/memory-storage/optane-dc-persistentmemory.html.
- [63] Timothy Prickett Morgan, InfiniBand innovation is about more than bandwidth and latency, 2021, https://www.nextplatform.com/2021/09/30/infinibandinnovation-is-about-more-than-bandwidth-and-latency/.
- [64] Luiz Barroso, Mike Marty, David Patterson, Parthasarathy Ranganathan, Attack of the killer microseconds, Commun. ACM 60 (4) (2017) 48–54.
- [65] Samsung, SmartSSD, 2000, https://semiconductor.samsung.com/ssd/smart-ssd/.
- [66] ScaleFlux, ScaleFlux computational storage drive, 2022, https://www.scaleflux. com/.
- [67] Mellanox Technologies, NVIDIA mellanox BlueField SmartNIC for InfiniBand & ethernet, 2020, https://network.nvidia.com/files/doc-2020/pb-bluefield-vpismart-nic.pdf.
- [68] Intel, Intel Tofino[™] programmable ethernet switch ASIC, 2022, https: //www.intel.com/content/www/us/en/products/network-io/programmableethernet-switch/tofino-series.html.
- [69] Jaeyoung Do, Yang-Suk Kee, Jignesh M Patel, Chanik Park, Kwanghyun Park, David J DeWitt, Query processing on smart SSDS: Opportunities and challenges, in: Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, 2013, pp. 1221–1230.
- [70] Boncheol Gu, Andre S. Yoon, Duck-Ho Bae, Insoon Jo, Jinyoung Lee, Jonghyun Yoon, Jeong-Uk Kang, Moonsang Kwon, Chanho Yoon, Sangyeun Cho, et al., Biscuit: A framework for near-data processing of big data workloads, ACM SIGARCH Comput. Archit. News 44 (3) (2016) 153–165.
- [71] Insoon Jo, Duck-Ho Bae, Andre S Yoon, Jeong-Uk Kang, Sangyeun Cho, Daniel DG Lee, Jaeheon Jeong, YourSQL: A high-performance database system leveraging in-storage computing, Proc. VLDB Endow. 9 (12) (2016) 924–935.
- [72] Mohammadamin Ajdari, Pyeongsu Park, Joonsung Kim, Dongup Kwon, Jangwoo Kim, CIDR: A cost-effective in-line data reduction system for terabitper-second scale SSD arrays, in: 2019 IEEE International Symposium on High Performance Computer Architecture, HPCA, IEEE, 2019, pp. 28–41.
- [73] Daehyeok Kim, Amirsaman Memaripour, Anirudh Badam, Yibo Zhu, Hongqiang Harry Liu, Jitu Padhye, Shachar Raindel, Steven Swanson, Vyas Sekar, Srinivasan Seshan, Hyperloop: Group-based NIC-offloading to accelerate replicated transactions in multi-tenant storage systems, in: Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, 2018, pp. 297–312.
- [74] Jongyul Kim, Insu Jang, Waleed Reda, Jaeseong Im, Marco Canini, Dejan Kostić, Youngjin Kwon, Simon Peter, Emmett Witchel, LineFS: Efficient SmartNIC offload of a distributed file system with pipeline parallelism, in: Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles, 2021, pp. 756–771.
- [75] Seung-seob Lee, Yanpeng Yu, Yupeng Tang, Anurag Khandelwal, Lin Zhong, Abhishek Bhattacharjee, MIND: In-network memory management for disaggregated data centers, in: Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles, SOSP '21, New York, NY, USA, 2021, pp. 488–504.
- [76] Jialin Li, Jacob Nelson, Ellis Michael, Xin Jin, Dan R.K. Ports, Pegasus: Tolerating skewed workloads in distributed storage with In-Network coherence directories, in: 14th USENIX Symposium on Operating Systems Design and Implementation, OSDI 20, 2020, pp. 387–406.
- [77] Zaoxing Liu, Zhihao Bai, Zhenming Liu, Xiaozhou Li, Changhoon Kim, Vladimir Braverman, Xin Jin, Ion Stoica, DistCache: Provable load balancing for largescale storage systems with distributed caching, in: Proceedings of the 17th USENIX Conference on File and Storage Technologies, FAST '19, USA, 2019, pp. 143–157.
- [78] Jialin Li, Naveen Kr. Sharma, Dan R.K. Ports, Steven D. Gribble, Tales of the tail: Hardware, OS, and application-level sources of tail latency, in: Proceedings of the ACM Symposium on Cloud Computing, SoCC '14, New York, NY, US, 2014, pp. 1–14.
- [79] Suli Yang, Jing Liu, Andrea C Arpaci-Dusseau, Remzi H Arpaci-Dusseau, Principled schedulability analysis for distributed storage systems using thread architecture models, in: 13th USENIX Symposium on Operating Systems Design and Implementation, OSDI 18, 2018, pp. 161–176.
- [80] Henry Qin, Qian Li, Jacqueline Speiser, Peter Kraft, John Ousterhout, Arachne: Core-aware thread management, in: 13th USENIX Symposium on Operating Systems Design and Implementation, OSDI'18, Carlsbad, CA, 2018, pp. 145–160.
- [81] Amy Ousterhout, Joshua Fried, Jonathan Behrens, Adam Belay, Hari Balakrishnan, Shenango: Achieving high {CPU} efficiency for latency-sensitive datacenter workloads, in: 16th USENIX Symposium on Networked Systems Design and Implementation, NSDI 19, 2019, pp. 361–378.
- [82] Kostis Kaffes, Timothy Chong, Jack Tigar Humphries, Adam Belay, David Mazières, Christos Kozyrakis, Shinjuku: Preemptive scheduling for μsecond-scale tail latency, in: 16th USENIX Symposium on Networked Systems Design and Implementation, NSDI 19, 2019, pp. 345–360.

- [83] Călin Iorgulescu, Reza Azimi, Youngjin Kwon, Sameh Elnikety, Manoj Syamala, Vivek Narasayya, Herodotos Herodotou, Paulo Tomita, Alex Chen, Jack Zhang, et al., {PerfIso}: Performance isolation for commercial {Latency-Sensitive} services, in: 2018 USENIX Annual Technical Conference, USENIX ATC 18, 2018, pp. 519–532.
- [84] Mingzhe Hao, Gokul Soundararajan, Deepak Kenchammana-Hosekote, Andrew A Chien, Haryadi S Gunawi, The tail at store: A revelation from millions of hours of disk and {ssd} deployments, in: 14th USENIX Conference on File and Storage Technologies, FAST 16, 2016, pp. 263–276.
- [85] Shiqin Yan, Huaicheng Li, Mingzhe Hao, Michael Hao Tong, Swaminathan Sundararaman, Andrew A Chien, Haryadi S Gunawi, Tiny-tail flash: Near-perfect elimination of garbage collection tail latencies in NAND SSDs, ACM Trans. Storage (TOS) 13 (3) (2017) 1–26.
- [86] Shine Kim, Jonghyun Bae, Hakbeom Jang, Wenjing Jin, Jeonghun Gong, Seungyeon Lee, Tae Jun Ham, Jae W Lee, Practical erase suspension for modern low-latency {SSDs}, in: 2019 USENIX Annual Technical Conference, USENIX ATC 19, 2019, pp. 813–820.
- [87] Jisoo Yang, Dave B. Minturn, Frank Hady, When poll is better than interrupt, in: FAST, Vol. 12, 2012, pp. 3–3.
- [88] R. Recio, B. Metzler, P. Culley, J. Hilland, D. Garcia, A remote direct memory access protocol specification, 2007, https://datatracker.ietf.org/doc/html/ rfc5040.
- [89] Hyeong-Jun Kim, Young-Sik Lee, Jin-Soo Kim, NVMeDirect: A User-space I/O Framework for Application-specific Optimization on NVMe SSDs, in: 8th USENIX Workshop on Hot Topics in Storage and File Systems, HotStorage'16, 2016.
- [90] Ziye Yang, James R. Harris, Benjamin Walker, Daniel Verkamp, Changpeng Liu, Cunyin Chang, Gang Cao, Jonathan Stern, Vishal Verma, Luse E. Paul, SPDK: A development kit to build high performance storage applications, in: 2017 IEEE International Conference on Cloud Computing Technology and Science, CloudCom, 2017, pp. 154–161.
- [91] Jing Liu, Andrea C Arpaci-Dusseau, Remzi H Arpaci-Dusseau, Sudarsun Kannan, File systems as processes, in: 11th USENIX Workshop on Hot Topics in Storage and File Systems, HotStorage 19, 2019.
- [92] Jing Liu, Anthony Rebello, Yifan Dai, Chenhao Ye, Sudarsun Kannan, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, Scale and performance in a filesystem semi-microkernel, in: Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles, SOSP '21, New York, NY, USA, 2021, pp. 819–835.
- [93] Shengwen Liang, Ying Wang, Youyou Lu, Zhe Yang, Huawei Li, Xiaowei Li, Cognitive SSD: A deep learning engine for in-storage data retrieval, in: 2019 USENIX Annual Technical Conferenc, USENIX ATC 19, USENIX Association, Renton, WA, 2019, pp. 395–410.
- [94] Vikram Sharma Mailthody, Zaid Qureshi, Weixin Liang, Ziyan Feng, Simon Garcia de Gonzalo, Youjie Li, Hubertus Franke, Jinjun Xiong, Jian Huang, Wen-mei Hwu, DeepStore: In-storage acceleration for intelligent queries, in: Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '52, New York, NY, USA, 2019, pp. 224–238.
- [95] Jaeyoung Do, Victor C Ferreira, Hossein Bobarshad, Mahdi Torabzadehkashi, Siavash Rezaei, Ali Heydarigorji, Diego Souza, Brunno F Goldstein, Leandro Santiago, Min Soo Kim, et al., Cost-effective, energy-efficient, and scalable storage computing for large-scale AI applications, ACM Trans. Storage (TOS) 16 (4) (2020) 1–37.
- [96] Amedeo Sapio, Marco Canini, Chen-Yu Ho, Jacob Nelson, Panos Kalnis, Changhoon Kim, Arvind Krishnamurthy, Masoud Moshref, Dan R.K. Ports, Peter Richtárik, Scaling distributed machine learning with in-network aggregation, 2019, arXiv preprint arXiv:1903.06701.
- [97] ChonLam Lao, Yanfang Le, Kshiteej Mahajan, Yixi Chen, Wenfei Wu, Aditya Akella, Michael Swift, ATP: In-network aggregation for multi-tenant learning, in: 18th USENIX Symposium on Networked Systems Design and Implementation, NSDI 21, 2021, pp. 741–761.
- [98] Jian Zhang, Yujie Ren, Sudarsun Kannan, FusionFS: Fusing I/O operations using CISCOps in firmware file systems, in: 20th USENIX Conference on File and Storage Technologies, FAST'22, Santa Clara, CA, 2022, pp. 297–312.
- [99] Lei Wang, Kaiyong Yang, Wanling Gao, Chunjie Luo, Chengxi Wang, Zhongxin Ge, Li Zhang, Fan Zhang, Jianfeng Zhan, WPC: Whole-picture Workload Characterization, Technical Report, Institute of Computing Technology, Chinese Academy of Sciences.
- [100] Paul H.J. Kelly, Advanced computer architecture: The stored program concept and the turing tax, 2020, https://www.doc.ic.ac.uk/~phjk/ AdvancedCompArchitecture/Lectures/pdfs/Ch01-part4-TuringTaxDiscussion. pdf.
- [101] Paul H.J. Kelly, "Turing Tariff" reduction: Architectures, compilers and languages to break the universality barrier, 2020, https://www.doc.ic.ac.uk/ ~phjk/Presentations/2020-06-24-DoCLunch-PaulKelly-TuringTaxV04.pdf.
- [102] Wikipedia, The network is the computer, https://en.wikipedia.org/wiki/The_ Network_is_the_Computer.
- [103] SunMicrosystemsInc, The network is the computer, 2007, https://www.youtube. com/watch?v=kWNtaUaDxZw.
- [104] Chetan Sharma, Correcting the IoT history, 2016, 14 March 2016, http://www. chetansharma.com/correcting-the-iot-history/.

W. Gao, L. Wang, M. Chen et al.

- [105] Verband Der, VDI/Vde-Richtlinien, 2005.
- [106] Ken Wallace, Capturing, storing and retrieving design knowledge in a distributed environment, in: International Conference on Computer Supported Cooperative Work in Design, Vol. 1, IEEE Computer Society, 2005, pp. 10–Vol.
- [107] Ragunathan Rajkumar, Insup Lee, Lui Sha, John Stankovic, Cyber-physical systems: The next computing revolution, in: Design Automation Conference, IEEE, 2010, pp. 731–736.
- [108] Luiz André Barroso, Urs Hölzle, The datacenter as a computer: An introduction to the design of warehouse-scale machines, in: Synthesis Lectures on Computer Architecture, vol. 4, (no. 1) Morgan & Claypool Publishers, 2009, pp. 1–108.

- [109] Chinese Academy of Sciences Strategic Research Group in the Field of Information, China's Information Technology Development Roadmap to 2050, Science Press, 2009.
- [110] Peter C. Evans, Marco Annunziata, Industrial internet: Pushing the boundaries, Gen. Electr. Rep. (2012) 488–508.
- [111] Jianqiang Li, F Richard Yu, Genqiang Deng, Chengwen Luo, Zhong Ming, Qiao Yan, Industrial internet: A survey on the enabling technologies, applications, and challenges, IEEE Commun. Surv. Tutor. 19 (3) (2017) 1504–1526.
- [112] Guojie Li, Zhiwei Xu, Judging new economy from perspective of information technology trend, Bull. Chin. Acad. Sci. (Chinese Version) 32 (3) (2017) 233–238.

Contents lists available at ScienceDirect



BenchCouncil Transactions on Benchmarks, Standards and Evaluations

journal homepage: www.keaipublishing.com/en/journals/benchcouncil-transactions-onbenchmarks-standards-and-evaluations/



Research Article

Understanding hot interconnects with an extensive benchmark survey

Yuke Li^a, Hao Qi^a, Gang Lu^b, Feng Jin^b, Yanfei Guo^c, Xiaoyi Lu^{a,*}

^a University of California, Merced, 5200 North Lake Rd., Merced, 95343, CA, USA

^b Tencent, No. 33, Haitian Second Road, Shenzhen, 518054, Guangdong, China

^c Argonne National Laboratory, 9700 S. Cass Avenue, Lemont, 60439, IL, USA

ARTICLE INFO

Keywords: Benchmarks Interconnects RDMA

ABSTRACT

Understanding the designs and performance characterizations of hot interconnects on modern data center and high-performance computing (HPC) clusters is a fruitful research topic in recent years. The rapid and continuous growth of high-bandwidth and low-latency communication requirements for various types of data center and HPC applications (such as big data, deep learning, and microservices) has been pushing the envelope of advanced interconnect designs. We believe this is high time to investigate the performance characterizations of representative hot interconnects with different benchmarks. Hence, this paper presents an extensive survey of state-of-the-art hot interconnects on data center and HPC clusters and the associated representative benchmarks to help the community to better understand modern interconnects. In addition, we characterize these interconnects by the related benchmarks under different application scenarios. We provide our perspectives on benchmarking data center interconnects based on our survey, experiments, and results.

1. Introduction

The scales of data center and high-performance computing (HPC) clusters grow rapidly with the increasingly large volume of data and the high demand for distributed computing capabilities [1]. This trend has led to various designs of modern data center interconnects and made their performance characterizations a rewarding research topic. To continuously improve the performance and scalability of data movement or communication across a large number of nodes in modern data center or HPC clusters, different types of advanced interconnects have been designed to meet the requirements of high-bandwidth and low-latency communications in popular data center applications, such as deep learning, big data, microservices, etc.

To upgrade the conventional Ethernet (~10 Gbps) network and accelerate the efficiency of data center applications, hardware vendors have demonstrated multiple types of advanced data center interconnects. For example, NVIDIA (Mellanox) has produced 200 Gbps InfiniBand (IB) [2] with well-optimized Remote Direct Memory Access (RDMA) subsystems to speedup the inter-node communication in applications. Cray has the Slingshot interconnect [3] and the Aries interconnect [4] as high-speed interconnects for modern HPC systems. RIKEN (Japanese Institute of Physical and Chemical Research) and Fujitsu developed the Tofu interconnect [5] family to be equipped on their designed supercomputers. Meanwhile, the Ethernet network speed has improved from 10 Gbps to 100 Gbps [6] and even above [7] during the decades of development.

With the trend of hardware evolution and the new interconnects being created, there are several issues that the application developers need to pay attention to. With the hardware upgrading, the developers need to re-evaluate the performance of different generations of hardware to design the proper systems software based on the improved data transfer rates. Also, many new interconnects are emerging with the development of novel hardware features. These features may potentially impact application performance and need to be systematically investigated.

On the other hand, different types of data center applications represent various performance characterizations, like HPC workloads, deep learning training and inferences, big data analytics, and cloud-based microservice. The impacts of new interconnects on these different workloads should be evaluated separately and carefully. Therefore, we believe this is high time to investigate the performance characterizations of modern data centers and HPC interconnects via standard benchmarking experiments under different application scenarios. This observation motivates us to extensively survey hot interconnects on modern data centers and HPC clusters and the associated representative benchmarks to help the community better understand these advanced interconnects.

There exist some surveys to summarize benchmarking experiences with different workloads. For example, Han et al. [8] surveyed ten big data benchmarks to discuss benchmarking challenges. Zhang et al. [9]

* Corresponding author.

E-mail addresses: yli304@ucmerced.edu (Y. Li), hqi6@ucmerced.edu (H. Qi), gateslu@tencent.com (G. Lu), ronyjin@tencent.com (F. Jin), yguo@anl.gov (Y. Guo), xiaoyi.lu@ucmerced.edu (X. Lu).

https://doi.org/10.1016/j.tbench.2022.100074

Received 19 September 2022; Received in revised form 21 October 2022; Accepted 21 October 2022

Available online 28 October 2022

2772-4859/© 2022 The Authors. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).



Fig. 1. An overview of data center interconnects and benchmarks.

investigated fourteen deep learning benchmarks. Zhou et al. [10] discussed seven microservice benchmarks. Gao et al. [11] compared fifteen big data and AI (Artificial Intelligence) benchmarks. However, we did not find such a survey that can extensively cover a broad range of the latest advanced interconnects in modern data centers and the associated representative benchmarks for different application scenarios. Therefore, this paper addresses the need to survey different hot interconnects deployed in modern data centers and the corresponding benchmarks to expose their performance characteristics.

Fig. 1 shows an overview of this paper's surveying scope, including various kinds of hot interconnects and the associated representative and popular benchmarks in the community. The following sections will introduce each component in Fig. 1 with a bottom-up approach. In Section 2, we survey the features and characteristics of hot interconnects in modern data centers and HPC clusters. In Section 3, we survey well-used micro benchmarks for evaluating these hot interconnects with their network primitives and mechanisms. Section 4 will survey application-level benchmarks with diverse evaluation granularity, as shown in Fig. 1. In Section 5, we choose several representative benchmarks, which include Netperf [12]. Perftest [13], and OSU Micro-Benchmarks (OMB) [14] for MPI (Message Passing Interface) [15] and PGAS (Partitioned Global Address Space) [16] applications, and interconnects, which include IB, Omni-Path [17], and Ethernet to run experiments. We present the results to show performance characterizations of these hot interconnects as examples or reference numbers. Section 6 will discuss some of our observations and perspectives on benchmarking data center interconnects based on our survey, experiments, and results. Section 7 discusses more related studies and Section 8 concludes the paper.

The main contributions of this paper are as follows:

- We perform an extensive survey on advanced hot interconnects in current-generation and emerging data centers and HPC clusters.
- We also comprehensively survey the associated representative benchmarks from both micro benchmarking and application-level benchmarking perspectives.
- We perform a set of benchmarking experiments on real interconnects hardware with well-used benchmarks and discuss their performance characterizations.
- We share our observations on improvable aspects of existing benchmarks, such as performance stability, reference number, experimental instructions, etc., to help the community to design better ones.

2. Overview of modern interconnects

As an indispensable part of HPC and data center systems, interconnects play an essential role in achieving higher scalability and performance for modern clusters. In recent years, the community has witnessed the development of conventional interconnects like Ethernet and InfiniBand, and the birth of proprietary interconnects such as Fugaku Tofu [5] and BXI (Bull eXascale Interconnect) [18]. This section will briefly overview some representative state-of-the-art modern interconnects, and their features [1]. After we go through these interconnects one by one, Table 1 shows a brief comparison of these hot interconnects.

2.1. Ethernet

Ethernet is one of the most traditionally utilized interconnects for HPC and data center clusters. At the early stage, 1 Gb/s Ethernet (1-GigE) was widely used. However, with the advancement of CPU performance and I/O speed, the 1-GigE has become the bottleneck. With the demand for higher bandwidth and data transfer rate, Ethernet with 10-GigE, 25-GigE, 50-GigE, and even 100-GigE, has been developed. As of June 2022, 25-GigE is the most widely used interconnect in the Top500 list, and the Ethernet interconnect family is the majority in the list, taking up nearly 50% [19].

Taken the advantages of RDMA, RDMA over Converged Ethernet (RoCE) [20] is developed, which is a network protocol that allows RDMA to operate over Ethernet networks. RoCE is designed to support RDMA over Ethernet on layer 2 networks, and its extended version RoCE v2 enables transportation on layer 3 networks. Traditionally, Ethernet has left the congestion control to the TCP (Transmission Control Protocol) layer. With the development, the first algorithm proposed for the Ethernet network is pause frame [21] in 1996. Congestion control on RoCE uses an extension to the TCP/IP protocol called ECN (Explicit Congestion Notification) [22]. Other techniques, such as the QCN (Quantized Congestion Notification) [23], were developed afterward. Both traditional Ethernet and RoCE are available for various interconnect topologies. In 2019, Amazon announced EFA (Elastic Fabric Adapter) [24] for its EC2 (Elastic Compute Cloud) instance. The libfabric [25] interface on EFA provides up to 100 Gbps speed and reduces overhead with techniques like operating system bypass.

2.2. InfiniBand

Provided by NVIDIA, InfiniBand (IB) is an industry-standard switch fabric and the second most popular interconnect family in the Top500 list [19]. As of June 2022, 32.4% of the Top500 clusters are interconnected by IB, especially for Top10 clusters such as Summit [26] and Sierra [27]. Besides the higher bandwidth (up to 400 Gbps) and lower latency (<1 μ s), IB also supports advanced features like RDMA, which allows the software to read/write data from/to the memory in remote

Comparison of interconnects.

Name	Ethernet			InfiniBand	Omni-Path	Slingshot	Aries	TofuD	BXI
	25–100 Gbps	200–400 Gbps	RoCE						
Manufacturer	Many	Many	Many	NVIDIA/Mellanox	Intel/Cornelis	Cray	Cray	Fujitsu	Atos
Commodity	Public	Public	Public	Public	Public	Proprietary	Proprietary	Proprietary	Proprietary
Unidirectional Bandwidth (Gbps)	25-100 [37]	200-400 [37]	100 [38]	400 [39]	100 [37]	200 [37]	40 [4]	56 [<mark>40</mark>]	100 [37]
End to End Latency (µs)	10-30	N/A	~1 [37]	<1 [37]	<1 [37]	<2 [37]	~1 [4]	0.5-1 [5]	<1 [37]
Congestion Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Topology	Various	Various	Various	Fat-tree, Dragonfly+	Fat-tree	Dragonfly	Dragonfly	Torus	Various
RDMA	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	2014	2017	2010	1999	2015	2019	2012	2018	2015

nodes without any CPU involvement from the remote side. IB provides reliable or unreliable, and connected or datagram data transport types [28]. The reliable transport can guarantee the packet delivery in order but it spends extra time to wait for the acknowledgment from the receiving side. The unreliable transport cannot ensure the packet is received, but it does not need extra time for waiting the acknowledgment. The queue pairs for connected transports are connected in the one-to-one mapping, while the queue pairs for datagram transports are connected in the one-to-any mapping. The connected transport is more suitable for applications with a small number of connections. The datagram transport usually performs better in large-scale applications because fewer connection contexts need to be maintained in memory [29].

Specifically, RDMA-capable networks (like InfiniBand) typically support four types of transport modes: Reliable Connection (RC), Reliable Datagram (RD), Unreliable Connection (UC), and Unreliable Datagram (UD). SEND and RECV operations are supported by all modes, while the RDMA WRITE operation is unsupported by UD, and the RDMA READ operation is unsupported by UD and UC. The most commonly used network topology for IB is fat-tree [30], but it also supports other topologies like dragonfly+ [31]. The IB standard includes a congestion control mechanism to detect and resolve congestion by using two relay messages: FECN (Forward Explicit Congestion Notification) [32] and BECN (Backward Explicit Congestion Notification) [33]. When applying IB to GPU, CUDA 5.0 first introduced GDR (GPUDirect RDMA) [34]. GDR allows IB adapters to directly access the GPU memory while also bypassing the host. GDR can significantly increase data communication performance among GPUs, which further benefits the increasing number of redesigned classical HPC and machine/deep learning applications.

2.3. Omni-Path

Omni-Path was first released by Intel in 2015 as a part of Intel's Scalable System Framework with the purpose of increasing HPC workload scalability and aiming for low communication latency, low power consumption, and high throughput [17]. Omni-Path mainly includes the network card, switch, and network manager components. It is built on Intel technology with multiple features, such as traffic flow optimization and packet integrity protection. It is mainly designed to support fat-tree topology, and its CCA (Congestion Control Architecture) has been updated continuously since its first release. The first generation Omni-Path delivers 100 Gbps bandwidth per port and is integrated into some CPU architectures like Skylake and Knights Landing (KNL) [35]. Although Intel stopped the development of the second-generation Omni-Path in 2019, it still takes 7.8% of Top500 clusters as of June 2022 [19]. In late 2020, Intel announced its spin-off to Cornelis Networks [36] to continue the business as a successor to the Omni-Path product.

2.4. Slingshot

In 2019, Cray launched its new generation of HPC interconnect technology called Slingshot [3]. Slingshot uses protocols on standard

Ethernet while also being compatible with proprietary HPC networks when needed. It offers key features like adaptive routing, quality of service guarantee, and advanced congestion control fully implemented in hardware. The slingshot switch is equipped with 64 ports, and each port is running at 200 Gbps. Slingshot also supports multiple interconnect topologies such as fat-tree and dragonfly [41]. As Cray's eighth major high-performance interconnection network technology, Slingshot is deployed on a variety of clusters like pre-exascale cluster Perlmutter [42] and exascale cluster Frontier [43], which is currently the top 1 supercomputer in the world. Slingshot is also planned to be deployed on upcoming exascale clusters like Aurora [44] and EI Capitan [45]. Slingshot is taking up 4.8% of the clusters in Top500 list as of June 2022 [19].

2.5. Aries interconnect

As Cray's third-generation interconnect architecture, Aries was introduced as part of the Cray XC system with the dragonfly topology, and it has been widely used in the HPC field [4]. A single Aries device with four NICs (Network Interface Card) and a 48-port tiled router can provide a network connection for all four nodes on a Cray XC blade. The NIC and switch in Aries are closely coupled in the dragonfly network to provide cost-effective and scalable global bandwidth. The system is configurable according to users' global bandwidth requirement, and its optical connection number can be adjusted according to the cost constraint. It also provides technologies such as adaptive routing, communication mechanisms, and synchronization mechanisms. Aries adopts the dragonfly topology and achieves congestion control by implementing Valiant's routing algorithm [46]. As of June 2022, 5% of the Top500 clusters use Aries, including Piz Daint [47] and Cori [48].

2.6. Tofu interconnect D

As one of the representatives of proprietary interconnects, Tofu [49] is an interconnect family developed by RIKEN and Fujitsu that is used for the K computer [50]. In 2018, TofuD (Tofu Interconnect D) was introduced as a new member of the Tofu family. Its main features are just as indicated by the name. The Tofu represents "torus fusion" and the letter D stands for high "density" node and "dynamic" packet slicing for "dual-rail" transfer [5]. TofuD is a proprietary torus-based [5] six-dimensional network, and it mainly supports congestion control with the family's virtual channel scheduling algorithm. Compared to previous Tofu and Tofu2 [51], TofuD has a much higher communication resource density, such as 48 cores per node. It also introduced dynamic packet slicing for the dual-rail transfer technique to solve the latency and fault tolerance issue in Tofu2. TofuD is adopted by the Fugaku [52], which was the top 1 cluster in the Top500 list at the time built in 2020 and ranked 2nd as of June 2022 [19].

2.7. Bull eXascale interconnect (BXI)

In 2015, Atos designed BXI as a new interconnect for HPC [18]. BXI is based on the scalable and reliable Portals4 [53] network programming interface and decouples computation and communication

The summary of micro-benchmarks.

	Perftest [13]	RDMA-bench [56]	NetPerf [12]	iPerf [57–60]	qperf [61]	GPCNeT [62]
Link layer Programming Models	IB, Eth (RoCE) RDMA	IB, Eth (RoCE) RDMA	Eth Socket	Eth Socket	IB, Eth (RoCE) RDMA/Socket	IB, Eth (RoCE), etc. MPI
Transport Protocols	RC/UC/UD DCT, SRD	RC/UC/UD	TCP, UDP, SCTP	TCP, UDP, SCTP	TCP, UDP, SCTP, SDP RC/UC/UD, RDS	Any protocol that can be used by MPI
Main metrics	Throughput, Average latency, Tail latency	Throughput, Average latency, Tail latency, WQE cache misses	Throughput, Average latency	Throughput, Average latency, Tail latency	Throughput, Average latency	Throughput, Average latency, Tail latency, Congestion Impact
Language	С	С	С	С	С	С
Thread Model	Single-thread	Single-thread	Single-thread	Multi-thread	Single-thread	Multi-process (MPI)
Communication pattern	P2P	P2P	P2P	P2P; multicast	P2P	P2P, collective communication
Real scenario	Ν	Y (w/ real applications)	Ν	Ν	Ν	Y
Real workload or trace	Ν	N	Ν	Ν	Ν	Y
Parameters of protocol internals	Ν	Ν	Ν	Y	Ν	Ν
Year of last update	2022	2018	2021	2022	2018	2021

by hardware offloading. It consists of two ASIC (Application-Specific Integrated Circuit)-based components: BXI NIC and BXI switch. The BXI NIC provides functions like OS bypass, communication offload, and reliability. Each BXI switch is equipped with forty-eight 100 Gbps ports and provides power saving and network performance monitoring functions. BXI supports multiple network topologies such as fat-tree, butterfly [54], and torus. BXI implements efficient fine-grain adaptive routing on each port basis to minimize the possibility of congestion. It also provides reliability and stability guarantees by some optimizations like deadlock-avoidance and load-balancing mechanisms. As of June 2022, Tera-1000-2 adopts BXI 1.2 and it is ranked 45th in the Top500 list [19].

2.8. Summary

We survey the above hot interconnects because of their popularity for clusters in the Top500 list. Table 1 summarizes a brief comparison of these hot interconnects. They show a huge diversity in aspects, such as bandwidth, latency, congestion control mechanism, and network topology, which motivates us to investigate their performance characteristics. Due to the lack of access to proprietary interconnects, we mainly focus on evaluating Ethernet, RoCE, IB, and Omni-Path in this paper. InfiniBand is an essential interconnect for native RDMA designs. 10/25 Gbps Ethernet networks are the majority (27.2%) of interconnects used in data center and HPC clusters. RoCE and TCP/IP can be deployed on 10/25 Gbps Ethernet [55]. We evaluate these different interconnects and show the results in Section 5.

3. Survey of micro-benchmarks

The community has designed many benchmarks to evaluate various types of interconnects. In this section, we survey six micro-benchmarks designed to measure low-level performance metrics such as latency and bandwidth. We introduce their features and discuss their pros and cons. As shown in Table 2, six publicly-available micro-benchmarks are included for comparison. The rest of this section discusses these micro-benchmarks one by one.

3.1. Perftest

Perftest [13] was developed by Mellanox and has been well maintained since 2005. The RDMA community widely uses it for latency and bandwidth performance evaluation on InfiniBand and RoCE networks. The included micro benchmarks adopt a single-thread and ping-pong communication pattern to evaluate the throughput and latency of basic RDMA operations. We can also use it to compare different transports by specifying the transport as RC, UC, UD, Raw Ethernet, and even Mellanox DCT (Dynamic Connected Transport) [63] and AWS SRD (Scalable Reliable Datagram) [64] transports that are not specified in the standard IB specification [65]. Besides the basic operations and transports, Perftest also supports the GPUDirect feature for direct inter-GPU communication through GPUDirect RDMA and the AESXTS [66] feature for data encryption and decryption scenarios using RDMA. Perftest is designed without emulating any real application traffic or traffic probability distribution. It does not allow users to choose the traffic pattern but only with a parameter to specify the message size in each test. These tests are mainly helpful for hardware or software tuning as well as for functional testing.

3.2. RDMA-bench

RDMA-bench [56] was developed by Carnegie Mellon University in 2016. Unlike Perftest, RDMA-bench is a new benchmark suite used to understand the RDMA performance in a few scenarios extracted from real applications. With the guidelines obtained from running RDMA-bench, the authors of RDMA-bench succeeded in developing a networked sequencer and a key–value store far superior to others [67].

The benchmarks in RDMA-bench can be classified into several categories: (1) application benchmarks which include HERD [68] and MICA [69] as RDMA-based key–value store systems, and DrTM-KV [70] as an RDMA-based in-memory transaction processing system; (2) micro-benchmarks which measure the throughput of outbound and inbound RDMA operations; (3) micro-benchmarks that emulate an echo server, in which users can choose different RDMA operations for the requests and responses; (4) micro-benchmarks which emulate an RPC (Remote Procedure Calls) based sequencer server using different RDMA transports and operations; (5) micro-benchmarks which emulate a complex communication scheme with configurable thread-QP ratios to the scalability evaluation; (6) micro-benchmarks which help understand low-level factors that affect RDMA performance, such as WQE cache misses of outbound READs and WRITEs, etc.

3.3. Netperf

Netperf [12] was developed by Hewlett-Packard in 2005. It is widely used to measure the performance of BSD Sockets [71] for TCP, UDP, or SCTP (Stream Control Transmission Protocol) [72] using IPv4 and IPv6, Unix domain sockets [73], and DLPI (Data Link Provider Interface) [74]. Netperf adopts a simple client–server model without multi-threading support. The main parameters include the socket buffer size, the message size, the TCP_NODELAY option, and the test mode. There are two test modes supported in Netperf: (1) the STREAM mode, which transfers bulk data through a TCP or UDP socket; (2) the RR (Request/Response) mode, which emulates iterative requester–response transactions between the client and server. The data transmitted is synthetic. Neither different probability distributions nor real-world data trace is supported. Hewlett-Packard made a plan of version 4.x of Netperf, which aimed to support synchronized and multi-threaded benchmarking.

3.4. iPerf

iPerf [57,58] is used to evaluate the performance of TCP, UDP, and SCTP traffic with IPv4 and IPv6. It provides abundant features [57,58]: (1) iPerf adopts a multi-threaded design that can scale with the number of CPUs within a system; (2) iPerf supports tuning of various parameters that are rarely supported in Netperf, such as timing, buffers, and most importantly, the internal parameters of the protocols; (3) iPerf supports multicast tests and bidirectional tests; (4) iPerf can run on many platforms which include Linux and Windows; (5) users can get various forms of outputs in iPerf; (6) iPerf provides the libiperf library, which is an straightforward way to use and customized the functionality of iPerf.

iPerf has evolved into two incompatible active branches. One branch is iPerf2 [57] which is the newer version of the original iPerf. The other branch is iPerf3 [58] which is a redesign of the original iPerf and was now principally developed by ESnet and Lawrence Berkeley National Laboratory. Either of them contains several options and functions that are not present in the other. Generally, for TCP and UDP in Ethernet, iPerf2 and iPerf3 are about the same if running with the default configuration. However, users should check the detailed comparison in [59,60] to avoid misuse.

3.5. qperf

qperf [61] was initially developed by QLogic in 2007 and then maintained by the Linux community. qperf can measure the bandwidth and latency between two hosts using TCP, UDP, SCTP, RDMA, SDP (Sockets Direct Protocol), and RDS (Reliable Datagram Sockets). It adopts a single-threaded client–server model similar to Netperf. For RDMA, we can test the bandwidth and latency of RC, UC, and UD transports. All the operations can be measured for each transport in the tests. Compared to Perftest, qperf supports fewer transports and features from the perspective of evaluating RDMA performance. For non-RDMA protocols, the option of qperf can only change the message size. Evaluations of the internal features of the protocols cannot be done by using qperf. Even though qperf only reports average latency and fails to perform precise tail latency measurements, it is still popular as it is a handy and tool. The release of qperf is stable and the light-weight update was four years ago.

3.6. GPCNeT

The Global Performance and Congestion Network Test (GPCNeT) [62,75] was developed by Cray in 2019 to evaluate the network performance of MPI-based systems with the MPI-3.0 specification [76]. GPCNeT is compromised of two benchmarks: *network_test* and *network_load_test*.

network_test characterize the latency and bandwidth of an MPI application when it runs without network congestion. It builds the natural ring and random ring pattern such that all communication occurs over the network rather than within local groups. The communication patterns include two-sided peer-to-peer (8 bytes latency and 128K bytes bandwidth, natural and random rings), one-sided remote

memory access (8 bytes latency and 128K bytes bandwidth, random ring), allreduce (8 bytes latency, random ring), and alltoall (128 bytes bandwidth, random ring).

network_load_test measure the performance of an MPI application with network congestion. This simulates the scenario when running on multi-tenant HPC networks. Each congestor has a unique random ring, and the communication patterns include Point-to-point Incast, All-to-all, One-sided RMA Incast, and One-sided RMA Broadcast. Two measurements execute in the random ring infrastructure: Point-to-point Latency measurement by sending and receiving 8 bytes messages from and to two sides, Point-to-point Bandwidth with Synchronization by sending and receiving eight 128K bytes messages from two sides.

The default settings are intended to be utilized in general production scenarios. It reports the mean and 99th percentile latencies as well as the bandwidth per rank. With congestors, it also reports the Congestion Impact metric, which is defined as the ratio of congested latency or bandwidth divided by the uncongested latency or bandwidth. The Congestion Impact metric is an indicator to study the impact of congestion across systems with different networks.

3.7. Summary

The above micro-benchmarks are surveyed because of their popularity in the community. In Table 2, we show a summary of these micro-benchmarks. Among the six micro-benchmarks, we will test the interconnects with Perftest and NetPerf in this survey. They are both widely used and well maintained since their first release. Besides, Perftest is provided by Mellanox, the most popular manufacturer of InfiniBand. Hence, we believe Perftest and NetPerf can represent defacto standard benchmarks for RDMA-based and socket-based programming models on various interconnects, respectively. We show the related results in Section 5.

4. Survey of application-level benchmarks

There are diverse types of workloads running across machines in a data center, from parallel computing to microservice, from GPU applications for deep learning workloads to Key–Value Store for big data workloads. The same issue these workloads share is that they all need efficient data communication through the interconnects. As mentioned above, different interconnects may show different characterizations on the same application. Therefore, researchers need to use benchmarks to characterize the application that runs on a specific interconnect. This section surveyed application-level benchmarks with diverse evaluation granularity for different application scenarios in data centers that involve cross-node communication via interconnects. To save space, we put detailed descriptions of these benchmarks in tables.

4.1. MPI benchmarks

MPI [15] is a message-passing standard and widely used in HPC where many processes or cores are organized to run parallel program simultaneously for acceleration. Using a benchmark to characterize MPI libraries on different interconnects can help developers understand the characteristics of interconnects and design applications in efficient ways.

We surveyed three popular MPI benchmarks. The OSU MPI Micro-Benchmarks [14] provided by Ohio State University (OSU) consist of point-to-point MPI operations, blocking/non-blocking collective MPI operations, and one-sided MPI operations. Table 3 shows the description details. The NAS Parallel Benchmarks (NPB) [77], provided by NASA, are derived from CFD (computational fluid dynamics) applications and are designed with MPI programming. Its description details are shown in Table 4. The Intel MPI Benchmarks (IMB) [78] provided by Intel perform MPI 1.0 \sim 3.0 measurements for communication operations for a range of message sizes, which are shown in Table 5.

The details of OSU MPI Micro-Benchmarks

		•		
Categories	Metrics Benchmarks		Description	
Point-to-	Latency	osu_latency, osu_latency_mt, osu_latency_mp, osu_multi_lat	The test in ping-pong fashion with single-/multi- thread, multi- process, and multi-pair operations.	
Point	Bandwidth	osu_bw, osu_bibw, osu_mbw_mr	The test in sending/receiving back-to-back messages with uni- /bi-directional operations.	
Collective Latency		osu_allgather, osu_allgatherv, osu_allreduce, osu_allroall, osu_allroallv, osu_barrier, osu_bcast, osu_gather, osu_scather, osu_reduce, osu_scatter, osu_scatterv	The test of blocking MPI collective operations.	
Non - Blocking Collective	Latency	osu_iallgather, osu_iallgatherv, osu_iallfeduce, osu_ialltoall, osu_ialltoall, osu_ialltoallw, osu_ialtoallw, osu_iagather, osu_igatherv, osu_ireduce, osu_iscatter, osu_iscatterv	The test of non-blocking collectives.	
One-sided	Latency	osu_put_latency, osu_get_latency, osu_acc_latency osu_cas_latency osu_fop_latency, osu_get_acc_latency	The test for one-sided MPI Put, Get, Accumulate, Compare and Swap, Fetch and Op, and Get_accumulate.	
	Bandwidth	osu_put_bw, osu_get_bw, osu_put_bibw,	The test for one-sided MPI Put and Get, in uni-/bi-direction.	

Table 4

The details of NAS Paralle	l Benchmarks (NPB).		
Benchmarks	Description		
Five kernels: IS, EP, CG, MG, FT	The original eight benchmarks to		
Three pseudo applications: BT, SP, LU	movement in CFD applications.		
BT-MZ, SP-MZ, LU-MZ	Benchmarks with multiple levels and hybrid parallelism.		
UA, BT-IO, DC, DT	Benchmarks for unstructured computation, parallel I/O, and data movement		

4.2. PGAS benchmarks

PGAS (Partitioned Global Address Space) is a parallel programming model in the HPC community. PGAS is defined by communications on a shared memory space that every Processing Element (PE) can access without permission issues. Many programming languages and libraries are designed from the PGAS model, e.g., Unified Parallel C (UPC) [79] and OpenSHMEM [80]. Communication happens when the processes transfer data from the global memory or to the global memory space, including within and across a node. OSU Micro-Benchmarks also provides benchmarks on PGAS model: OpenSHMEM benchmark is shown in Table 6; UPC and UPC++ benchmarks with point-to-point (*put* and *get*) and collective communications.

4.3. RPC benchmarks

RPC is a method when a process on a machine calls procedures on other machines where the execution of the procedure happens [81]. It is a client–server interaction where data is transferred frequently over the interconnects to call (from the client) and respond (from the server) procedures. Therefore, the characteristics of the interconnect can have a direct impact on the RPC performance.

We surveyed three benchmarks for RPC applications in data centers: Apache Thrift Benchmarks (ATB), TF-gRPC-Bench, and RPC-perf. ATB is proposed in [82], which evaluates the Apache Thrift [83] based

 Table 5

 The details of Intel MPI Benchmarks.

MPI-1						
Categories	Metrics	Вс	enchmarks	Description		
Single Transfer	Latency, Throughput	Pii Pii Pii	ngPong, ngPongSpecificSource, ngPongAnySource	Using non- blocking/blocking MPI send/recv to transfer data in ping-pong pattern.		
Parallel Transfer	Latency, Message Rate, Throughput	Se Ui Pii M Ex M	ndrecv, Exchange, niband, Biband, Multi- ngPong, Multi-PingPing, ulti-Sendrecv, Multi- tchange, Multi-Uniband, ulti-Biband	Using blocking or non- blocking MPI send/recv to transfer data among multiple processes in uni- /bi-direction.		
Collective	Latency, Throughput	Beast/multi-Beast, Allgather/multi-Allgather, Allgather/multi-Allgathery, Alltoall/multi-Alltoall, Scatter/multi-Seatter, Scatter/multi-Seatter, Gather/multi-Gathery, Reduce_scatter/multi-Gathery, Reduce_scatter/multi-Gathery, Reduce_scatter, Allreduce/multi-Allreduce, Barrier/multi-Allreduce, Barrier/multi-Barrier		Using blocking MPI collective operations to transfer data among multiple processes.		
IMB-P2P	IMB-P2P Latency, Bandwidth, Message Rate		ngPong, PingPing, nirandom, Birandom, orandom	Using blocking or non- blocking MPI send/recv to transfer data in the shared memory.		
			MPI-2			
Categories	Metrics		Benchmarks	Description		
IMB-EXT	Latency, Throughp	ut	Unidir_Put, Unidir_Get, Bidir_Put, Bidir_Get, Accumulate, Window	Benchmarks in one- sided communications with uni-/bi-direction.		
IMB-IO Timing, Throughpu		ut	S_[ACTION}_indv/expl P_[ACTION}_indv/expl /shared/priv, C_[ACTION}_indv/expl /shared, Open_Close	' Input/Output (I/O) benchmarks with blocking or non- blocking I/O.		
			MPI-3			
Categories	Metrics	Be	nchmarks	Description		
IMB-NBC	Timing, Overlap Percentage		cast, Iallgather, Iallgatherv, hther, Igatherv, Iscatter, atterv, Ialltoall, Ialltoallv, duce, Ireduce_scatter, Ireduce, Ibarrier. Add pure'' to measure pure time	Using non-blocking collective operations to measure the overlap of communication and computation or the pure communication time.		
IMB-RMA Throughput		Ac All Bio Ex Fer Ge Ge On Pu Un	cumulate, All_get_all, _put_all, Bidir_get, dir_put, Compare_and_swa change_Get, Exchange_Put tch_and_op, t_accumulate, Get_all_loc: t_local, One_put_all, le_get_all, Put_all_local, t_local, Truly_passive_put idir set, Unidir put	^{up} Using one-sided communication in remoti al, memory access (RMA) benchmarks in a uni-/bi- directional manner.		

RPC performance and consists of three categories: the RPC latency evaluation benchmark, the RPC throughput evaluation benchmark, and the mixed RPC latency and throughput evaluation benchmark. Table 7 shows the details of TF-gRPC-Bench, which evaluates the communication performance between parameter server and worker process. Twitter maintains RPC-perf [84]. It is designed to evaluate the RPC's performance for caching systems regarding latency and message rate.

4.4. Storage benchmarks

With the development of hardware technology, much new storage hardware is produced, such as the NVMe SSD [85]. Storage systems rely on different drivers and libraries in data center, with interactions between the processors and the storage devices via interconnects. Intel SPDK [86] provides NVMe perf [87] as an NVMe SSDs benchmarking tool with minimal overhead in benchmarking. NVMe perf provides several runtime options to support the most common workload. Users

The details	of	OSU	0	penSHMEM	Micro-Benc	hmarks.

Categories	Metrics	Benchmarks	Description
	Latency	osu_oshm_put, osu_oshm_put_nb, osu_oshm_get, osu_oshm_get_nb	Benchmarks for PUT and GET with blocking and non-blocking routines The users can select whether the memory allocation mode should be in global memory (non-symmetry) or heap memory (symmetry). All benchmarks use 2 PEs and the uni-directional operation. The osu_oshm_put_overlap provides the statistics for communication and computation overlap measurement.
Point-to- Point	Message Rate	osu_oshm_put_mr, osu_oshm_put_mr_nb, osu_oshm_get_mr_nb, osu_oshm_put_overlap	
Collective	Latency	osu_oshm_collect, osu_oshm_fcollect, osu_oshm_broadcast, osu_oshm_reduce, osu_oshm_barrier	Benchmarks for collective operations. The users can specify the number of processes to run, and the results give the minimum, maximum, and average latency.
Atomics	Latency, Operation Rate	osu_oshm_atomics	Benchmarks for Atomics Routines. The users can select the memory allocation mode between global and heap modes.

Table 7

The details of TF-gRPC-Bench.

Categories	Description				
Point-to-Point latency	Benchmarks for measuring the latency and	The benchmarks can run on single or multiple			
Point-to-Point bandwidth	the bandwidth of payload transmission between a parameter server and a worker process.	nodes. Users can specify IP, port, warm-up period, total running period, buffer, payload generation			
Parameter Server Throughput	Benchmark for measuring the throughput of parameter server over gRPC.	scheme, payload serialization mode, and the number of parameter servers and workers.			

Table 8

The details of IOzone and Iometer.

Benchmark	Description	Metrics
IOzone	Evaluating the file system.	(Random/backwards/normal) Read/Write, Mmap, Async I/O
Iometer Evaluating the I/O subsystem.		I/O speed, response time (avg, max), CPU utilization, error count

can configure the NVMe perf in many aspects like the workload characterizations (e.g., the percentage of Read/Write, with/without random Read/Write), the data movement protocols (e.g., PCIe, RDMA, TCP), and the execution time [87].

Besides the hardware, modern data centers also have different storage systems. Two surveyed benchmarks for storage systems are shown in Table 8. IOzone [88] is a filesystem benchmark to measure the file operations in storage systems. Iometer [89] is an I/O subsystem measurement tool for single and clustered systems. And one benchmark for EC (Erasure Coding) coder on distributed storage systems. EC-Bench [90] is an erasure coding scheme benchmark for storage architectures with description details in Table 9.

4.5. GPU applications benchmarks

GPU has been becoming incredibly popular for compute-intensive workloads in data centers and HPC clusters in recent years. Two popular deep learning frameworks, TensorFlow [91] and PyTorch [92], provide benchmarks to evaluate deep learning models, such as PerfZero [93] and TorchBench [94]. PerfZero is a benchmark framework

Table 9

The	details	ot	EC-Bench.
		_	

Categories	Description	Metrics
Encoding	A large in-memory file is split into data blocks. Each encoding operation encodes one data clock into new data and parity chunks.	Latency, Throughput, CPU Utilization,
Decoding	Recover the encoded data and the parity chunks.	Cache Pressure

Table 10

The details of PerfZero and TorchBench.

Benchmark	Platform	Metrics	Description
PerfZero	TensorFlow	wall time, CPU time, throughput, accuracy, extra information about system and environment, etc.	Evaluating and profiling for TensorFlow. Users can specify the behavior of the test, like number of iterations to run, the values to feed for each iteration, and the memory usage.
TorchBench	PyTorch	wall time, GPU time, extra information about system and environment, etc.	Evaluating and profiling for PyTorch. Users can specify the behavior of the test, like devices, modes, runs with evaluation/training.

for debugging and tracking the TensorFlow performance regression and change. TorchBench includes a collection of open-source benchmarks to evaluate models and workloads with PyTorch. More details about PerfZero and TorchBench are shown in Table 10. The PARAM benchmark [95] from Meta Platforms (Facebook formerly) can both evaluate the performance of communication components in the Py-Torch deep learning framework, and evaluate the application-level workloads, like deep learning recommendation models [96,97]. Deep-Bench [98] produced by BaiduResearch is another benchmark for evaluating deep learning operations on different platforms. NCCL (NVIDIA Collective Communications Library) [99] and Gloo [100] provide their benchmarks on collective communication libraries, which are NCCL Tests [101] and Gloo Benchmarking [102] to evaluate the performance on collective operations.

OSU Micro-Benchmarks also provide several extensions for GPU programming models and libraries, such as CUDA [103], ROCm [104], and OpenACC [105] extensions by configuring with --enable-cuda, --enable-rocm, and --enable-openacc in the runtime [14].

4.6. Key-Value Store benchmarks

Key–Value Store holds a data storage model that stores associations between keys and values. Keys are primitives, and values can be primitive or complex. It is popular in the big data community and widely used in NoSQL databases in data centers because of its high efficiency and scalability. We surveyed two benchmarks for Key– Value Store. YCSB [106] (Yahoo! Cloud Serving Benchmark) is used for evaluating the performance of key–value and cloud serving stores. YCSB provides five workloads with different percentages of database operations and evaluates three metrics of performance: the latency of requests, the database performance when increasing machines, and the database performance with increasing machines while the system is running. OSU HiBD-Benchmarks [107] provide benchmarks for evaluating Memcached and HBase based Key–Value Store.

4.7. Microservice benchmarks

Microservice is a type of cloud service architecture. Unlike traditional monolithic applications, microservice consists of multiple services working together to finish a workload. Therefore, communications happen frequently among services via interconnects in data centers. We surveyed two benchmarks for microservice workloads.

The details of DeathStarBench.

Microservice / Benchmark	Comm. Protocol	Unique Micro- services	Description	Data Storage
Social Network	RPC	36	A broadcast-style social network with posting posts, accounts and posts management, and friend recommendation functionalities.	Memcached, MongoDB
Movie Reviewing	RPC	38	A movie information management system with searching, browsing, reviewing, and rating the movies.	Memcached, MongoDB, MySQL database
E-commerce Website	REST, RPC	41	An e-commerce site for clothing with searching, browsing, ordering, and paying functionalities.	Memcached, MongoDB
Banking System	RPC	34	A secure banking system with processing payments, requesting loans, and balances credit card functionalities.	Memcached, MongoDB, Relational database
Swarm Cloud	REST, RPC	25	It runs on the cloud and edge devices with a swarm of	Cassandra
Swarm Edge	REST	21	programmable drones to perform image recognition and obstacle avoidance.	database

Table 12

The details of	pource.
Microservice / Benchmark	Description
HDSearch	HDSearch is a content-based image similarity search application service that searches the images whose vectors are near the query image's feature vector in a large image repository. The front-end sends the query to the backend, then the backend returns the results. It evaluates the performance of image search functionality.
Router	Router is a distributed Memcached service. It routes client requests to suitable Memcached servers. It evaluates the performance of GETs and SETs
Set Algebra	Set Algebra is a set intersection processing service. It performs information retrieval on posting lists by searching through millions of documents. It evaluates the performance of that searching process.
Recommend	Recommend is a recommendation system in the fields of e- commerce and behavior prediction. It predicts the users' rating for an item by using their overall preference data. It evaluates the performance of each {user, item} query.

DeathStarBench [108] is an open-source benchmark suite for microservices on cloud and edge systems, and the details are shown in Table 11. μ Suite [109] can be used for evaluating the influence of OS and network on microservices, and the details are shown in Table 12.

4.8. Summary

The above-mentioned application-level benchmarks can represent a broad range of data center applications, including HPC, big data, AI, and cloud computing. In this survey, we choose MPI and PGAS based benchmarks as application examples and run them on different interconnects. MPI and PGAS based benchmarks have been designed and maintained for many years with a lot of contributed optimizations from the community. Our experience also reveals that they are easy to deploy and convenient to run. The experiment results are shown in Section 5.

5. Experiment

This section presents performance characterizations with the selected benchmarks on various hot interconnects.

Table 13

The details of the testbeds in the experiments. PADSYS and Pinnacles [110] clusters are used in Section 5.2 and Section 5.3, Bebop [111] and JLSE [112] clusters are used in Section 5.4.

Testbed	Interconnect	Intel Xeon	RAM	Communication
(Nodes)	(Gbps)	CPU	(GB)	Subsystem
PADSYS	InfiniBand	Calil 6220	256	OFED
(2)	(200)	G010 0330	200	MLNX-5.5
Pinnacles	InfiniBand	Cold 6220	256	RHTL 8.6
(40)	(100)	G010 0550	230	IB Driver
Bebop	Omni-Path	Broadwell	100	Libfabric [25]
(36)	(100)	E5-2695v4	120	v1.15.1
JLSE	InfiniBand	Platinum	768	UCX [113,114]
(13)	(100)	8180M/8176	/00	v1.13.0



Fig. 2. The latency of Perftest on 200 Gbps InfiniBand, Perftest on RoCE (25 Gbps Ethernet), Netperf on 10/25 Gbps Ethernet, and Netperf on IPoIB (200 Gbps InfiniBand).

5.1. Benchmarking setup

We run benchmarks on different clusters with various interconnects and Table 13 shows the details of each cluster. We try to keep the comparisons of the experiment results as fair as possible by: (1) allocating nodes in the same rack across different experiments; (2) tuning the number of iterations (Perftest) or time duration (Netperf) of benchmark options until getting the relatively stable results.

5.2. Micro-benchmark evaluation

We organize the following experiments using two programming models, RDMA and socket, with two Perftest and Netperf microbenchmarks on three 10/25 Gbps Ethernet and 200 Gbps InfiniBand interconnects. We discuss the experiment results in three aspects: (1) the latency comparison; (2) the bandwidth usage comparison; (3) the impact on the performance using two different InfiniBand interconnects.

5.2.1. Latency

Fig. 2 shows the latency of benchmarks based on different network. The Perftest benchmark on 200 Gbps InfiniBand, the fastest interconnect in our experiments, has the lowest latency because of the nature of kernel-bypass and high-performance protocol in RDMA. Although RoCE also supports RDMA, the hardware it uses is 25 Gbps Ethernet on our testbed which is lower than 200 Gbps InfiniBand, so the Perftest on RoCE are slower than those on InfiniBand. The Netperf [12] is a TCP benchmark running on IPoIB and 10/25 Gbps Ethernet. Due to the well-known heavy overhead of TCP [29,115], the latency numbers of these three are far slower than the native RDMA designs, and the latency becomes larger with the decrease of the network bandwidth.



Fig. 3. The bandwidth of Perftest on 200 Gbps InfiniBand, Perftest on RoCE (25 Gbps Ethernet), Netperf on 10/25 Gbps Ethernet, and Netperf on IPoIB (200 Gbps InfiniBand).

5.2.2. Bandwidth

Fig. 3 shows the bandwidth comparisons using different benchmarks on different interconnects. The RDMA benchmark, Perftest, runs on 200 Gbps InfiniBand and can achieve the highest bandwidth of Infini-Band when the message size is large enough (4K bytes) because of the 4K bytes MTU setting on InfiniBand. The same benchmark running on 25 Gbps Ethernet, shown as RoCE in the figure, shows the same behavior but its MTU is 1K bytes so RDMA RoCE saturates the bandwidth earlier than the one on RDMA IB. The Netperf benchmark running on 10/25 Gbps Ethernet shows a similar behavior when the message size is larger than one MTU (1K bytes) but does not show the same on 200 Gbps InfiniBand. The reason is the TCP protocol stack overhead by deploying IPoIB on InfiniBand. We also observed the unstable results on Netperf benchmark evaluations and the reason could come from the performance fluctuation nature of TCP. Therefore, we ran the experiment five times for each one and took the average results to show in the figure.

5.2.3. InfiniBand EDR VS. HDR

200 Gbps InfiniBand (HDR) is emerging as a replacement of the widely-used 100 Gbps InfiniBand (EDR) in data center and HPC clusters. Therefore, it is high time to use benchmarks to compare the performance characteristics between EDR and HDR. In this experiment, we run the same benchmark, Perftest, on these two kinds of InfiniBand interconnects. Fig. 4 shows the bandwidth comparison and Fig. 5 shows the throughput comparison. As we expect, the saturated bandwidth (when message size is larger than 4K bytes, which equals to one MTU) of IB HDR is around two times that of IB EDR. We observe that the throughput of IB EDR is lower than the throughput of IB HDR all the time, and the IB HDR throughput numbers are 1.5X-2X times of the EDR numbers, which corresponds to the bandwidths differences between IB EDR and HDR. For read in 4K bytes message size, its performance is poorer than send and write. The reason is the read requests need to be maintained with more context overhead to wait the responses arrive [68].

5.3. MPI benchmark evaluation

This section gives the evaluation results (latency, bandwidth, and throughput) with OSU MPI Micro-Benchmarks on IB EDR and HDR.

5.3.1. Latency

The latency evaluation is shown in Fig. 6. MPI adds an extra software layer over low-level RDMA verbs. Therefore, the latency of MPI is slightly higher than that of Perftest in Section 5.2. The first observation is that running MPI on IB HDR has lower latency than on IB EDR, as expected. The second observation is that the average latency will increase with more processes usage, which scenario is closer to the real world because of the more communication overhead.



Fig. 4. Bandwidth evaluations on IB EDR and HDR.



Fig. 5. Throughput evaluations on IB EDR and HDR.



Fig. 6. OSU Micro-Benchmarks for MPI latency on multiple process pairs scenario. The first number in the bracket means the number of process pairs, and the second number is the highest bandwidth of InfiniBand.



(a) Bandwidth comparison on multiple processes scenario. The first number in the bracket means the number of processes, and the second number is the highest bandwidth of InfiniBand.



(b) Throughput on multiple processes scenario. (The legend is same as Figure 7a)

Fig. 7. The OSU Micro-Benchmarks for MPI bandwidth and throughput on multiple processes scenario.

5.3.2. Bandwidth

We also see the same trend as in Section 5.2.2 when the MPI benchmark runs on different IB interconnects in Fig. 7(a): The saturated bandwidth on IB HDR is two times the saturated bandwidth on IB EDR. The more processes are used, the earlier the bandwidth will be saturated.

5.3.3. Throughput

Although in Section 5.2.3 we can see that when the message size is the same, the throughput of Perftest on IB EDR is lower than that on IB HDR, we do not get the exact same behavior on the throughput of OSU Micro-Benchmarks for MPI which is shown in Fig. 7(b). The reason comes from different aspects. When the message size is small, MPI cannot saturate the bandwidth, so the throughput is almost the same at that stage. When the message size becomes larger, the throughput decreases rapidly and starts to saturate the bandwidth. We can observe that the throughput of IB HDR is around two times that of IB EDR, corresponding to the bandwidth ratio between two InfiniBand interconnects.

5.4. PGAS benchmark evaluation

We use OSU Micro-Benchmarks for OpenSHMEM to characterize the performance of running OpenSHMEM benchmarks on different interconnects. We use Sandia-OpenSHMEM (SOS) [116] because SOS



Fig. 8. OSU Micro-Benchmarks for OpenSHMEM for point-to-point communication on Omni-Path and InfiniBand.

is one of the native OpenSHMEM implementations. To evaluate how interconnect influences the OpenSHMEM performance, we evaluate the latency of point-to-point communication on 2 nodes (1 PE per node) and collective communication on 8 nodes (1 PE per node) on two different interconnects: Omni-Path Fabric and InfiniBand with the same 100 Gbps bandwidth. Fig. 8 shows the latency performance comparison that are divided into two parts, point-to-point operations: *put* and *get*, and Fig. 9 shows the comparison of collective operations: *broadcast* (one-to-all) and *alltoall* (all-to-all).

For the point-to-point communication in Fig. 8, the latency results on Omni-Path and InfiniBand are comparable in most cases. The latency of *get* operation on Omni-Path is slightly better than that on IB with medium message size (\sim 1K bytes). When it goes to large messages (16– 512K bytes), the latency of IB is better than that of Omni-Path for both *put* and *get* operations. Although SOS has some specific optimizations on Omni-Path, we do not observe the corresponding optimizations compared with IB for *put* and *get*, which could attribute to the heterogeneous hardware configurations, like the cache size or CPU, on two clusters.

The collective communication performance is shown in Fig. 9. InfiniBand shows a lower latency number than Omni-Path in most cases, except for *alltoall* operation in small message size. Especially, the latency number of *alltoall* operation on Omni-Path is much slower than the one on InfiniBand. In addition to the heterogeneous hardware configurations, the different network frameworks (Libfabric on Omni-Path and UCX on InfiniBand) could also be why the performance is different.

6. Discussion

This section summarizes some improvable aspects of existing benchmarks as we have observed after showing the example experiments and performance characterization results. (1) *It is not always easy to get stable numbers.* As we mentioned earlier, many benchmarks are not stable, and we may need to tune many parameters carefully or take the average of multiple rounds of running to get stable numbers. (2) *Not many benchmarks provide reference numbers.* Some benchmarks do not give reference numbers so that the users cannot evaluate whether their results are reasonable or not. Hence, we encourage our community to publish more reference numbers with the surveyed benchmarks in this paper on various interconnects as guidance. (3) *Some benchmarks lack clear instructions or specifications.* Some benchmarks assume that users are experts and do not provide clear instructions or specifications



Fig. 9. OSU Micro-Benchmarks for OpenSHMEM with Broadcast and Alltoall on Omni-Path and InfiniBand.

about benchmark installation, configuration, and usage information. (4) Not many benchmarks provide detailed warnings. We find that the proper warning messages are helpful for benchmarking. For example, Perftest can warn the users when 'CPU frequency is not max'. (5) Benchmarks usually use different ways to calculate and present numbers. Some benchmarks may use number of iterations to characterize performance, while some use time duration. Some benchmarks present the best performance numbers, while some benchmarks give the average performance numbers. Therefore, users need to compare the numbers across benchmarks carefully.

7. Related work

Besides the related survey studies and benchmarks discussed in Section 1, 3, and 4, this section summarizes more benchmarking studies.

Benchmarking distributed storage systems: The distributed storage benchmarks evaluate how the storage system serves requests for reading and writing files and objects. For example, SKB [117] supports performance benchmarking of 43 distributed storage systems. The Cloud Object Storage Benchmark [118] is for benchmarking cloud object storage services. Acquaviva et al. [119] developed a benchmark to evaluate different Cloud Distributed File Systems.

Benchmarking big data systems: Many benchmarks are proposed to evaluate the big data systems with the big data boom. HiBench [120] and MRBench [121] are designed for evaluating MapReduce systems. TextBenDS [122] is applied to evaluate the performance of Hive, Spark, and MongoDB on a textual corpus. The TPC [123] organization designed benchmark standards what were data-centric benchmark and disseminated verifiable data to the industry. Wang et al. [124] discussed the challenges of using the widely-used benchmarks (TPC-C and YCSB) for systems evaluation. DCQCN [125] and DSCP-BASEDPFC [126] introduce how to benchmark and monitor the RDMA traffic on data centers with RoCEv2 networks.

Benchmarking AI systems: Both AIBench [127–129] and MLPerf [130–132] cover a broad diversity of scenarios to evaluate the AI systems. DataPerf [133] benchmarks the datasets in machine learning and the algorithms in processing these datasets. HPC AI500 [134] is a benchmark suite to evaluate HPC systems that run real-world workloads.

Benchmarking computing systems: SPEC (Standard Performance Evaluation Corporation) [135] designed standardized benchmarks and tools to evaluate performance and energy efficiency for computing systems. PARSEC (Princeton Application Repository for Shared-Memory

Computers) [136] benchmarks the workloads and shared-memory programs for chip-multiprocessors and contains thirteen programs in different areas.

To the best of our knowledge, we are the first to extensively survey the benchmarks for hot interconnects.

8. Conclusion

This paper presents an extensive survey on hot interconnects in modern data centers and HPC clusters and associated benchmarks to help the community understand these advanced interconnects better. After introducing some representative modern interconnects and their features, we survey some commonly used micro-benchmarks and application-level benchmarks that can be used on these interconnects to measure their performance. Based on the micro-/application-level benchmarks survey, we conduct experiments on some kinds of real interconnects with the corresponding benchmarks, illustrate performance characteristics of these interconnects, and provide our interpretation of the experiment results. Considering the continuous evolution of data center interconnects and benchmarks in the future, we also discuss existing benchmarks' improvable aspects and our insights for future benchmark design and development.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by the NSF research grant CCF #2132049. Part of this research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration, and by the U.S. Department of Energy, Office of Science, under Contract DE-AC02-06CH11357. We gratefully acknowledge the computing resources provided on Bebop, an HPC cluster operated by the Laboratory Computing Resource Center at Argonne National Laboratory, USA. Part of this research was conducted using Pinnacles (NSF MRI, #2019144) at the Cyberinfrastructure and Research Technologies (CIRT) at University of California, Merced.

References

- D.K. Panda, X. Lu, D. Shankar, High-Performance Big Data Computing, The MIT Press, 2022.
- [2] Mellanox, Introducing 200G HDR InfiniBand Solutions, 2022, https: //network.nvidia.com/sites/default/files/pdf/whitepapers/WP_Introducing_ 200G HDR InfiniBand Solutions.pdf.
- [3] Slingshot, 2022, https://www.hpe.com/us/en/compute/hpc/slingshotinterconnect.html.
- [4] B. Alverson, E. Froese, L. Kaplan, D. Roweth, Cray XC Series Network, 2012, Cray Inc., White Paper WP-Aries01-1112.
- [5] Y. Ajima, T. Kawashima, T. Okamoto, N. Shida, K. Hirai, T. Shimizu, S. Hiramoto, Y. Ikeda, T. Yoshikawa, K. Uchida, et al., The Tofu Interconnect D, in: 2018 IEEE International Conference on Cluster Computing, CLUSTER, IEEE, 2018, pp. 646–654.
- [6] 100 Gigabit Ethernet, 2022, https://en.wikipedia.org/wiki/100_Gigabit_ Ethernet.
- [7] Terabit Ethernet, 2022, https://en.wikipedia.org/wiki/Terabit_Ethernet.
- [8] R. Han, X. Lu, J. Xu, On Big Data Benchmarking, in: Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware, Springer, 2014, pp. 3–18.
- [9] Q. Zhang, L. Zha, J. Lin, D. Tu, M. Li, F. Liang, R. Wu, X. Lu, A Survey on Deep Learning Benchmarks: Do We Still Need New Ones? in: International Symposium on Benchmarking, Measuring and Optimization, Springer, 2018, pp. 36–49.
- [10] X. Zhou, X. Peng, T. Xie, J. Sun, C. Xu, C. Ji, W. Zhao, Poster: Benchmarking Microservice Systems for Software Engineering Research, in: 2018 IEEE/ACM 40th International Conference on Software Engineering: Companion (ICSE-Companion), IEEE, 2018, pp. 323–324.

- [11] W. Gao, J. Zhan, L. Wang, C. Luo, D. Zheng, X. Wen, R. Ren, C. Zheng, X. He, H. Ye, et al., BigDataBench: A Scalable and Unified Big Data and AI Benchmark Suite, 2018, arXiv preprint arXiv:1802.08254.
- [12] Netperf, 2022, https://github.com/HewlettPackard/netperf.
- [13] Open Fabrics Enterprise Distribution (OFED) Performance Tests, 2022, https: //github.com/linux-rdma/perftest.
- [14] D.K. Panda, H. Subramoni, C.-H. Chu, M. Bayatpour, The MVAPICH project: Transforming Research into High-performance MPI Library for HPC Community, J. Comput. Sci. 52 (2021) 101208, http://dx.doi.org/10. 1016/j.jocs.2020.101208, URL https://www.sciencedirect.com/science/article/ pii/S1877750320305093. Case Studies in Translational Computer Science.
- [15] M.P.I. Forum, MPI: A Message-Passing Interface Standard Version 4.0, 2021, URL https://www.mpi-forum.org/docs/mpi-4.0/mpi40-report.pdf.
- [16] G. Almasi, PGAS (Partitioned Global Address Space) Languages, in: D. Padua (Ed.), Encyclopedia of Parallel Computing, Springer US, Boston, MA, 2011, pp. 1539–1545, http://dx.doi.org/10.1007/978-0-387-09766-4_210.
- [17] OmniPath, 2022, https://www.cornelisnetworks.com/products/.
- [18] S. Derradji, T. Palfer-Sollier, J.-P. Panziera, A. Poudes, F.W. Atos, The BXI Interconnect Architecture, in: 2015 IEEE 23rd Annual Symposium on High-Performance Interconnects, IEEE, 2015, pp. 18–25.
- [19] TOP 500 Supercomputer Sites, 2022, http://www.top500.org.
- [20] I.T. Association, et al., Supplement to InfiniBand Architecture Specification, Release 1 (2) (2010) 1.
- [21] IEEE std 802.3x-1997 and IEEE std 802.3y-1997 (supplement to ISO/IEC 8802-3: 1996, in: ANSI/IEEE Std 802.3, 1996 ed., IEEE Standards for Local and Metropolitan Area Networks: Supplements to Carrier Sense Multiple Access with Collision Detection (CSMA/CD).
- [22] S. Floyd, TCP and Explicit Congestion Notification, ACM SIGCOMM Comput. Commun. Rev. 24 (5) (1994) 8–23.
- [23] M. Alizadeh, B. Atikoglu, A. Kabbani, A. Lakshmikantha, R. Pan, B. Prabhakar, M. Seaman, Data Center Transport Mechanisms: Congestion Control Theory and IEEE Standardization, in: 2008 46th Annual Allerton Conference on Communication, Control, and Computing, IEEE, 2008, pp. 1270–1277.
- [24] Elastic Fabric Adapter Amazon Elastic Compute Cloud, 2022, https://docs. aws.amazon.com/AWSEC2/latest/UserGuide/efa.html.
- [25] OpenFabrics Interfaces Working Group, Libfabric OpenFabrics, 2022, URL https: //ofiwg.github.io/libfabric/.
- [26] Summit, 2022, https://www.olcf.ornl.gov/summit/.
- [27] Sierra, 2022, https://hpc.llnl.gov/hardware/compute-platforms/sierra.
- [28] InfiniBand Trade Association, InfiniBand Architecture Specification: Release 1.2.1, 2007, https://www.afs.enea.it/asantoro/Vlr1_2_1.Release_12062007.pdf.
- [29] X. Lu, D. Shankar, S. Gugnani, H. Subramoni, D.K. Panda, Impact of HPC Cloud Networking Technologies on Accelerating Hadoop RPC and HBase, in: 2016 IEEE International Conference on Cloud Computing Technology and Science (CloudCom), IEEE Computer Society, Los Alamitos, CA, USA, 2016, pp. 310–317, http://dx.doi.org/10.1109/CloudCom.2016.0057, URL https:// doi.ieeecomputersociety.org/10.1109/CloudCom.2016.0057.
- [30] M. Al-Fares, A. Loukissas, A. Vahdat, A Scalable, Commodity Data Center Network Architecture, ACM SIGCOMM Comput. Commun. Rev. 38 (4) (2008) 63–74.
- [31] A. Shpiner, Z. Haramaty, S. Eliad, V. Zdornov, B. Gafni, E. Zahavi, Dragonfly+: Low Cost Topology for Scaling Datacenters, in: 2017 IEEE 3rd International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (HiPINEB), IEEE, 2017, pp. 1–8.
- [32] J. Jiang, R. Jain, C. So-In, An Explicit Rate Control Framework for Lossless Ethernet Operation, in: 2008 IEEE International Conference on Communications, IEEE, 2008, pp. 5914–5918.
- [33] P. Newman, Backward Explicit Congestion Notification for ATM Local Area Networks, in: Proceedings of GLOBECOM'93. IEEE Global Telecommunications Conference, IEEE, 1993, pp. 719–723.
- [34] GPUDirect RDMA, 2022, https://docs.nvidia.com/cuda/gpudirect-rdma/index. html.
- [35] A. Sodani, Knights Landing (KNL): 2nd Generation Intel[®] Xeon Phi Processor, in: 2015 IEEE Hot Chips 27 Symposium, HCS, IEEE, 2015, pp. 1–24.
- [36] Cornelis Networks, 2022, https://www.cornelisnetworks.com/.
- [37] A. Tekin, A. Tuncer Durak, C. Piechurski, D. Kaliszan, F. Aylin Sungur, F. Robertsén, P. Gschwandtner, State-of-The-Art and Trends for Computing and Interconnect Network Solutions for HPC and AI, 2021, Partnership for Advanced Computing in Europe, Available Online At www.Praceri.Eu.
- [38] QSFP28 Adapter, 2022, https://www.ibm.com/docs/en/power9?topic=adpcie3-2-port-100-gbe-nic-roce-qsfp28-adapter-fc-ec31-ec3m-ccin-2cec.
- [39] NVIDIA Mellanox InfiniBand NDR 400G Architecture, 2022, https: //applieddatasystems.com/wp-content/uploads/2020/11/br-ndr-architecturebrochure.pdf.
- [40] About Fugaku, 2022, https://www.r-ccs.riken.jp/en/fugaku/about/.
- [41] J. Kim, W.J. Dally, S. Scott, D. Abts, Technology-driven, Highly-scalable Dragonfly Topology, in: 2008 International Symposium on Computer Architecture, IEEE, 2008, pp. 77–88.
- [42] Perlmutter, 2022, https://www.nersc.gov/systems/perlmutter/.
- [43] Frontier, 2022, https://www.olcf.ornl.gov/frontier/.
- [44] Aurora, 2022, https://www.alcf.anl.gov/aurora.

- [45] EI Capitan, 2022, https://www.hpe.com/us/en/compute/hpc/cray/doe-elcapitan-press-release.html.
- [46] L.G. Valiant, A Scheme for Fast Parallel Communication, SIAM J. Comput. 11 (2) (1982) 350–361.
- [47] Piz Daint, 2022, https://www.cscs.ch/computers/piz-daint/.
- [48] Cori, 2022, https://docs.nersc.gov/systems/cori/.
- [49] Y. Ajima, S. Sumimoto, T. Shimizu, Tofu: A 6D Mesh/Torus Interconnect for Exascale Computers, Computer 42 (11) (2009) 36–40.
- [50] Y. Ajima, T. Inoue, S. Hiramoto, T. Shimizu, Tofu: Interconnect for the K Computer, Fujitsu Sci. Tech. J. 48 (3) (2012) 280–285.
- [51] Y. Ajima, T. Inoue, S. Hiramoto, S. Uno, S. Sumimoto, K. Miura, N. Shida, T. Kawashima, T. Okamoto, O. Moriyama, et al., Tofu Interconnect 2: System-on-chip Integration of High-performance Interconnect, in: International Supercomputing Conference, Springer, 2014, pp. 498–507.
- [52] Fugaku, 2022, https://www.fujitsu.com/global/about/innovation/fugaku/.
- [53] B.W. Barrett, R.B. Brightwell, R.E. Grant, K.S. Hemmert, K.T. Pedretti, K.B. Wheeler, K.D. Underwood, R. Riesen, A.B. Maccabe, T. Hudson, The Portals 4.0. 2 Networking Programming Interface, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States); Intel Corp ..., 2014.
- [54] D.-G. Sun, Z.-H. Weng, Butterfly Interconnection Implementation for an n-bit Pparallel Ripple Carry Full Adder, Appl. Opt. 30 (14) (1991) 1781–1785.
- [55] A. Kashyap, X. Lu, NVMe-OAF: Towards adaptive NVMe-of for IO-intensive workloads on HPC cloud, in: Proceedings of the 31st International Symposium on High-Performance Parallel and Distributed Computing, HPDC '22, Association for Computing Machinery, New York, NY, USA, 2022, pp. 56–70, http://dx.doi.org/10.1145/3502181.3531476.
- [56] RDMA-bench: A Framework to Understand RDMA Performance, 2022, https: //github.com/efficient/rdma bench.
- [57] iPerf2 A Means to Measure Network Responsiveness and Throughput, 2022, https://sourceforge.net/projects/iperf2/.
- [58] iPerf3 The Ultimate Speed Test Tool for TCP, UDP and SCTP, 2022, https: //iperf.fr/.
- [59] Comparison Table of iPerf2 and iPerf3, 2022, https://iperf2.sourceforge.io/ IperfCompare.html.
- [60] Change between iPerf 2.0, iPerf 3.0 and iPerf 3.1, 2022, https://iperf.fr/iperfdoc.php.
- [61] qperf: Measure RDMA and IP performance, 2022, https://github.com/linuxrdma/qperf.
- [62] GPCNet, 2022, https://github.com/netbench/GPCNET.
- [63] H. Subramoni, K. Hamidouche, A. Venkatesh, S. Chakraborty, D.K. Panda, Designing MPI Library with Dynamic Connected Transport (DCT) of InfiniBand: Early Experiences, in: International Supercomputing Conference, Springer, 2014, pp. 278–295.
- [64] L. Shalev, H. Ayoub, N. Bshara, E. Sabbag, A Cloud-Optimized Transport Protocol for Elastic and Scalable HPC, IEEE Micro 40 (6) (2020) 67–73.
- [65] Transport Modes RDMA Aware Programming User Manual v1.7, 2022, https://docs.nvidia.com/networking/display/RDMAAwareProgrammingv17/ Transport+Modes#.
- [66] IEEE Standard for Cryptographic Protection of Data on Block-Oriented Storage Devices, IEEE Std 1619-2007 (2008) 1–40, http://dx.doi.org/10.1109/IEEESTD. 2008.4493450.
- [67] A. Kalia, M. Kaminsky, D.G. Andersen, Design Guidelines for High Performance RDMA systems, in: 2016 USENIX Annual Technical Conference (USENIX ATC 16), 2016, pp. 437–450.
- [68] A. Kalia, M. Kaminsky, D.G. Andersen, Using RDMA Efficiently for Key-Value Services, in: Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14, Association for Computing Machinery, New York, NY, USA, 2014, pp. 295–306, http://dx.doi.org/10.1145/2619239.2626299.
- [69] H. Lim, D. Han, D.G. Andersen, M. Kaminsky, MICA: A Holistic Approach to Fast In-Memory Key-Value Storage, in: 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14), 2014, pp. 429–444.
- [70] X. Wei, J. Shi, Y. Chen, R. Chen, H. Chen, Fast In-memory Transaction Processing Using RDMA and HTM, in: Proceedings of the 25th Symposium on Operating Systems Principles, 2015, pp. 87–104.
- [71] Wikipedia contributors, Berkeley Sockets Wikipedia, the free encyclopedia, 2022.
- [72] Wikipedia contributors, Stream Control Transmission Protocol Wikipedia, the free encyclopedia, 2022, https://en.wikipedia.org/w/index.php?title= Stream_Control_Transmission_Protocol&oldid=1091520392 (Online; Accessed 12 September 2022).
- [73] W. contributors, Unix Domain Socket Wikipedia, the free encyclopedia, 2022, https://en.wikipedia.org/w/index.php?title=Unix_domain_socket&oldid= 1100609125 (Online; Accessed 12 September 2022).
- [74] The Open Group, Data Link Provider Interface (DLPI), 2022, https://pubs. opengroup.org/onlinepubs/009638599/toc.htm (Online; accessed 12 September 2022).
- [75] S. Chunduri, T. Groves, P. Mendygral, B. Austin, J. Balma, K. Kandalla, K. Kumaran, G. Lockwood, S. Parker, S. Warren, N. Wichmann, N. Wright, GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks, SC '19, Association for Computing Machinery, New York, NY, USA, 2019, http://dx.doi.org/10.1145/3295500.3356215.

- [76] M.P.I. Forum, MPI: A Message-Passing Interface Standard Version 3.0, 2012, URL https://www.mpi-forum.org/docs/mpi-3.0/mpi30-report-book.pdf.
- [77] NAS Parallel Benchmarks, 2022, https://www.nas.nasa.gov/software/npb.html.
 [78] Intel[®] MPI Benchmarks, 2022, https://www.intel.com/content/www/us/en/ developer/articles/technical/intel-mpi-benchmarks.html
- [79] U. Consortium, D. Bonachea, G. Funck, UPC Language and Library Specifications, Version 1.3, 2013, http://dx.doi.org/10.2172/1134233, URL https: //www.osti.gov/biblio/1134233.
- [80] OpenSHMEM Application Programming Interface, v1. 5 Final, 2022.
- [81] M. Van Steen, A.S. Tanenbaum, Distributed Systems, Maarten van Steen Leiden, The Netherlands, 2017.
- [82] T. Li, H. Shi, X. Lu, HatRPC: Hint-Accelerated Thrift RPC over RDMA, in: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '21, Association for Computing Machinery, New York, NY, USA, 2021, http://dx.doi.org/10.1145/3458817. 3476191.
- [83] Apache Thrift, 2022, https://thrift.apache.org/.
- [84] rpc-perf, 2022, https://github.com/twitter/rpc-perf.
- [85] Wikipedia contributors, NVM Express Wikipedia, the free encyclopedia, 2022, https://en.wikipedia.org/w/index.php?title=NVM_Express&oldid= 1105967834 (Online; accessed 13-September-2022).
- [86] Z. Yang, J.R. Harris, B. Walker, D. Verkamp, C. Liu, C. Chang, G. Cao, J. Stern, V. Verma, L.E. Paul, SPDK: A Development Kit to Build High Performance Storage Applications, in: 2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom), IEEE, 2017, pp. 154–161.
- [87] SPDK: NVMe Driver, 2022, https://spdk.io/doc/nvme.html.
- [88] IOzone Filesystem Benchmark, 2022, https://www.iozone.org/.
- [89] Iometer project, 2022, http://www.iometer.org/.
- [90] H. Shi, X. Lu, D.K. Panda, EC-Bench: Benchmarking Onload and Offload Erasure Coders on Modern Hardware Architectures, in: C. Zheng, J. Zhan (Eds.), Benchmarking, Measuring, and Optimizing, Springer International Publishing, Cham, 2019, pp. 215–230.
- [91] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al., TensorFlow: A System for Large-Scale Machine Learning, in: 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), 2016, pp. 265–283.
- [92] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., PyTorch: An Imperative Style, High-Performance Deep Learning Library, Adv. Neural Inf. Process. Syst. 32 (2019).
- [93] PerfZero, 2022, https://github.com/tensorflow/benchmarks/tree/master/ perfZero.
- [94] TorchBench, 2022, https://github.com/pytorch/benchmark.
- [95] Meta, PARAM, 2022, https://github.com/facebookresearch/param.
- [96] M. Naumov, D. Mudigere, H.M. Shi, J. Huang, N. Sundaraman, J. Park, X. Wang, U. Gupta, C. Wu, A.G. Azzolini, D. Dzhulgakov, A. Mallevich, I. Cherniavskii, Y. Lu, R. Krishnamoorthi, A. Yu, V. Kondratenko, S. Pereira, X. Chen, W. Chen, V. Rao, B. Jia, L. Xiong, M. Smelyanskiy, Deep Learning Recommendation Model for Personalization and Recommendation Systems, 2019, CoRR arXiv:1906.00091.
- [97] Deep Learning Recommendation Model for Personalization and Recommendation Systems, 2022, https://github.com/facebookresearch/dlrm.
- [98] BaiduResearch, DeepBench, 2022, https://github.com/baidu-research/ DeepBench.
- [99] S. Jeaugey, Nccl 2.0, in: GPU Technology Conference, GTC, vol. 2, 2017.
- [100] Gloo, 2022, https://github.com/facebookincubator/gloo.
- [101] NCCL Tests, 2022, https://github.com/NVIDIA/nccl-tests.
- [102] Gloo, 2022, https://github.com/facebookincubator/gloo.
- [103] D. Kirk, et al., NVIDIA CUDA Software and GPU Parallel Computing Architecture, in: ISMM, vol. 7, 2007, pp. 103–104.
- [104] ROCm, 2022, https://www.amd.com/en/graphics/servers-solutions-rocm.
- [105] OpenACC, 2022, https://www.openacc.org/.
- [106] B.F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, R. Sears, Benchmarking Cloud Serving Systems with YCSB, in: Proceedings of the 1st ACM Symposium on Cloud Computing, Association for Computing Machinery, New York, NY, USA, 2010, pp. 143–154, http://dx.doi.org/10.1145/1807128.1807152.
- [107] The High-Performance Big Data Project, 2022, https://hibd.cse.ohio-state.edu/.
- [108] Y. Gan, Y. Zhang, D. Cheng, A. Shetty, P. Rathi, N. Katarki, A. Bruno, J. Hu, B. Ritchken, B. Jackson, K. Hu, M. Pancholi, Y. He, B. Clancy, C. Colen, F. Wen, C. Leung, S. Wang, L. Zaruvinsky, M. Espinosa, R. Lin, Z. Liu, J. Padilla, C. Delimitrou, An Open-Source Benchmark Suite for Microservices and Their Hardware-Software Implications for Cloud & Edge Systems, in: Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, Association for Computing Machinery, New York, NY, USA, 2019, pp. 3–18, http://dx.doi.org/10.1145/ 3297858.3304013.
- [109] A. Sriraman, T.F. Wenisch, μSuite: A benchmark suite for microservices, in: 2018 IEEE International Symposium on Workload Characterization, 2018, pp. 1–12, http://dx.doi.org/10.1109/IISWC.2018.8573515.
- [110] Pinnacles cluster, 2022, https://ucmerced.github.io/hpc_docs/#/Pinnacles.
- [111] Bebop, 2022, https://www.lcrc.anl.gov/systems/resources/bebop/.

- [112] Joint Laboratory for System Evaluation (JLSE), 2022, https://www.jlse.anl.gov/.
- [113] The Unified Communication X Library, 2022, http://www.openucx.org.
- [114] P. Shamis, M.G. Venkata, M.G. Lopez, M.B. Baker, O. Hernandez, Y. Itigin, M. Dubman, G. Shainer, R.L. Graham, L. Liss, et al., UCX: An Open Source Framework for HPC Network APIs and Beyond, in: 2015 IEEE 23rd Annual Symposium on High-Performance Interconnects, IEEE, 2015, pp. 40–43.
- [115] X. Lu, D. Shankar, S. Gugnani, D.K. Panda, High-Performance Design of Apache Spark with RDMA and Its Benefits on Various Workloads, in: 2016 IEEE International Conference on Big Data (Big Data), 2016, pp. 253–262, http: //dx.doi.org/10.1109/BigData.2016.7840611.
- [116] Sandia-OpenSHMEM/SOS, 2022, https://github.com/Sandia-OpenSHMEM/SOS.
- [117] K. Munegowda, N. Sanjay Kumar, Design and Implementation of Storage Benchmark Kit, in: Emerging Research in Computing, Information, Communication and Applications, Springer, 2022, pp. 45–62.
- [118] Q. Zheng, H. Chen, Y. Wang, J. Duan, Z. Huang, COSBench: A Benchmark Tool for Cloud Object Storage Services, in: 2012 IEEE Fifth International Conference on Cloud Computing, IEEE, 2012, pp. 998–999.
- [119] L. Acquaviva, P. Bellavista, A. Corradi, L. Foschini, L. Gioia, P.C.M. Picone, Cloud Distributed File Systems: A Benchmark of HDFS, Ceph, GlusterFS, and XtremeFS, in: 2018 IEEE Global Communications Conference, IEEE, 2018, pp. 1–6.
- [120] S. Huang, J. Huang, J. Dai, T. Xie, B. Huang, The HiBench Benchmark Suite: Characterization of the MapReduce-based Data Analysis, in: 2010 IEEE 26th International Conference on Data Engineering Workshops, ICDEW 2010, IEEE, 2010, pp. 41–51.
- [121] K. Kim, K. Jeon, H. Han, S.-g. Kim, H. Jung, H.Y. Yeom, MRBench: A Benchmark for Mapreduce Framework, in: 2008 14th IEEE International Conference on Parallel and Distributed Systems, IEEE, 2008, pp. 11–18.
- [122] C.-O. Truică, E.-S. Apostol, J. Darmont, I. Assent, TextBenDS: a Generic Textual Data Benchmark for Distributed Systems, Inform. Syst. Front. 23 (1) (2021) 81–100.
- [123] The Transaction Processing Performance Council (TPC), 2022, https://www.tpc. org/.
- [124] Y. Wang, M. Yu, Y. Hui, F. Zhou, Y. Huang, R. Zhu, X. Ren, T. Li, X. Lu, A Study of Database Performance Sensitivity to Experiment Settings, Proc. VLDB Endow. 15 (7) (2022) 1439–1452, http://dx.doi.org/10.14778/3523210.3523221.
- [125] Y. Zhu, H. Eran, D. Firestone, C. Guo, M. Lipshteyn, Y. Liron, J. Padhye, S. Raindel, M.H. Yahia, M. Zhang, Congestion Control for Large-Scale RDMA Deployments, in: Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, Association for Computing Machinery, New York, NY, USA, 2015, pp. 523–536, http://dx.doi.org/10.1145/2785956. 2787484.
- [126] C. Guo, H. Wu, Z. Deng, G. Soni, J. Ye, J. Padhye, M. Lipshteyn, RDMA over Commodity Ethernet at Scale, in: Proceedings of the 2016 ACM SIGCOMM Conference, Association for Computing Machinery, New York, NY, USA, 2016, pp. 202–215, http://dx.doi.org/10.1145/2934872.2934908.
- [127] AIBench, 2022, https://www.benchcouncil.org/aibench/.
- [128] W. Gao, F. Tang, J. Zhan, X. Wen, L. Wang, Z. Cao, C. Lan, C. Luo, X. Liu, Z. Jiang, Aibench Scenario: Scenario-distilling AI Benchmarking, in: 2021 30th International Conference on Parallel Architectures and Compilation Techniques, IEEE, 2021, pp. 142–158.
- [129] F. Tang, W. Gao, J. Zhan, C. Lan, X. Wen, L. Wang, C. Luo, Z. Cao, X. Xiong, Z. Jiang, et al., AlBench Training: Balanced Industry-standard AI Training Benchmarking, in: 2021 IEEE International Symposium on Performance Analysis of Systems and Software, IEEE, 2021, pp. 24–35.
- [130] P. Mattson, V.J. Reddi, C. Cheng, C. Coleman, G. Diamos, D. Kanter, P. Micikevicius, D. Patterson, G. Schmuelling, H. Tang, et al., MLPerf: An Industry Standard Benchmark Suite for Machine Learning Performance, IEEE Micro 40 (2) (2020) 8–16.
- [131] P. Mattson, C. Cheng, G. Diamos, C. Coleman, P. Micikevicius, D. Patterson, H. Tang, G.-Y. Wei, P. Bailis, V. Bittorf, et al., MLPerf Training Benchmark, Proc. Mach. Learn. Syst. 2 (2020) 336–349.
- [132] C. Banbury, V.J. Reddi, P. Torelli, J. Holleman, N. Jeffries, C. Kiraly, P. Montino, D. Kanter, S. Ahmed, D. Pau, et al., MLPerf tiny benchmark, 2021, arXiv preprint arXiv:2106.07597.
- [133] M. Mazumder, C. Banbury, X. Yao, B. Karlaš, W.G. Rojas, S. Diamos, G. Diamos, L. He, D. Kiela, D. Jurado, D. Kanter, R. Mosquera, J. Ciro, L. Aroyo, B. Acun, S. Eyuboglu, A. Ghorbani, E. Goodman, T. Kane, C.R. Kirkpatrick, T.-S. Kuo, J. Mueller, T. Thrush, J. Vanschoren, M. Warren, A. Williams, S. Yeung, N. Ardalani, P. Paritosh, C. Zhang, J. Zou, C.-J. Wu, C. Coleman, A. Ng, P. Mattson, V.J. Reddi, DataPerf: Benchmarks for Data-Centric AI Development, 2022, http://dx.doi.org/10.48550/ARXIV.2207.10062, URL https://arxiv.org/ abs/2207.10062.
- [134] Z. Jiang, W. Gao, L. Wang, X. Xiong, Y. Zhang, X. Wen, C. Luo, H. Ye, X. Lu, Y. Zhang, S. Feng, K. Li, W. Xu, J. Zhan, HPC AI500: A Benchmark Suite for HPC AI Systems, in: C. Zheng, J. Zhan (Eds.), Benchmarking, Measuring, and Optimizing, Springer International Publishing, Cham, 2019, pp. 10–22.
- [135] The Standard Performance Evaluation Corporation (SPEC), 2022, https://www. spec.org/.
- [136] C. Bienia, Benchmarking Modern Multiprocessors (Ph.D. thesis), Princeton University, 2011.

Contents lists available at ScienceDirect

KeA1 CHINESE ROOTS GLOBAL IMPACT

BenchCouncil Transactions on Benchmarks, Standards and Evaluations

journal homepage: www.keaipublishing.com/en/journals/benchcouncil-transactions-onbenchmarks-standards-and-evaluations/



Review article

A review of Blockchain Technology applications for financial services

Mohd Javaid^{a,*}, Abid Haleem^a, Ravi Pratap Singh^b, Rajiv Suman^c, Shahbaz Khan^d

^a Department of Mechanical Engineering, Jamia Millia Islamia, New Delhi, India

^b Department of Industrial & Production Engineering, Dr B R Ambedkar National Institute of Technology, Jalandhar, Punjab, India

^c Department of Industrial & Production Engineering, G.B. Pant University of Agriculture & Technology, Pantnagar, Uttarakhand, India

^d Institute of Business Management, GLA University, Mathura, UP, India

ARTICLE INFO

Keywords: Blockchain Finance Customer Banking Transaction

ABSTRACT

Financial service providers find blockchain technology useful to enhance authenticity, security, and risk management. Several institutions are adopting blockchain in trade and finance systems to build smart contracts between participants, improve efficiency and transparency, and open up newer revenue opportunities. Blockchain's unique recording capabilities make the existing clearing and settlement process redundant. Banks and other financial entities are adopting blockchain-enabled IDs to identify people. Better results come from organisations' capacity to foresee emerging trends in financial blockchain applications and develop blockchain functionality. The transfer of asset ownership and addressing the maintenance of a precise financial ledger. Measurement, communication, and analysis of financial information are three significant areas to be focussed on by accounting professionals. Blockchain clarifies asset ownership and the existence of obligations for accountants, and it has the potential to improve productivity. This paper identifies and studies relevant articles related to blockchain for finance. This paper focuses on Blockchain technology and its importance for financial services. Further takes up various tools, strategies, and featured services in Blockchain-based financial services. Finally, the paper identifies and evaluates the significant applications of Blockchain technology in financial services. Credit reports significantly impact the financial lives of customers. Recent data breaches demonstrate the superior security of blockchain-based credit reporting over conventional server-based reporting. Blockchainbased systems enable the faster, more cost-effective, and more customised issuance of digital securities. With its adoption, the market for investors can be expanded, costs for issuers can be reduced, and counterparty risk can be reduced due to the ability to customise digital financial instruments to the demands of investors. It uses mutualised standards, protocols, and shared procedures to give network users a single common source of truth. Participants in the business network can now more easily collaborate, manage data, and agree with this technology's application.

1. Introduction

Blockchain offers a decentralised system in which users can update the blockchain network. Blockchain networks are devoid of interference from financial institutions. Information can be stored on blockchains, and the digital ledger system facilitates information sharing. It can be utilised to communicate information with network users directly. A secure network for performing transactions is provided by Blockchain. Because of its robust security mechanism, blockchain technology appeals to various businesses. Each company's accounting functions are now carried out independently, and the data reconciliation process requires time and personnel [1–3]. Blockchain technology can address this issue by allowing for the real-time recording of transactional, contractual, and other information in a shared ledger. It implies that automatic verification of legal compliance will take place. The effectiveness of the organisation's operations will be significantly increased. The consumer experience might be enhanced, making data transactions and identities more secure. Blockchain is based on a distributed ledger concept that logs every transaction and maintains the timeline and veracity of that information on a secure, tamper-proof worldwide network [4,5].

As the digital revolution advances, this technology can help to maintain the balance between technology, user data, and privacy. The emphasis on confidentiality may increase while data management may also benefit. The audit process is more transparent and faster when accounting documents between counterparties are trustworthy and current. Auditor attention might be focused on more complicated and divisive problems rather than reviewing many everyday

* Corresponding author.

E-mail addresses: mjavaid@jmi.ac.in (M. Javaid), ahaleem@jmi.ac.in (A. Haleem), singhrp@nitj.ac.in (R.P. Singh), dr.r.suman@gbpuat-tech.ac.in (R. Suman), shahbaz.me12@gmail.com (S. Khan).

https://doi.org/10.1016/j.tbench.2022.100073

Received 29 September 2022; Received in revised form 17 October 2022; Accepted 17 October 2022

Available online 22 October 2022

^{2772-4859/© 2022} The Authors. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

transactions. As a result, neither auditors nor accountants were eliminated due to process automation [5,6]. Artificial intelligence and Blockchain are two very different technologies with exceptionally diverse applications. Contrarily, artificial intelligence relies on secure data that cannot be accessed or replicated and is a highly centralised service. Numerous advantages stem from their collaboration, especially in financial assistance. Blockchain technology allows for seamless communication between the parties involved in transactions, eliminating the need for recordkeeping in the order-to-cash, record-to-report, and procure-to-pay processes [7–9].

Smart contracts enabled by blockchain technology can help all parties create legally binding financial agreements that they will execute with a guarantee once all prerequisites have been satisfied. Like traditional contracts, smart contracts enforce the terms in real-time and without ambiguity on a blockchain, cutting out the intermediary and enhancing responsibility for all parties in ways regular contracts cannot [10-12]. Since a decentralised network of computers handles intermediary duties through the internet, the distributed ledger solution does not require a reliable third party. Every transaction is documented in a digital ledger, disseminated to every network member, and publicly available. The network can confirm asset ownership and transparent transactions since each network member has a legitimate copy of the ledger, making it a more secure mechanism than the existing central ledger approach [13–15].

Digital technology has created new opportunities for increased collaboration. Financial accounting procedures have been altered by cloud-based applications with analytics specialised for particular use cases such as account payables and receivables, contract administration, reporting, and others. The safest forms of payment are cash, cashier's checks, and wire transfers. However cannot tacks wire transfers, which take time, and cash. Payments made using blockchain technology do away with these problems and increase customer confidence. Technology makes real-time cash transfers between financial institutions possible, lowering friction and expediting settlement. This technology offers the potential for automation and is ideal for tracking transactions. Financial service providers can use smart contracts to monitor buyer payments and seller deliverables [16–19]. This article discusses blockchain technology, its need, tools, features and significant applications of Blockchain for finance services.

The remaining paper is structured into ten sections. Section 2 develops the understanding of blockchain, whereas Section 3 highlights the importance of blockchain in finance; Section 4 outlines the research objectives. Section 5 highlights strategies in blockchain for financial services, Section 6 discusses the Services of Blockchain Technology in the Financial area, and Section 8 provides the potential applications of blockchain in finance. Section 9 addresses the limitation of the study, Section 10 provides the future scope of the study, and the last Section 11 concludes the study.

2. Blockchain

A blockchain comprises of blocks, chains, nodes, and master nodes. Nodes are in-charge of the network's blocks. Adding blocks to the Blockchain is a challenging operation requiring mathematical problemsolving. The blockchain network's capacity to expand endlessly is constrained by the task of solving challenging mathematical puzzles. Hacking, cheating, or otherwise altering the blockchain network is virtually impossible due to the uniqueness of hash codes. Blockchain is a distributed ledger in which a copy of the ledger is kept on each connected computer. The network is called the Blockchain because it consists of interconnected blocks serving transaction records. The idea and functioning of cryptocurrencies depend on the blockchain network [20–22].

A blockchain is a digital transactional ledger. Its structure, in which separate data, known as blocks, are connected in a single list known as a chain, gives rise to the name. Blockchains have numerous uses besides keeping track of monetary transactions like those involving Bitcoin. A blockchain manages and stores data, making it hard or impossible to alter, hack, or defraud the network [23,24]. A blockchain is a network of computer systems that duplicates and distributes copies of a digital transaction record. Modern technology has long been employed in the financial industry to guarantee data and process security. Blockchain has already gained popularity in the banking sector. Blockchains allow for the safe, dependable, and verifiable conduct of financial transactions, as demonstrated by the emergence of cryptocurrencies [25, 26].

Blockchain is a digital database that enables simultaneous storage of certain operation records across numerous machines. Digital data on transactions, contracts, and contact databases are stored using this technology as a series of interconnected blocks. The absence of transparent and unambiguous financial system regulations exposes the business to common mistakes and inaccurate information interpretation [27,28]. Blockchain technology addresses the majority of these problems and dramatically lowers financial risk. The importance of Blockchain technology is becoming more widely known. It is surrounded by a small number of people trying to figure out how to adopt and use this technology's advantages in their companies. The main goal of founding banks was to unite the population and make it possible for them to engage safely and efficiently through trade and commerce. A creation that makes it easier to complete various activities on a global scale is the blockchain platform [29–31].

3. Need of blockchain

The global financial system provides services to billions of people daily while managing trillions of cash. Such ambitious objectives come with several difficulties that the finance sector has been coping with for a long time. These issues include the expenditure of having numerous stakeholders, delays, extra paperwork, and data breaches, resulting in enormous losses the business endures each year. The issues facing the global financial system may be resolved by blockchain technology [32,33]. In addition, the cost of the current stock market is increased by the presence of organisations like regulators, brokers, and the stock exchange. System effectiveness can be increased by using a decentralised management strategy for stock exchanges. There is no need for external regulators because smart contracts can be created on Blockchain. Equity markets are getting ready to decentralise as a result. Blockchain technology makes it possible to conduct every type of investor-company interaction securely and without intermediaries, lowering expenses [34-36].

The financial sector has suffered several difficulties for a very long time. Numerous issues have been solved due to tremendous technological advances, yet some innovations have brought forth new issues. Since so many fintech options are available today, it can be challenging for financial service providers to choose the one that would work best for them. As a result, they look for a comprehensive solution that can handle all the pressing problems. The use of Blockchain for financial services is quite exciting and has the potential to address substantial business issues [37-39]. Due to centralisation, the financial industry must spread a considerable sum of money over numerous businesses. Financial service providers must invest in accounting, database upkeep, central database procurement, value transfer systems, database security, labour costs, and commissions for intermediaries. Financial service providers also need to budget for each of these assets regularly because they are all recurrent. A financial service system can become expensive due to all the additional expenses [40,41].



Fig. 1. Several tools and methods in blockchain for financial domain.

4. Research objectives

Blockchain-based technologies may potentially aid in the development of capital markets. Traditional trade financing techniques have long been a source of annoyance for firms, as the lengthy processes frequently disrupt operations and make liquidity challenging to manage. Blockchain can ease cross-border operations and streamline trade finance transactions. It facilitates business transactions beyond regional or geographic boundaries in a secure manner. Blockchain is particularly suited to tasks like real-time tracking commodities as they move and change hands across the supply chain due to its immutable record. Using a blockchain gives businesses that deliver various items and possibilities. Events in a supply chain, such as allocating arriving items to different shipping containers, can be queued up using entries on a blockchain. A novel and flexible method of organising and utilising tracking data are provided by blockchain technology [42–45]. The primary research objectives of this paper are as under:

RO1: - to brief about Blockchain technology and its need for financial service;

RO2: - to discuss the tools and strategies in Blockchain for financial services;

RO3: - to study the various featured services of blockchain technology in the financial domain;

RO4: - to identify and study the significant applications of Blockchain technology in finance service.

5. Tools and strategies in blockchain for financial services

Several featured tools and methods have been observed in the broad domain of blockchain technology for financial services and its structure. Fig. 1 reflects the various tools and strategies in blockchain applications for financial services, found impressive over time. These methods and tools are pretty smart and practical for handling real-time financial issues through the concepts of Blockchain. The highlighted soft tools are parity, geth, solc, mtyhx, truffle, infura, metamask, etc. These smart and advanced tools further ensure the future of blockchain practices towards strengthening financial services and their domains [46–48].

The financial services industry has speculated about Blockchain's possibilities for the last ten years. Blockchain is essentially a ledger of recorded financial transactions. Several locations disseminate, publish, and store this ledger. When a transaction occurs, it is recorded in each ledger copy through block creation. This helps to ensure that transactions are accurately recorded [49,50]. Blockchain is practically unchangeable and incredibly secure because there are multiple copies of the ledger; to alter or falsify any section of the record, a hacker

would have to alter every copy of the ledger simultaneously, which is exceedingly challenging to do. Blockchain promotes confidence among commercial partners and allows for safe, straightforward transactions. It makes creating and using deterministic smart contracts tamper-proof programmes that automate business logic, boost efficiency, and promote trust. At every stage of the software stack, it provides market-leading technologies for granular data privacy, enabling selective data sharing in corporate networks [51–53].

Compared to regular securities, digital securities can be issued faster and more efficiently. Customised digital financial instruments can be created by issuers and directly matched to investor demand. These are fractionalised ownership of real-world assets, tokenised microeconomies, safe, scalable, and rapid asset transfers, and more. Due to these benefits, governance systems are more transparent and accountable, businesses are run more effectively, and stakeholder incentives are better aligned [54–56]. Venture capital, private equity, real estate funds, and specialist markets are under pressure to strengthen liability risk management, implement more dynamic decision-making frameworks, and address the increasing complexity of ever-changing rules. Blockchain can enhance stakeholder and asset management greatly. Blockchain applications in finance are among the most promising because digital currencies were the first thing to be stored on them [57,58].

For instance, smart contracts might be used by an insurance firm to speed up the claim's procedure. The codes built into the Blockchain will automatically assess claims when a client submits one. The smart contract will be carried out, and the client will be compensated if it is valid. Most financial institutions demand that their clients undergo an identity verification process to prevent fraud and money laundering. A digital ledger is produced when a new block is created for each transaction and added to the chain [59–61]. The potential for blockchain use in finance has increased with the significant benefits of blockchain ledgers over conventional digital ledgers. A distributed digital ledger can be created using blockchain technology. As a result, processing and storing transaction data are not needed by a single third party. Due to the absence of a centralised repository for keeping transaction data with a unique security mechanism, using Blockchain can eliminate the potential of transaction data hacking [62,63].

Blockchain applications in banking may be easier to use and less expensive. Security with blockchain technology is among the many aspects that encourage using this technology in banking. Blockchain secures its transaction ledger via encryption. As a result, the data was only accessible to those with a unique key code. Many different fintech solutions are currently available in the financial sector. As a result, financial service providers typically struggle to find the correct answer to their problems. Blockchain applications in finance can help solve some of the industry's biggest problems [64-66]. All around the world, financial services are still run in a conventional, centralised, and multilayered fashion. Most financial data is kept in centralised systems and must go via several intermediaries, and transparency is compromised. Furthermore, database security and intermediaries are the only factors that affect data security. On the other hand, even databases with the highest levels of security are susceptible to hacking and data breaches. Because no one is aware of any disparities until a data breach or other system error is found, a lack of transparency usually results in complex security issues [67-70].

Policymakers might support the creation of teaching materials on blockchain technology. Users might be able to avoid frequent blockchain frauds, and businesses might find additional capacity to deploy the technology. Policymakers may use blockchain technology to accomplish their own unique goals. This could help organisations in public, and private sectors decide whether the technology can help solve particular issues [71,72]. Organisations attempting to integrate blockchain technology with their current systems may find this to be more accessible as a result. Based on blockchain technology, policymakers could explain current laws and regulations or create new ones.



Fig. 2. Specific and typical services of blockchain in financial sectors.

This would lessen the ambiguity surrounding the potential regulation of various technological implementations, increasing the comfort level of businesses and others in embracing blockchain solutions [73–75].

Blockchain applications use two types of security keys: private and public keys. All network users have access to the public key, but only the participants in a transaction can access the private key. As a result, users inside a network can see the transaction, while participants can only access the transaction's specifics. Blockchain can preserve financial system transparency while safeguarding the private financial data of transaction stakeholders [76,77]. Nearly every industry in the world could experience a fundamental shift in how business is conducted due to blockchain technology. As the technology and its use cases develop and advance, Blockchain enables businesses to create better transparency, traceability, and operational efficiency for various business transactions and contracts. Financial institutions are looking at ways to use Blockchain to its full potential, including identifying product opportunities, resolving regulatory issues, and overcoming challenges in recognising/assessing risks and corresponding controls [78–80].

6. Various featured services of blockchain technology in the financial domain

Apart from the various developments and advancements made in the scope of blockchain practices for the finance sector, there are several featured services, too, for making financial services impactful in real-time applications. Fig. 2 exemplifies the several featured services, such as; cross-border financial transactions, trade finance platforms, proper reporting of credits, clearing and settlements, and digital identity verification. These featured services and developments in the blockchain sector will offer a capable blockchain-based financial sector [81–83].

Authorities in the financial sector and Blockchain specialists claim that by bringing visibility and reducing friction along the lengthy list of transactions that typically precede financial interactions, Blockchain is enhancing security, reducing risk, and saving money. These blockchain benefits save financial institutions expenses to some extent. Financial institutions have typically acted as a bridge between different parties, involving labour-intensive, complicated processes that slow down transactions [84,85]. Given the immutability of Blockchain, it is simple to understand why the technology is perfect for financial applications since it allows for safe, simple transactions and promotes trust between participants. Technology can automate and optimise services while lowering fraud, so even banks stand to gain significantly. Through Blockchain, a financial institution can secure identity information, and financial institutions can boost consumer confidence while preventing fraud and accelerating the verification process [86–88].

Blockchain technology is a tamper-proof log of sensitive activities that are efficiently and securely created. Therefore, it is perfect for money transfers and international payments. One can automate the entire procedure on the Blockchain, increasing the process efficiency while reducing the number of intermediaries traditionally needed in these transactions. Blockchain technology can lower the cost of payments by removing the requirement for banks to settle transactions [89, 90]. The majority of regulatory supervision relies on recordkeeping, but there is no disputing that the repercussions of not maintaining records are much harsher. As a result, firms cannot compromise on compliance. By using Blockchain, regulators and corporations may access real-time record updates, reducing delays and making it easier to spot irregularities. Blockchain's central encryption is particularly beneficial for record administration because it eliminates duplication, fraudulent entries, and other issues [91–93].

Additionally, when banks collaborate on a Blockchain, the overall costs of the Blockchain and its supporting ecosystem could be higher than the separate costs associated with handling transactions at a particular bank. However, there is a significant cost decrease because the expenditures are split among all participating institutions. Because smart contracts automatically execute after specific pre-established conditions are satisfied, they increase contractual term performance when utilised by banks and other financial organisations [94-96]. These smart contracts must be solidly grounded in the law and adhere to all applicable regulations, including cross-jurisdictional compliances if necessary. Blockchain can be helpful in complex financial asset transfers controlled by an unchangeable set of business rules that can automatical certain types of disputes. Peer-to-peer transactions are made possible by Blockchain, which is one of its main advantages because it does away with the necessity for a reliable middleman. Blockchain technology may make obsolete fee-charging intermediaries like custodian banks and clearers in the financial services sector. Blockchain offers excellent capital optimisation because banks' operational costs have been significantly reduced [97,98].

Blockchain transactions are currently automatable and programmable. These fundamental components set the stage for the financial services industry to achieve preferred-of levels of efficiency, transparency, and security. The capital market, which allows for trading financial securities like shares and bonds, is used by businesses to raise money and might be significantly impacted by Blockchain. This technology is feasible to issue and exchange assets quickly and do so for less money, and the new blockchain technology makes no middlemen. The ability to more fully customise blockchain technology [99-101]. Substituting streamlined, automated processes with laborious, paper-based ones has the potential to boost trade efficiency. A public blockchain can be a fantastic tool for collaboration because it is decentralised and cannot be held by one single entity. Bank transactions may be settled directly and tracked more efficiently with a distributed ledger technology like Blockchain than current methods [102,103].

High-tech securities are replacing conventional securities in the banking sector. By simulating recent asset transactions on the Blockchain, the business has started to test the Blockchain. However, there is still some room for the success of the blockchain solution. One of the most important advantages of Blockchain is the history of unchangeable transactions. This will help to lower the number of crimes against financial institutions. Smart contracts have been made possible by using blockchain technology. Agreements based on tamperproof algorithmic executions and decentralised consensus are known as smart contracts. A group of digital agreements contains the terms and conditions pledged by contract participants. With its programmable protocol, the smart contract makes it possible to execute and automate contract terms [104–106].

Smart contracts can save costs for information gathering and processing, contract formulation and negotiation, agreement monitoring & enforcement, and relationship management, sometimes enabling more market-based governance structures. Due to the safe storage system and assurance that actions are carried out automatically without human error or the involvement of intermediaries in the payment process, smart contracts generally have the potential to boost data trust [107-109]. Smart contracts can potentially improve open account trading parties' confidence, promote trade transaction transparency, ensure data veracity, lower the risk of errors or fraud, and simplify the exchange of payments. Blockchain applications go beyond Bitcoin and other cryptocurrencies. Any transaction or ownership information, including tangible assets (like real estate) and intangible assets (like intellectual property), can be recorded and tracked using blockchain technology. Additionally, it may automate contracts, making creating and carrying them out much simpler [110,111].

7. Blockchain technology applications in finance service

Blockchain technology shows potential applications for financial services. Transaction fees, which traditional financial institutions profit from, could be reduced or eliminated by blockchain technology. Consumers must rely on banks or outside organisations to conduct transactions involving money transfers. The implementation of blockchain technology may avoid intermediaries like banks, thereby removing fees and other expenses related to these services [112,113]. As a result, banks can experience problems with volume and transactionbased revenue. It makes it possible for private and public chains to communicate. By enabling previously unheard-of degrees of connectedness and programmability among goods, services, assets, and holdings, the digitisation of financial instruments, which includes digital assets, smart contracts, and programmable money, extends the advantages of blockchain technology. Digitisation permits asset provenance and complete transaction history in a single shared source of truth while guaranteeing data integrity. Increased automation improves operational effectiveness overall. The real-time settlement, auditing, and reporting are made possible, and processing durations, the chance of error and delay, and the number of stages and intermediaries involvement are required to reach the same levels of trust as conventional processes are all decreased [114-116]. Table 1 discusses the significant applications of blockchain technology in financial services.

Blockchain is a technology that makes a readily available, secure, and impenetrable record of online transactions. Like the internet, a blockchain is a shared record of transactions dispersed across an extensive network of users and lacks a central authority. It is made up of several data blocks, each of which contains a collection of transactions. The blocks are said to be connected and protected by cutting-edge cryptography. Major stock exchanges are looking into how Blockchain could enable almost instantaneous stock settlements by lowering transaction times and overhead [276,277]. It increases security and transparency while automating compliance with smart contracts. The financial services sector uses blockchain technology more frequently; this invention has revolutionised the global financial system and improved its efficiency and security. Numerous ways blockchain technology is enhancing the global financial services sector. The principle of building a worldwide network utilising Blockchain that is both cost-effective and possibly transparent is known as "cross-border settlements", and it is the major advantage of Blockchain. Costs are reduced while service seekers receive value additions [278-280].

Blockchain technology can reduce costs for financial services providers and end users while enhancing payment transparency, efficiency, trust, and security. Before the advent of blockchain technology, payments between banks could take up to a week to transfer. Through digital currencies and distributed ledger technologies, payments are quicker, less expensive, and more convenient [281,282]. Central banks are testing the possibility of incorporating distributed ledger technology into redesigned payments. Leaving a digital trace on the Blockchain will also help items whose provenance determines their worth. A platform for truth and trust is an immutable, unhackable, distributed ledger of digital assets. The consequences are enormous for practically every sphere of society, not just the financial services sector. There is a significant demand for blockchain software engineers, which drives up the cost of creating and maintaining blockchain-based products. Although there are many benefits to using blockchain apps, it can be challenging to integrate them with older systems and off-chain data. There are some obstacles that developers must overcome before they can link their applications with services offered by different blockchains [283–286].

8. Discussion

The financial services sector could profit significantly from the development of blockchain-based solutions. Decentralised finance was made possible by the use of Blockchain in financial services. It is a form of financing that employs smart contracts and blockchain technology to do away with middlemen from the financial services industry. Various financial institutions and organisations can benefit from Blockchain to build trust, promote transparency [287,288], and cut expenses. Blockchain technology can be used by businesses in many vital areas, including financial software and systems [289-293]. Banks are reluctant to discuss possible blockchain applications in public, but some have just ordered studies to figure out where they can. Financial technology companies have developed into a sizable segment of the financial services sector by enabling investors to open accounts with virtual advisors and make their own financial decisions. The importance of fintech in the global financial system and its relationship to Blockchain will grow together. Because investors get more value for their money and there is a balance between automation of financial services and cheaper costs, this innovation may benefit consumers.

All around the world, financial services are still centralised and multi-layered. Financial data is frequently kept in centralised databases and must go via several intermediaries, such as the front office, back office, and other places. The system lacks transparency, and data protection depends only on middlemen and database security. There is a big chance of data leaks and server hacking even if the databases are well-protected. Digital currency-based blockchain technology has the potential to be used for both domestic and international fund transfers. Because they have already made significant investments in centralised systems, banks are likely reluctant to adopt blockchain technology domestically, but they would greatly profit from it globally. International transfers benefit from the vast differences in rules and regulations and IT systems between banks in different countries.

The banking industry cannot function without auditing. Lack of openness hinders financial auditing, a time-consuming and expensive process. Without openness, intermediaries might divulge private data while conducting the audit. It is possible to see anticipated blockchainrelated legislation as another obstacle to integrating Blockchain into financial services. To decide if blockchain technology is appropriate for financial institutions and what implications it will have for businesses and consumers, regulators are examining its advantages and disadvantages. As Blockchain transforms the financial services sector, there are several possibilities for investors to get involved. One choice is to invest in businesses that run their operations using blockchain technology. Currency can be transferred with the assurance that the transaction is secure and dependable using this technology in the finance sector. The ledger is replicated many times around the network. Everybody on the network gets copies of any new transactions or added blocks.

Blockchain technology applications in financial services
--

S. no.	Applications	Description	References
1.	Fraud prevention	Blockchain technology can circumvent conventional fraud prevention techniques that require several parties to validate transactions. Blockchain is one of the best technology for any sector that benefits from the speedy movement of verifiable, fraud-free information and transactions due to its peer-to-peer network and anti-tampering features. Contracts, financial procedures, and transactions are essential to the financial sector. The use of blockchain technology can dramatically increase the efficiency of this enormous number of moving documents. Blockchain presents a problem due to its decentralised nature. The significant advancements in the financial sector are that blockchain technology can lower fraud, assure swift and safe exchanges, and ultimately aid in risk management inside the networked global financial system. Fraudsters nearly always target financial institutions. There is a chance that information will be stolen when digital payments go through payment processors and banks. Blockchains use cryptographic algorithms to process and store transaction blocks. Financial institutions might find using this cryptography less risky when processing transactions. Companies spend much time on contracts since they are a crucial component of finance self-executing contracts could considerably increase the effectiveness of this procedure self-executing contract.	[117–120]
2.	Banks and other financial institutions	Banks and other financial institutions are already using blockchain financial institutions to enhance their offerings, minimise fraud, and lower client fees. Because systems often pass through many banks on their approach to the payment's eventual destination, international money transfers have been delayed and expensive. Blockchain has the potential to make international transactions more efficient, precise, and affordable. The financial sector is aggressively adopting blockchain technology. Numerous financial companies, from smaller businesses to the biggest names in the sector, invest in blockchain stocks and support the usage of blockchain technology. For the best outcomes, blockchains need to be widely adopted. This is particularly true in the financial services sector, where numerous businesses cooperate and need a standardised way to handle transactions. For instance, each bank engaged in the transfer must have adopted Blockchain before it can be used for money transfers. In the financial industry, peer-to-peer transactions are possible with Blockchain. Because smart contracts will be able to manage transactions successfully, it means the elimination of intermediaries. Instant payment settlements are made possible as the system's layers are decreased. Cross-border payments can also be instantaneously undertaken by using blockchain payment systems.	[121-125]
3.	Calculate credit scores	The Blockchain enables new banking and finance products and services, shared operating models, more effective processes, cheaper costs, and business networks that are more open, inclusive, and secure. Financial service providers' accountants and compliance officers can provide detailed information during audits. It promotes unethical conduct, dishonesty, inconsistent compliance, and protracted auditing periods. With Blockchain, the auditing process in financial services can be expedited. Due to the immutability of blockchain data, auditors can utilise them to determine whether compliance requirements are being met and what is happening within a certain. Blockchain might also make it possible for companies to calculate credit scores using non-traditional criteria. The system transparency could be achieved by managing credit scores on a blockchain. In order to assess a person's creditworthiness, lenders can use immutable blockchain records of financial transactions. Personal information about an applicant is never compromised or made public using smart contracts. Blockchain technology in finance allows providers to keep the user's legal, personal, and public information. Using an immutable smart contract, the fund investment companies can instantly monitor users and the purpose and identity of the data's users. Blockchain in financial services can thereby increase transparency in the process of investing in funds.	[126-129]
4.	Maintaining privacy and confidentiality	This maintains privacy and confidentiality while enhancing transparency, trust, and efficiency. Its private and hybrid networks are built to manage frequent spikes in network activity and hundreds of transactions per second. In the current banking system, certain payments can take up to a week to settle. Because of the multiple layers of the current financial system, each transaction must go through at least two intermediaries before it can be resolved. In the case of cross-border payments, these intermediates could be the front and back offices of a bank or outside firms like currency exchanges. One option to guarantee security and authenticity in a centralised system is to have a large number of intermediaries, but this has a variety of drawbacks, including slow settlement times and higher costs. Banks can dramatically reduce the number of workers needed for this operation by using smart contracts to automate approval workflows and clearing computations, which will assist shorten processing times. In the worldwide financial sector, there are trillions of bank records, from personal account information to ledgers of stock market transactions because they are immutable and guard against fraud.	[130-133]
5.	Keeps track of transactions	A blockchain is a decentralised ledger that keeps track of transactions. Automated contracts, quicker and less expensive transactions, and improved security for financial service providers could all result from this technology. Although there is still a long way to go before blockchain technology is widely employed, several financial institutions are already utilising it. Blockchain-based financial organisations might be able to offer quicker money transactions. They want this techno to switch from conventional computer systems to Blockchain-whether they want it or not, based on payment gateways. Banks eventually need to switch over because traditional payment mechanisms are less secure. Automating transactions necessary for trade finance after adopting blockchain solutions would be highly advantageous to the financial services sector transactions' decentralised nature and enhance bank record security or security for banks. Conflicts about omitted or inaccurate transactions would be history because the other participants in the transaction would also obtain a record. Blockchain technology might enhance current systems in investment banks' clearing and settlement procedures. Banks must rapidly and securely record many transactions since they must detail all loans and securities on their accounts.	[134–139]

S. no.	Applications	Description	References
6.	Assurance of security and transparency	Blockchain in financial services allows for the simultaneous assurance of security and transparency. Security risks are made more likely by the system's lack of transparency because nobody knows what is happening until something goes wrong or data is compromised. Even though nobody wants their financial information to be made public, having some degree of transparency in the system is advantageous and essential for financial service providers and their clients. Ownership can be more easily tracked because a distributed ledger is practically challenging to modify. The ledger can validate information for ownership transfers and liens, improving trust. Blockchain technology enables automation, reducing the cost, complexity, and time required for transactions. Smart contracts may record when a buyer pays and delivers and take care of any problems that might occur during the transaction. Automated systems function continuously and reduce human mistakes. Startups can compete with established banks using Blockchain's low costs, encouraging financial inclusion. Due to limitations like minimum balance requirements, access restrictions, and fees associated with using a bank, many consumers are looking for alternatives to banks. Blockchain can offer a substitute for traditional banking by utilising digital identification and mobile devices.	[140–144]
7.	Helpful in money transaction	Blockchain enables people to transfer and receive money without needing several third intermediaries. Blockchain will uphold transparency, guaranteeing the integrity and morality of those who supply financial services. It is simple to spot any shady transaction activity. Auditing procedures will go more quickly because all information will be accessible. Prior to Blockchain, mediators were required to build trust and execute transactions. By employing immutable smart contracts, Blockchain in finance enables borrowers to bargain directly with lenders about the interest rate, payment schedule, and transaction length. Smart contracts enable negotiations between lenders and borrowers. The smart contract increases the total amount owed to the lender by late payment fines if borrowers do not adhere to the requirements. Banks and other financial organisations require an applicant's credit score before proceeding with a loan application. The lack of credit rating mobility is one drawback of the existing credit management system. The current credit score of an individual is no longer valid in another nation. Blockchain keeps track of information in a ledger, with each block including details on transactions and a different hash that points to the block before it. Additionally, copies of the transactions are sent to every user on the network. These features immune blockchain technology to distributed denial-of-service assaults, hackers, and other types of fraud.	[145-149]
8.	Boosts stakeholder's confidence in the transaction	By utilising sophisticated cryptography immune to hacking, Blockchain accomplishes this and boosts confidence in the transaction environment. Numerous financial applications of Blockchain exist, including managing trades and transactions. Investors would be wise to understand how Blockchain is transforming the system and how to get and manage exposure to this development as our global financial system becomes more integrated with the digital age. The financial industry must take significant risks, like building intermediary trust, to provide loans. The failure of a counterparty to fulfil its obligations and the credit risk brought on by knowledge asymmetry are additional concerns. Furthermore, commercial banks focus on loan tracking, and monitoring ultimately depends on intermediaries. Financial service companies do not have any tested risk-management techniques or methods. The financial industry is where Blockchain's most well-known uses are used. Blockchain may be connected to digital currency like Bitcoin. Cryptocurrencies use blockchain technology, a digital currency, to serve as an investment or money. These applications might lead to financial product accessibility improvements, cost savings, and other radical adjustments.	[150–153]
9.	Banking operations	In banking and financial services, the usage of Blockchain may make peer-to-peer transactions possible. Financial service providers can therefore put to rest any worries about the function of intermediaries in peer-to-peer transactions. Blockchain technology can be utilised with data immutability, improves accuracy, and smart contracts speed up transaction settlement. Most significantly, keeping track of every network transaction can help to lower credit and financial management risks. Therefore, integrating Blockchain into the finance function may help financial service providers better manage risk. Market participants see Blockchain as a way to improve reporting and compliance through real-time access to immutable asset-level information, achieve superior execution, and create new asset classes and structures through facilitated innovation and lower barriers to entry. Blockchain is a way to reduce the need for reconciliation and due diligence. Using Blockchain to execute transactions in real-time, distribute and enforce business rules, cut costs by doing away with middlemen and streamlining infrastructure, lower the risk of data loss, and increase end-to-end predictability and transparency. Blockchain technology is fundamentally easy to understand. The system consists of a shared database, where each entry has to go through peer-to-peer networks to be verified and encrypted.	[154–158]
10.	Improve client affordability	Blockchain has the potential to improve client affordability, reduce fraud risk, and increase transparency in the financial services sector. Blockchain can increase the transparency of the financial sector because users conduct transactions on a public ledger. This openness can reveal fraud and other inefficiencies, leading to problem-solving that can reduce the risk for financial institutions. The online world is becoming a haven for scammers as customers grow active there. Blockchain technology has the potential to allay this concern. Blockchain payments and transfers are quicker and more traceable than traditional banking. Blockchain might likely be helpful in the real estate industry, given the volume of activity there. Immediately confirming finances would hasten the sale of homes, prevent fraud via encryption, and ensure openness throughout the purchasing and selling process. According to proponents of blockchain technology for identity management, people would only need to disclose the bare minimum to establish their identities. Voting could become more widely available and more secure using blockchain technology. Blockchain technology would render hackers helpless since, even if they managed to access the terminal, they would be unable to influence other nodes. Blockchain technology can significantly improve the labour-intensive and prone-to-human-error process of filing taxes. This is provided that there is enough data saved on the Blockchain.	[159–163]

(continued on next page)

Financial or technological firms who view Blockchain as a disruptive technology and wish to be specialists in it can sell services to customers. This includes businesses specialising in the development of blockchain technologies. This also provides services to assist companies in integrating Blockchain for improved productivity, scalability, and expansion. Many computers in a network use a distributed, decentralised

S. no.	Applications	Description	References
11.	Useful for international payments	Everything could change as a result of Many large banks have embraced blockchain technology by many large banks for international payments, which saves time and money. Blockchain money transfers allow users to send and receive money electronically using their mobile devices, doing away with the need to travel to a money transfer location, wait in line, and pay transaction fees. Most money transfers occur through financial institutions like banks or credit card processing businesses. Worldwide, blockchain technology is witnessing revolutions. Whatever we look at, there are blockchain applications in every field. By utilising Blockchain, many firms are growing in the supply chain, healthcare, logistics, finance, and other industries. Blockchain applications' primary objective is to make business operations more transparent and effective. Businesses are starting to understand how blockchain technology may benefit them and help them expand. The need for new blockchain platforms is steadily growing as businesses experiment with various platforms by creating blockchain applications. Blockchain technology is well-known because of its openness and decentralisation. Many businesses are thinking about developing cutting-edge financial apps on blockchain systems. Immutability, security, and decentralisation are not issues with blockchain-based financial applications. Many companies are focusing their efforts on this area.	[164–168]
12.	Minimise expediting of the transfer procedure	Blockchain payments minimise or completely do away with fees by expediting the transfer procedure. When customers write a bad check to pay for goods or services, the business suffers a loss, must pay additional fees, and may need legal action to recoup the money. With blockchain-based payments, businesses can be sure that the transaction will be finished in seconds or minutes. Blockchain allows it to do away with financial intermediaries, cut expenses, and streamline several operations. It can be used by banks, for instance, to expedite document reconciliation during factoring. Additionally, it gives banks access to a pooled database of scammers, preventing money laundering. Blockchain technology enables the automation of several processes and the removal of mediators from financial operations. The financial system's efficiency is increased through lower costs. Numerous worldwide norms and regulations that control importer and exporter operations are necessary for trading financing. It involves completing forms and entering information into registries. For all parties involved in international trade, blockchain technology promotes transaction transparency. Importers and other stakeholders can save time and money by using blockchain technology to simplify the complex world of trade finance.	[169–174]
13.	Speedup transaction system	Blockchains offer a distributed, unchangeable record of transactions, which financial institutions can use for recordkeeping and reporting to regulatory bodies. Blockchain technology's speedier transaction settlements can improve a range of financial services. Vendors will receive payments sooner, lenders will be able to fund loans more quickly, and stock exchanges will be able to settle purchases and sales of stocks instantly. A long-standing issue for banks may eventually be resolved with the use of blockchain technology. Verifying the identities of its clients is the responsibility of the banks. Several businesses are already developing blockchain technology to help banks and other financial institutions establish identities. This innovative technology protects information transfers while they take place. Blockchain aims to lower transaction costs while also increasing its effectiveness and speed. The technology offers investors a wide range of prospects with its diverse applications that may be integrated into other businesses. The risk of intercepting information through different financial intermediaries increases the fraud likelihood. The cryptographic techniques used by the Blockchain allows customers to save money on traditional financial services as investors turn away from financial advisors to avoid paying more significant fees.	[175-179]
14.	Enable digital currencies	The most recent generation of blockchain-based assets is digital currencies. Although digital currency is already in use, blockchain businesses are lowering the entry barrier and offering a seamless exchange of the most well-known cryptocurrencies as a banking substitute. The promise of blockchain technology and cryptocurrencies is being recognised by many financial organisations even though banking is subject to several rules and regulations. By removing reconciliations and providing assurance over transaction history, Blockchain could broaden the scope of accounting, considering more factors now viewed as being too difficult or unreliable to measure, such as the worth of a company's data. Due to the distributed, unchangeable transaction records that blockchains offer, financial institutions can utilise them to maintain records and books while still adhering to legal requirements. Finance blockchain applications' quicker transaction settlement times can enhance current financial services. For instance, lenders will be able to fund loans more quickly, suppliers will get paid faster, and stock exchanges will be able to settle purchases and sales of securities faster. Applications for the property blockchain eliminate the need for individual and paper-based communication, cutting expenses and human error while accelerating the process. Borrowers and financial institutions gain from eliminating mediators since it enables them to provide more competitive pricing and decrease staff costs.	[180-184]
15.	Ease in the Auditing process	The auditing process could be more straightforward using Blockchain in the finance industry. The immutable blockchain records enable auditors to verify that compliance requirements are being met while offering total transparency into the financial organisation's events. Blockchain can help maintain transparency by guaranteeing financial service providers' integrity and moral conduct. With Blockchain, it is easier to follow any suspicious transaction. Additionally, Blockchain offers immediate access to all financial data, reducing the time needed for auditing procedures. By establishing connections with other organisations, markets, and economies around the world based on transparency, security, and trust, dedicated blockchain enterprise platforms have the potential to transform the way businesses function on a global scale. Through computer code, blockchain technology enables the development and automation of business logic, confirming each stage of the business process with accuracy, security, and time control. Modern private blockchain networks are built to manage spikes in network activity and hundreds of transactions per second. Numerous companies have been researching the advantages of Blockchain for years. Blockchain technology is used in letter of credit transactions to connect banks and businesses for an end-to-end digital letter of credit independent of other systems. Blockchain, an emerging technology, will significantly impact the commercial and financial industries. It sheds light on how professionals will need to advance their skills and get ready to adopt cutting-edge technology that will influence the future of business and finance.	[185-190]

S. no.	Applications	Description	References
16.	Financial services	Several financial services are being developed using blockchain technology. Investment in blockchain-based financial applications has significantly expanded. Many firms are gradually showing progress with Blockchain in the financial services ecosystem. With identical copies stored on every network computer, Blockchain Technology is a vast, distributed ledger that runs on millions of devices and records anything. Most of the nodes in a blockchain network must run some algorithms when a new transaction or an edit to an existing transaction is received to evaluate and verify the history of the proposed transaction and reach a consensus that the history and signature are valid. Only then is the new transaction record. The ledger and its existing transactions are presumed to be of high integrity since each transaction is digitally signed to confirm its legitimacy and that no one tampers with it. Blockchain enables transactions within a closed group funded by a cryptocurrency, unlike traditional value exchange systems where two or more parties must agree on the value and other economic considerations for a transaction to occur.	[191–194]
17.	Tokenisation	Blockchain technology plays a crucial role in tokenisation, the process of creating tokens on a blockchain that reflect tangible assets. The application of Blockchain in banking is growing along with the launch of Central Bank Digital Currencies. Additionally, in order to streamline fund management, financial service companies are exploring blockchain technology. Financial service companies may find it easier to manage cost management pressures if blockchain technology is used in fund administration. The development of decentralised applications is made possible by the Blockchain. Even between banks and numerous external services, including blockchains, there is a chance for interoperability. All of these aspects share the security and dependability of the Blockchain. Products that benefit from Blockchain and AI technology can be made by combining their most delicate features. A straightforward token system might also improve a platform's appeal to users and banking services. The Blockchain is mainly used to store client requests and make AI based judgments. Blockchain technology has gradually made its way into the payments industry, changing the nature of transactions. It revolutionised financial services by removing incorruptibility and fostering efficiency and simplicity by adopting new financial processes and infrastructure.	[195–198]
18.	Smart contract	A crucial component of blockchain applications is smart contracts. Users of blockchain technology must ensure that smart contract code is accurate, secure, uniform, and effective. It is crucial to test both the functionality and the controls surrounding it. Users of blockchain technology must continuously test their performance. Using smart contracts to regulate loan terms and conditions, distributed ledger technology to address communication and transaction tracking, transparency and immutable data to shed light on time-consuming reconciliations and incorrect payments, and other techniques will improve execution and servicing efficiency across the syndicated loan ecosystem. Combining machine learning data capabilities with blockchain-based smart contracts enables this dynamic. While machine learning can seek abnormalities and warn humans when identified, smart contacts can automate processes by looking for them. Financial transactions would be secure, transparent, and effective using the beneficial infrastructure. Many people view Blockchain, a decentralised technology, as a competitive alternative to historically centralised organisations like banks. As long as artificial intelligence attracts significant investment, machine learning seems to be in a position for significant expansion. The technology is projected to permeate various sectors and applications because of its versatility. Integrations between Blockchain and machine learning can potentially provide tremendous value, especially for the financial services industry. These complimentary technologies enhance security, boost performance, and regulate automation to bring about transformational change in the financial sector.	[199-202]
19.	Reduce time and expenses of financial institutions	By eliminating friction and thereby cutting down on time and expenses required by financial institutions, Blockchain can streamline various procedures, especially reconciliation, clearing, and settlement. Similar to this, the financial sector can use Blockchain to do away with the manual procedures needed to gather and distribute the various types of documents that are frequently needed for transactions, including custom forms, insurance policies, and other documents of many different kinds gathered by banks and financial services companies. Blockchain technology speeds up transaction processing. There is no need for mediators to approve financial transactions between customers due to the distributed existence. This offers a less expensive and more practical method of exchanging currencies than banks. It is the safest way to clear promises, fraud, and money laundering. In the upcoming years, blockchain adoption in the financial sector will multiply. The sector is also investigating how rapidly Blockchain instances are being used. Companies that produce consumer packaged goods have made significant investments in blockchain technology because of the promise of increased supply chain transparency. Their consumers, retailers, and food safety organisations call for more openness, and Blockchain seems to be a potential solution for this complicated industry.	[203–206]
20.	Identity management	Identity management is another great application for blockchain technology. Users can use this technology to construct their own tamper-proof, dependable, and secure digital identity. The use of passwords and user names for online accounts that are increasingly susceptible is predicted to be soon replaced with Blockchain-based IDs. People will be able to sign digital documents and perform other simple actions like logging onto websites and applications using their Blockchain identities. Any business activity dependent on transactions that occur on conventional corporate databases, which serve as the foundation for almost every financial service function, has the potential to be replaced by blockchain technology. Applications based on Blockchain is a distributed shared ledger that allows all participants to view and access information about business transactions in an unbreakable chain. Blockchain technology has the potential to upend financial industry applications because it offers permanent and tamper-proof recording of transactions in a distributed network. The advantage of Blockchain is that financial transactions may be quickly validated, cleared, and settled without the need for a central authority. Blockchain technology will significantly affect capital markets and other financial services. Blockchain will change the financial sector in the following years.	[207-210]

S. no.	Applications	Description	References
21.	Securely store financial transactions.	This technology can also be used to permanently and securely store financial transactions. Any other information can also be stored with it, creating an incorruptible distributed record that is more secure than conventional databases. There are several uses for this application. It can be used in clinics and hospitals to compile a patient's medical history. It can also be used to safeguard creative digital items like e-books, music, pictures, and intellectual property. Additionally, it can be used to register real estate or automobiles. People are generally optimistic about using Blockchain in the financial sector. Like the Internet did for offline commerce, many industry professionals think blockchain technology has the potential to revolutionise business and financial services. By redefining interactions along the value chain, reducing operational complexity, and bringing down transaction costs, blockchain technology has the potential to revolutionise corporate processes. Blockchain technology is made up of a distributed database that independently manages a constantly expanding list of transactions that are recorded in units called blocks and are safe from modification and tampering. Most blockchain networks aspire to create a database system where decentralised institutions or agents can work together to record and retain information without any entity continuously exerting market power or control. Decentralising data storage such that it cannot be owned, controlled, or altered by a single party is at the core of blockchain technology.	[211-215]
22.	Traceability of data during the transaction	The potential use of blockchain applications to promote increased efficiency and traceability of data in transactions is one of the fundamental components of Blockchain's success. Another aspect is smart contracts, which are crucial to eliminating the need for human intervention in transaction performance. Technology thus can transform financial services in general completely. The areas of clearing and settlement, loan syndication, and financial transactional mechanics like trade finance are where Blockchain is most frequently used operationally. The Blockchain is a revolutionary new currency, banking system, and transaction mechanism revolutionising how we conduct financial transactions and the entire planet. In a nutshell, a blockchain is a distributed ledger maintained across tens of thousands of computers and keeps a running log of every transaction made on every network. This changes how banking is conducted and makes hacking impossible. Financial organisations frequently implement blockchain and Al Solutions to enhance customer service. These organisations have a motivation to reduce expenses and increase value, just like any other company. By utilising both technologies in corporate power processes, those who supply financial services can increase client value while maximising their profits.	[216–219]
23.	Assist in enhancing capital markets	Systems built on the Blockchain might potentially assist in enhancing capital markets. Traditional trade finance techniques have been a significant source of annoyance for companies because the drawn-out procedures usually cause operations to be disrupted and make managing liquidity challenging. Cross-border trading involves many variables and generates much paperwork when sharing information like the place of origin and product specifics. Blockchain has the potential to ease cross-border procedures and trade finance transactions. It makes it easier for firms to transact with one another beyond regional or geographic boundaries. Because everyone involved in a blockchain transaction must agree on a transaction before it can be completed, and anybody can inspect the updated ledger following a transaction, blockchain payments are also incredibly safe. Most investment bankers request credit and financial information before investing. They have to guarantee that their money is secure. Maintaining a ledger of investments and validating accounts is relatively easy using cryptocurrency. There are several ways businesses can leverage Blockchain to attract funding, even without investment firms. Blockchain provides a wide range of intriguing financial applications. In order to offer speedy and transparent financial services, many blockchain companies are also creating cryptocurrencies and blockchain applications. Blockchain can revolutionise how banks conduct business by facilitating quicker payments, more detailed audits, and more thorough identification despite some concerns.	[220-224]
24.	Digital currency transactions	Blockchain is most frequently used for bitcoin transfers, which are digital currency transactions. Its independence from a nation or institution lowers the risk of currency inflation or devaluation. The necessity for transaction verification by a central authority is removed when using Blockchain for financial transactions. Blockchain offers a wide range of possibilities and difficulties. Applications for blockchain technology are now accessible to complete financial transactions and clear the exchange of many different financial assets. Within a short period, the transaction was finished. In contrast, a comparable transaction using a paper-based system can take a week to complete. Blockchain's distributed-ledger design has the potential to help banks in several commercial sectors, including payments, asset management, loyalty, and lending, by enhancing security, speed, and operational efficiency. Blockchain technology can significantly impact these issues, which can help regulators, financial institutions, and people in many ways. Decentralisation and immutability, two fundamental aspects of blockchain technology, are essential in this situation. Blockchain technology can increase private regulatory compliance and help regulators by making financial services more transparent. The auditing process can be made simpler and less expensive for financial organisations by obtaining a comprehensive perspective and a single source of truth for their assets and transactions.	[225-228]
25.	Facilitate trading	Trade finance, which refers to the financial goods and mechanisms that facilitate trading, is becoming more and more significant in today's world as global trade increases. By recreating the entire process on a blockchain, this technology has the potential to increase trade security, efficiency, and transparency. As a result, quicker procedures are swiftly automated, and human errors are removed, leading to the development of trust through open transparency. Open, distributed ledger technology, such as Blockchain, transactions between two parties can be efficiently and permanently recorded. A blockchain is composed of discrete data blocks that each includes a collection of related transactions linked to one another in a specific sequence. All parties can exchange a digital ledger across a computer network without a centralised authority or intermediaries. The speed of blockchain technology has numerous potential advantages for the financial industry. Greater security and transparency are equally crucial as increased productivity. Blockchain transactions' payment processing costs are lower since they do not need bank resources or authorisation from a third party. Transactions and their impact on the many contractual partners use blockchain technology's unchanging and explanatory nature. This information is unchangeable, which promotes transparency and precise analysis.	[229-235]

26. Promotes data confidentiality The financial and banking sectors could undergo a major transformation by using Blockchain. However, it might acconfidentiality have significant effects in other areas as well. It promotes data confidentiality and integrity, which raises the calibre of services. Some odd use cases include voting, sharing electricity, charging an electric automobile, and streaming music. With its promise of safe and quick transactions, Blockchain is expected to play a significant role in the digital world. Banks worldwide are now utilising blockchain technology to its full potential after realising that it is not just for cryptocurrencies. Blockchain will play a vital role because banks and financial organisations seek to make their goods and services quicker and more accessible. Banks will be unable to carry out online financial transactions without identity verification. Customers dislike the lengthy authentication	36–239]
process of several steps, however. It is possible to employ face-to-face verification, authentication, or authorisation. Every new service provider must go through these procedures for security reasons. Blockchain will facilitate quicker verification procedures for consumers and enterprises. This is due to the secure reuse of identity verification through blockchain technology for other services. Blocks of blockchain data are kept on nodes. Larger servers, laptops, and desktop computers are a few instances of nodes. They are responsible for maintaining, distributing, and storing blockchain data.	
27. Increases stock trading increases transparency in stock trading since it offers a decentralised platform. Blockchain will end the convoluted multi-person engagement process. Smart contracts can be used to record transactions on the Blockchain. Blockchain technology has the potential to restructure banks, speed up transactions, and modernise stock exchanges in the financial services sector while preserving appropriate security. The digitisation of accounting has been relatively slow. This is due, in part, to the necessity of meeting strict regulatory standards for the legitimacy and integrity of data. As a result, another industry that Blockchain has the potential to revolutionise is accounting. This technology will streamline double-entry bookkeeping procedures while also simplifying compliance. Instead of maintaining separate records based on transaction receipts, companies can input transactions directly into a joint register. The fundamental goal of blockchain technology must be kept in mind to realise this potential. Furthermore, high-quality data must be given to machine learning programmes to succeed. These technologies, whether they exist together or separately, are here to stay despite the uncertainties. Integrating blockchain technology and machine learning may be the following major forces transforming the financial sector.	.0-244]
28. Maintain financial ledger Blockchain is an accounting technology. It is focused on the ownership transfer of assets and the maintenance of a precise financial ledger. Accountants' measurement, sharing, and analysis of financial information are significant concerns. Primarily professionals are concerned with managing financial resources best or calculating or quantifying property rights and duties. Blockchain clarifies asset ownership and whether or not accountants have obligations, and it can increase productivity dramatically. Blockchain adoption must be widespread if transactions are to be speedy and easy. Because businesses in the financial industry interact with one another and need a framework to utilise when conducting transactions, it is particularly crucial. For instance, all banks engaged in the transaction must adopt blockchain technology if a bank wishes to start a fund transfer utilising the platform. The use of blockchain technology in the financial services industry is still in its infancy. We expect interoperability and advancements in transaction processing to be two major future advances. The benefits of technology for financial industry will expand due to these improvements. Implementing Blockchain in the banking and financial industry will reduce transaction processing time, decrease paperwork, and establish a safe environment. In addition, Blockchain is anticipated to create new cost-saving options. It could enhance client experiences and encourage safer data transactions.	5-248]
29. Facilitates communication Blockchain can facilitate communication between parties engaged in transactions while also disseminating proof of transaction agreements to all parties. A "distributed ledger," or shared database duplicated and synchronised by a decentralised network, is what Blockchain refers to. Users can access blockchain apps through a browser or specialised desktop programmes, and a blockchain network is linked to the internet. Blockchain has the potential to alter several other industries, including manufacturing, sustainable energy, electronic health record management, and more; the banking and finance sectors are where Blockchain has found the most adopters. Distributed ledgers use independent computers to record, share, and synchronise transactions in many electronic ledgers instead of centralising data as in a traditional ledger. The financial industry could undergo a major change due to this technology, becoming more dependable, efficient, and robust. On a public ledger, users carry out operations using Blockchain. As a result, the industry is becoming more transparent. Such transparency can help reveal inefficiencies and prompt problem-solving as the risk to financial institutions reduces. Financial organisations can use Blockchain to administer bank guarantees and track various parties. It simplifies matters for both businesses and customers.	9-253]
30. Knowledge Blockchain will similarly revolutionise finance and accounting operations to how the internet revolutionised [25: sharing and collaboration professionals. Removing the necessity for accounting transaction management among businesses and related stakeholders like banks, Blockchain and distributed ledgers may someday be the way to merge the commercial world's recordkeeping. There are numerous uses for Blockchain in the financial, accounting, insurance, trade finance, payments, settlements, and auditing industries. Many organisations are trying to get from proof of concept to practical application. Blockchain proponents think alternatives to the current time-consuming and expensive banking operations can be made using this technology. Banks are decreasing the number of intermediaries and participants in transaction processing. The sectors of finance where Blockchain could revolutionise our industry are essential to financial professionals, the banking sector, and audit heads. To help senior management make decisions, they should take the initiative and steer the conversation. Accounting and accounting will no longer require transactions to be recorded using the conventional double-entry technique and will instead use a single record to maintain track of all transactions.	4-257]

39

Table 1 ((continued).
Table 1	сопшпиеа).

S. no.	Applications	Description	References
31.	Secure domestic and international payments	Blockchain technology can enable direct and secure domestic and international payments with few or no intermediaries. Blockchain can also simplify the procedure and hence considerably save costs. Using blockchain technology, users may control and store their data on a blockchain. The industry claims that Blockchain is appealing because it offers the possibility of cost reductions, new operational efficiency, and increased accuracy and transparency. Due to decentralisation and peer-to-peer trades made possible by Blockchain and distributed ledger technologies, financial transactions are becoming more efficient on a global scale. By their very nature, these technologies have shown to be quite successful in streamlining conventional transactional procedures and allowing fast global payment options. Blockchain can be set to work in a variety of ways to achieve transactional consensus and specify known participants in the chain while excluding everyone else. The most well-known application of Blockchain is bitcoin, which uses a public ledger open to participation and anonymous. Many organisations are adopting blockchains to regulate who participants.	[258–261]
32.	Increase supply chain traceability	Blockchain technology has the potential to increase supply chain traceability and transaction transparency. Paperless trade would significantly improve the supply chain by reducing costs, eliminating documentation errors, and speeding up the delivery of papers to customers. The use of blockchain technology for supply chain monitoring is straightforward. Businesses may uncover supply chain inefficiencies and spot products in real time by getting rid of paper-based trials. As goods move from their origin to the merchant, Blockchain enables businesses and even consumers to monitor their quality. By becoming the standard for loyalty benefits, Blockchain also helps to change the shopping experience. Blockchain is a system that holds everyone to the utmost level of accountability, as revolutionary as that may sound. There will be no more omitted transactions, mistakes made by people or machines, or transactions carried out without the parties' approval. Blockchain verifies a transaction's legitimacy by recording it on the central register and a networked distributed system of registers connected by a safe validation process. Decentralisation means that no single person or organisation controls a blockchain. It is unchangeable in any way. A node, which can be any smartphone, computer, or larger server, records each transaction or block, and there is nothing that connects the nodes.	[262-266]
33.	Boosting productivity	Blockchain technology has the potential to revolutionise the financial industry by boosting productivity, transparency, and security, cutting costs, and spurring a previously unheard-of wave of innovation. One of the most popular issues in the financial sector is Blockchain, the technology that underpins the cryptocurrency Bitcoin. Many of the largest banks in the world and other critical financial organisations have already started initiatives to look into the potential of Blockchain. The middlemen or centralised authorities that have historically handled, authorised, or verified transactions are no longer necessary because of the widespread control and total transparency that blockchain technology brings to the transaction process. Users may examine all trade-related data in one place in real-time using distributed data repository built on the Blockchain. Trade finance services are being developed from the intricate web of participants, complete with manual procedures and documentation. The sell-side and buy-side settlement processes in syndicated lending are intended to be sped up using Blockchain, saving time and millions of dollars. Blockchain technology is being embraced by all sectors of society, including healthcare and entertainment, for several uses. On every level, this technology is upending conventional approaches to data security. This technology can help maintain the interaction between technology, user data, and privacy as the digital revolution develops. It can aid in data management and emphasise privacy more.	[267-270]
34.	Management of digital assets	The comprehensive management of digital assets in a reliable, traceable, automated, and predictable manner is now possible by using blockchain technology. The unique feature of Blockchain is how each "block" is connected and secured with encryption. Blockchain technology promises to make it possible to handle foreign payments quickly, securely, and affordably by using encrypted distributed ledgers that offer reliable real-time verification of transactions without the need for middlemen like correspondent banks and clearinghouses. Blockchain offers distributed ledger architecture and transaction immutability, both of which are necessary for removing the need for a trust enforcer in the ecosystem. By establishing a setting where trust is not a concern, tamper-proof distributed data enables counterparties to carry out their operations with certainty that they always possess the exact version of the truth and that its history cannot be altered. Transparency among market participants will dramatically enhance by using blockchain technology. Blockchain solutions promote the construction of an open, real-time ecosystem activity log accessible to all market participants. Blockchain has preserved an immutable record of transactions and asset ownership since the asset first appears in a transaction on the network. Due to the decreased risk, many asset types no longer require the related mitigation operations.	[271-275]

ledger blockchain to record transactions. Due to its structure and inherent characteristics, Blockchain is safe, transparent, and almost hard to alter. All transactions are accurately and chronologically recorded on a blockchain. Since everyone on the network has a copy, it is complicated to change or remove transactions or add unverified data. A concerted attack on thousands or possibly hundreds of thousands of machines would be necessary for this to succeed, which is doubtful. Consumers and financial institutions encounter several issues and difficulties when sending money abroad.

Blockchain-based payments eliminate all of these problems, increasing confidence. This technology can transform the banking sector in ways other than money transfers. A blockchain is a terrific tool for tracking transactions and ensuring accurate, secure data. Although blockchain-based payments are swift and reversible, many consumers are worried about online scams. Big transactions, in particular, are also less expensive than using banking services. The safest payment options include cash, wire transfers, and cashier's checks; however cash cannot be tracked, wire transfers take time, and cashier's checks can be faked. Nearly every area, including finance, supply chain management, and healthcare, has been significantly impacted by the radical new tendencies that Blockchain helped to establish. The public first became aware of blockchain startups two to three years ago. Nowadays, almost every modern business is seeking methods to use blockchain technology. In its most basic form, a blockchain is a distributed ledger system that functions as a decentralised ledger.

9. Limitations

Switching to blockchain technology can be expensive and timeconsuming, especially given the scarcity of qualified blockchain developers. Smaller financial companies, in particular, could be reluctant to invest in modernising their current systems. Data on a blockchain cannot be altered. Although this is a benefit of employing Blockchain, financial companies that regularly need to change stored data may find it problematic. To implement Blockchain, companies would need to alter their current procedures. Both the development

M. Javaid, A. Haleem, R.P. Singh et al.

of blockchain technology and its use in the financial services sector are still in their infancy. The two most crucial blockchain innovations to keep an eye out for are our transaction processing and interoperability advancements, as both will increase the technology's utility for financial institutions. It is doubtful that blockchains will replace current financial systems in the foreseeable future. Instead, financial institutions will test out Blockchain to gauge its potential before implementing it gradually as an addition to their current systems.

The development of Blockchain is still in its early phases. It has several difficulties due to continuing changes. On the Blockchain, data updates are not authorised. Information from other blockchains cannot be exchanged or used by one Blockchain. They are unable to converse with one another as a result. Interoperability solutions must be prioritised in blockchain networks. It is expensive and time-consuming to make the switch to blockchain technology. This is true because there are not many skilled blockchain engineers. Smaller financial institutions can therefore be reluctant to make investments in addition to current system changes. These include the potential for technology to be used to facilitate illicit behaviour, hazards to users, and the financial system brought on by the current dearth of consumer protections. Blockchain may not adequately handle most of the significant issues related to each application.

10. Future scope

There are difficulties in putting blockchain technology into practice. Despite various challenges, it can be used by hundreds of financial institutions, and blockchain stocks are accepted forms of investment. It is evident that the financial sector is aware of the potential advantages of Blockchain and that it will play a more significant role in financial services in the future. Blockchain technology uses a decentralised ledger and is a form of distributed ledger technology that is secured with public and private security keys. The public key is available to all network users, and the transaction's stakeholders access the private key. As a result, the stakeholders and transaction details will only be visible to those who possess the private key, while the transaction will be visible to all network users with the help of the public key. It will guarantee system transparency while safeguarding the private financial data of the stakeholders.

Large institutions' intercompany relationships and transactions will be revolutionised by Blockchain's ability to maintain a single source of unchangeable truth, automate intercompany transactions using smart contracts and consensus mechanisms, provide visibility across disparate systems, significantly reduce intercompany imbalances, use near-realtime reporting rather than the past, and streamline and standardise intercompany supply chain processes. By ensuring that updated data is consistent across systems and by generating an audit record of changes made to client data, Blockchain can alleviate these problems in future. Financial leaders believe this blockchain use case will increase transparency, decrease friction, speed up transactions, save money, increase security, and reduce financial crimes. In future, Blockchain also makes it possible to use tools like "smart contracts", self-executing contracts built on the Blockchain with the ability to automate human operations ranging from compliance and claim processing to dispersing a will's contents.

The financial sector has attempted to test Blockchain by duplicating its current asset transactions. This gives some flexibility in how a blockchain solution will affect efficiency but ignores how it will affect the ecosystem. Blockchain software is used as infrastructure for realtime digital asset transfer between market participants. Blockchain enables the redrawing of procedures and challenging established business model orthodoxies. This technology will significantly increase transparency among market participants, thereby levelling the playing field. Blockchain's function in the banking sector in the following years is imperative to consider as blockchain applications proliferate across numerous industries.

11. Conclusion

Blockchain technology is being adopted by factories worldwide as they get more and more connected. The future factory will comprise a vast network of equipment, accessories, goods, and value-chain partners, like equipment suppliers and logistics companies. The main goal of this technology is to develop a tamper-proof ledger for digital assets like cryptocurrencies. Blockchain applications maintain data integrity, enabling marketers to target the relevant consumer segments and musicians to obtain fair royalties for their original compositions. This technology is gaining ground in banking payments. People exchange money mainly through their bank accounts; therefore, payments are crucial. Banks have long been at the forefront of the digital revolution, accepting disruptive developments in exchange for reliable payments and printing their digital currencies. Blockchain technology allows banks to track every transaction in real-time. This technology will enable banks to settle transactions on a public blockchain. Banking executives need to fulfil several requirements to become a widely used technology in the banking sector. Blockchain's ability to share information and temporarily make the property available to someone else would dramatically change our mobility. By utilising intelligent contracts over the Blockchain, it would be feasible to directly pay for and utilise a car while finding solutions to issues like electromobility. Smart contracts can be used by businesses using Blockchain in finance to upload invoices to the Blockchain. The Blockchain can contain data like payment due dates, amounts, and client information. The smart contract updates the invoice status to paid when the customer pays the bill and notifies the businesses that the payment has been received. Blockchain in financial services can assess a client's trustworthiness before trading. In the future, blockchain will play an important role and manage various activities in the finance sector.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- M. Kowalski, Z.W. Lee, T.K. Chan, Blockchain technology and trust relationships in trade finance, Technol. Forecast. Soc. Change 166 (2021) 120641.
- [2] S. Trivedi, K. Mehta, R. Sharma, Systematic literature review on application of blockchain technology in E-finance and financial services, J. Technol. Manag. Innov. 16 (3) (2021) 89–102.
- [3] P. Treleaven, R.G. Brown, D. Yang, Blockchain technology in finance, Computer 50 (9) (2017) 14–17.
- [4] V. Chang, P. Baudier, H. Zhang, Q. Xu, J. Zhang, M. Arami, How Blockchain can impact financial services-The overview, challenges and recommendations from expert interviewees, Technol. Forecast. Soc. Change 158 (2020) 120166.
- [5] M.J. Lahkani, S. Wang, M. Urbański, M. Egorova, Sustainable B2B E-commerce and blockchain-based supply chain finance, Sustainability 12 (10) (2020) 3968.
- [6] A.I.N.O. Nordgren, E.L.L.E.N. Weckström, M.I.N.N.A. Martikainen, O.M. Lehner, Blockchain in the fields of finance and accounting: a disruptive technology or an overhyped phenomenon, ACRN J. Finance Risk Perspect. 8 (2019) 47–58.
- [7] M. Poongodi, A. Sharma, V. Vijayakumar, V. Bhardwaj, A.P. Sharma, R. Iqbal, R. Kumar, Prediction of the price of ethereum blockchain cryptocurrency in an industrial finance system, Comput. Electr. Eng. 81 (2020) 106527.
- [8] A.V. Bogucharskov, I.E. Pokamestov, K.R. Adamova, Z.N. Tropina, Adoption of blockchain technology in trade finance process, J. Rev. Global Econ. 7 (2018) 510–515.
- [9] E. Hofmann, U.M. Strewe, N. Bosia, Discussion—how does the full potential of blockchain technology in supply chain finance look like? in: Supply Chain Finance and Blockchain Technology, Springer, Cham, 2018, pp. 77–87.
- [10] X. Zhu, D. Wang, Research on blockchain applications for E-commerce, finance and energy, IOP Conf. Ser.: Earth Environ. Sci. 252 (4) (2019) 042126.
- [11] M. Li, S. Shao, Q. Ye, G. Xu, G.Q. Huang, Blockchain-enabled logistics finance execution platform for capital-constrained E-commerce retail, Robot. Comput.-Integr. Manuf. 65 (2020) 101962.
- [12] T. Katona, Decentralized finance: the possibilities of a blockchain money lego system, Financ. Econ. Rev. 20 (1) (2021) 74–102.

- [13] J. Chen, T. Cai, W. He, L. Chen, G. Zhao, W. Zou, L. Guo, A blockchain-driven supply chain finance application for auto retail industry, Entropy 22 (1) (2020) 95.
- [14] A. Rijanto, Blockchain technology adoption in supply chain finance, J. Theor. Appl. Electron. Commerce Res. 16 (7) (2021) 3078–3098.
- [15] L. Dong, Y. Qiu, F. Xu, Blockchain-enabled deep-tier supply chain finance, Manuf. Serv. Oper. Manag. (2022).
- [16] R. Wang, Application of blockchain technology in supply chain finance in beibu gulf region, in: 2021 International Wireless Communications and Mobile Computing, IWCMC, IEEE, 2021, pp. 1860–1864.
- [17] B. Maurer, Re-risking in realtime. On possible futures for finance after the blockchain, Behemoth- J. Civil. 9 (2) (2016) 82–96.
- [18] L. Wang, X.R. Luo, F. Lee, J. Benitez, Value creation in blockchain-driven supply chain finance, Inf. Manag. (2021) 103510.
- [19] G. Hilary, L.X. Liu, Blockchain and other distributed ledger technologies in finance, in: The Palgrave Handbook of Technological Finance, Palgrave Macmillan, Cham, 2021, pp. 243–268.
- [20] B.M. Till, A.W. Peters, S. Afshar, J.G. Meara, From blockchain technology to global health equity: can cryptocurrencies finance universal health coverage? BMJ Global Health 2 (4) (2017) e000570.
- [21] X. Zhang, M. Aranguiz, D. Xu, X. Zhang, X. Xu, Utilizing blockchain for better enforcement of green finance law and regulations, in: Transforming Climate Finance and Green Investment with Blockchains, Academic Press, 2018, pp. 289–301.
- [22] B. Scott, J. Loonam, V. Kumar, Exploring the rise of blockchain technology: Towards distributed collaborative organizations, Strateg. Change 26 (5) (2017) 423–428.
- [23] V.S. Anoop, J. Goldston, Decentralized finance to hybrid finance through blockchain: a case-study of acala and current, J. Bank. Financ. Technol. (2022) 1–7.
- [24] M. O'Dair, R. Owen, Monetizing new music ventures through blockchain: Four possible futures? Int. J. Entrepreneurship Innov. 20 (4) (2019) 263–276.
- [25] A. Tapscott, D. Tapscott, How blockchain is changing finance, Harv. Bus. Rev. 1 (9) (2017) 2–5.
- [26] L. Zhang, Y. Xie, Y. Zheng, W. Xue, X. Zheng, X. Xu, The challenges and countermeasures of blockchain in finance and economics, Syst. Res. Behav. Sci. 37 (4) (2020) 691–698.
- [27] Y. Chen, C. Bellavitis, Blockchain disruption and decentralised finance: The rise of decentralized business models, J. Bus. Ventur. Insights 13 (2020) e00151.
- [28] G. Caldarelli, J. Ellul, The blockchain oracle problem in decentralized finance—A multivocal approach, Appl. Sci. 11 (16) (2021) 7572.
- [29] M. Du, Q. Chen, J. Xiao, H. Yang, X. Ma, Supply chain finance innovation using blockchain, IEEE Trans. Eng. Manage. 67 (4) (2020) 1045–1058.
- [30] C. Harwick, J. Caton, What's holding back blockchain finance? On the possibility of decentralized autonomous finance, Q. Rev. Econ. Finance (2020).
- [31] I. Eyal, Blockchain technology: Transforming libertarian cryptocurrency dreams to finance and banking realities, Computer 50 (9) (2017) 38–49.
- [32] Y. Tian, Z. Lu, P. Adriaens, R.E. Minchin, A. Caithness, J. Woo, Finance infrastructure through blockchain-based tokenization, Front. Eng. Manag. 7 (4) (2020) 485–499.
- [33] Q. Gan, R.Y.K. Lau, J. Hong, A critical review of blockchain applications to banking and finance: a qualitative thematic analysis approach, Technol. Anal. Strateg. Manag. (2021) 1–17.
- [34] K. Schulz, M. Feist, Leveraging blockchain technology for innovative climate finance under the Green Climate fund, Earth Syst. Gov. 7 (2021) 100084.
- [35] J. Li, S. Zhu, W. Zhang, L. Yu, Blockchain-driven supply chain finance solution for small and medium enterprises, Front. Eng. Manag. 7 (4) (2020) 500–511.
- [36] A. Poberezhna, Addressing water sustainability with blockchain technology and green finance, in: Transforming Climate Finance and Green Investment with Blockchains, Academic Press, 2018, pp. 189–196.
- [37] S. Karim, M.R. Rabbani, H. Bawazir, Applications of blockchain technology in the finance and banking industry beyond digital currencies, in: Blockchain Technology and Computational Excellence for Society 5.0, IGI Global, 2022, pp. 216–238.
- [38] C. Dong, C. Chen, X. Shi, C.T. Ng, Operations strategy for supply chain finance with asset-backed securitization: Centralization and blockchain adoption, Int. J. Prod. Econ. 241 (2021) 108261.
- [39] C. Fisch, M. Meoli, S. Vismara, Does blockchain technology democratize entrepreneurial finance? An empirical comparison of ICOs, venture capital, and REITs, Econ. Innov. New Technol. 31 (1–2) (2022) 70–89.
- [40] M. Campbell-Verduyn, M. Hütten, Beyond scandal? Blockchain technologies and the legitimacy of post-2008 finance, Finance Soc. 5 (2) (2019) 126–144.
- [41] R. Jiang, Y. Kang, Y. Liu, Z. Liang, Y. Duan, Y. Sun, J. Liu, A trust transitivity model of small and medium-sized manufacturing enterprises under blockchain-based supply chain finance, Int. J. Prod. Econ. 247 (2022) 108469.
- [42] T. Zhang, J. Li, X. Jiang, Analysis of supply chain finance based on blockchain, Procedia Comput. Sci. 187 (2021) 1–6.
- [43] A. Anjum, M. Sporny, A. Sill, Blockchain standards for compliance and trust, IEEE Cloud Comput. 4 (4) (2017) 84–90.
- [44] S. Demirkan, I. Demirkan, A. McKee, Blockchain technology in the future of business cyber security and accounting, J. Manag. Anal. 7 (2) (2020) 189–208.

- [45] Q. He, Application of blockchain technology in commercial banks, in: E3S Web of Conferences, vol. 235, EDP Sciences, 2021, p. 03070.
- [46] K. George, P.N. Patatoukas, The blockchain evolution and revolution of accounting, in: Information for Efficient Decision Making: Big Data, Blockchain and Relevance, 2021, pp. 157–172.
- [47] A. Abdelmaboud, A.I.A. Ahmed, M. Abaker, T.A.E. Eisa, H. Albasheer, S.A. Ghorashi, F.K. Karim, Blockchain for IoT applications: taxonomy, platforms, recent advances, challenges and future research directions, Electronics 11 (4) (2022) 630.
- [48] W. Riggs, V. Vyas, M. Sethi, Blockchain and distributed autonomous community ecosystems: opportunities to democratize finance and delivery of transport, housing, urban greening and community infrastructure, in: Housing, Urban Greening and Community Infrastructure (February 1, 2022), 2022.
- [49] S. Wang, P. Xiao, H. Chai, X. Tu, Q. Sun, H. Cai, F.Y. Wang, Research on construction of supply chain financial platform based on blockchain technology, in: 2021 IEEE 1st International Conference on Digital Twins and Parallel Intelligence, DTPI, IEEE, 2021, pp. 42–45.
- [50] Y.P.A. Sanjaya, M.A. Akhyar, Blockchain and smart contract applications can be a support for msme supply chain finance based on sharia crowdfunding, Blockchain Front. Technol. 2 (1) (2022) 44–49.
- [51] G. Chen, B. Xu, M. Lu, N.S. Chen, Exploring blockchain technology and its potential applications for education, Smart Learn. Environ. 5 (1) (2018) 1–10.
- [52] X. Meng, Risk assessment and analysis in supply chain finance based on blockchain technology, J. Sensors 2022 (2022).
- [53] H. Qiu, X. Wu, S. Zhang, V.C. Leung, W. Cai, ChainIDE: A cloud-based integrated development environment for cross-blockchain smart contracts, in: 2019 IEEE International Conference on Cloud Computing Technology and Science, CloudCom, IEEE, 2019, pp. 317–319.
- [54] K. Fanning, D.P. Centers, Blockchain and its coming impact on financial services, J. Corp. Account. Finance 27 (5) (2016) 53–57.
- [55] Y. Tian, Y. Zhang, R.E. Minchin, A. Asutosh, C. Kan, An innovative infrastructure financing instrument: Blockchain-based tokenization, in: Construction Research Congress, 2020, pp. 731–740.
- [56] Y. Wang, Research on supply chain financial risk assessment based on blockchain and fuzzy neural networks, Wirel. Commun. Mob. Comput. 2021 (2021).
- [57] A.K. Saha, S. Barua, S.D. Nath, Blockchain for financial technology: challenges and opportunities for India, in: The Impact of Artificial Intelligence on Governance, Economics and Finance, Vol. I, 2021, pp. 247–261.
- [58] K.W. Prewett, G.L. Prescott, K. Phillips, Blockchain adoption is inevitable— Barriers and risks remain, J. Corp. Account. Finance 31 (2) (2020) 21–28.
- [59] A. Tandon, P. Kaur, M. Mäntymäki, A. Dhir, Blockchain applications in management: A bibliometric analysis and literature review, Technol. Forecast. Soc. Change 166 (2021) 120649.
- [60] B.B. Sundaram, M.D.A. Wako, A. Pandey, M.A.S. Genale, M.V. Janga, P. Karthika, Supply chain management finance application in bank official website using blockchain, in: 2022 6th International Conference on Intelligent Computing and Control Systems, ICICCS, IEEE, 2022, pp. 812–817.
- [61] J. Abou Jaoude, R.G. Saade, Blockchain applications-usage in different domains, IEEE Access 7 (2019) 45360–45381.
- [62] T. Ankenbrand, A structure for evaluating the potential of blockchain use cases in finance, Perspect. Innov. Econ. Bus. 17 (2) (2017) 77–95.
- [63] W. Gao, C. Su, Analysis of earnings forecast of blockchain financial products based on particle swarm optimization, J. Comput. Appl. Math. 372 (2020) 112724.
- [64] T. Bhogal, A. Trivedi, Blockchain technology and trade finance, in: International Trade Finance, Palgrave Macmillan, Cham, 2019, pp. 303–312.
- [65] R.G. Wiatt, From the mainframe to the blockchain, Strateg. Finance 100 (7) (2019) 26–35.
- [66] T.C. Chen, Y.S. Liang, P.S. Ko, J.C. Huang, Optimization model of cross-border E-commerce payment security by blockchain finance, Wirel. Commun. Mob. Comput. 2021 (2021).
- [67] O. Alao, P. Cuffe, Structuring special purpose vehicles for financing renewable generators on a blockchain marketplace, IEEE Trans. Ind. Appl. 58 (2) (2021) 1478–1489.
- [68] L. Guo, J. Chen, S. Li, Y. Li, J. Lu, A blockchain and IoT-based lightweight framework for enabling information transparency in supply chain finance, Digit. Commun. Netw. (2022).
- [69] A. Gorkhali, R. Chowdhury, Blockchain and the evolving financial market: A literature review, J. Ind. Integr. Manag. 7 (01) (2022) 47–81.
- [70] D. Younus, A. Muayad, M. Abumandil, Role of smart contract technology blockchain services in finance and banking systems: concept and core values, in: Mohanad, Role of Smart Contract Technology Blockchain Services in Finance and Banking Systems: Concept and Core Values (April 8, 2022), 2022.
- [71] G. Marotta, C.D. Au, Blockchain in corporate finance: a review of use-case, opportunities and risks, Int. J. Appl. Res. Manag. Econ. 4 (1) (2021) 1–12.
- [72] A. Brunner, N. Abderrahmane, A. Muralidharan, P. Halfpap, O. Sume, S. Zimprich, Trade finance disrupted: a blockchain use case, J. Financ. Transform. 45 (2017) 41–48.

- [73] S. Kaur, N. Jaswal, H. Singh, Blockchain technology: principles, applications, and advantages of blockchain technology in the digital era, in: Applications, Challenges, and Opportunities of Blockchain Technology in Banking and Insurance, IGI Global, 2022, pp. 204–212.
- [74] W. Zheng, Z. Zheng, X. Chen, K. Dai, P. Li, R. Chen, NutBaaS: a blockchain-as-a-service platform, IEEE Access 7 (2019) 134422-134433.
- [75] L. Lotti, Contemporary art, capitalization and the blockchain: On the autonomy and automation of art's value, Finance Soc. 2 (2) (2016) 96–110.
- [76] J. Liu, J. Li, J. Wang, M.M. Uddin, B. Zhang, Research on the application of blockchain technology in coal supply chain finance, Sustainability 14 (16) (2022) 10099.
- [77] I.A. Ibrahim, J. Truby, Governance in the era of blockchain technology in Qatar: a roadmap and a manual for Trade Finance, J. Bank. Regul. (2021) 1–20.
- [78] B. Kucukaltan, R. Kamasak, B. Yalcinkaya, Z. Irani, Investigating the themes in supply chain finance: the emergence of blockchain as a disruptive technology, Int. J. Prod. Res. (2022) 1–20.
- [79] R. Patel, M. Migliavacca, M. Oriani, Blockchain in Banking and Finance: is the best yet to come? A bibliometric review, Res. Int. Bus. Finance (2022) 101718.
- [80] A.F. Aysan, B. Sadriu, H. Topuz, Blockchain futures in cryptocurrencies, trade and finance: a preliminary assessment, Bull. Monet. Econ. Bank. 23 (4) (2019) 525–542.
- [81] H. Sheth, J. Dattani, Overview of blockchain technology, Asian J. Convergence Technol. (2019) ISSN-2350-1146.
- [82] S. Manski, Building the blockchain world: Technological commonwealth or just more of the same? Strateg. Change 26 (5) (2017) 511–522.
- [83] B. Khan, T. Syed, Recent progress in blockchain in public finance and taxation, in: 2019 8th International Conference on Information and Communication Technologies, ICICT, IEEE, 2019, pp. 36–41.
- [84] M. Xu, S. Ma, G. Wang, Differential game model of information sharing among supply chain finance based on blockchain technology, Sustainability 14 (12) (2022) 7139.
- [85] N. Jain, R.R. Sedamkar, A blockchain technology approach for the security and trust in trade finance, in: 2020 14th International Conference on Innovations in Information Technology, IIT, IEEE, 2020, pp. 192–197.
- [86] E. Bulut, Blockchain-based entrepreneurial finance: success determinants of tourism initial coin offerings, Curr. Issues Tourism 25 (11) (2022) 1767–1781.
- [87] K. Zheng, L.J. Zheng, J. Gauthier, L. Zhou, Y. Xu, A. Behl, J.Z. Zhang, Blockchain technology for enterprise credit information sharing in supply chain finance, J. Innov. Knowl. 7 (4) (2022) 100256.
- [88] S. Paul, A. Adhikari, I. Bose, White knight in dark days? Supply chain finance firms, blockchain, and the COVID-19 pandemic, Inf. Manag. 59 (6) (2022) 103661.
- [89] Y. Alabbasi, Governance and legal framework of blockchain technology as a digital economic finance, Int. J. Innov. Digit. Econ. 11 (4) (2020) 52–62.
- [90] E. Hofmann, U.M. Strewe, N. Bosia, Concept—Where are the opportunities of blockchain-driven supply chain finance? in: Supply Chain Finance and Blockchain Technology, Springer, Cham, 2018, pp. 51–75.
- [91] D. Li, D. Han, N. Crespi, R. Minerva, K.C. Li, A blockchain-based secure storage and access control scheme for supply chain finance, J. Supercomput. (2022) 1–30.
- [92] N.R. Mosteanu, A. Faccia, Digital systems and new challenges of financial management–FinTech, XBRL, blockchain and cryptocurrencies, Quality-Access Success J. 21 (174) (2020) 159–166.
- [93] V.C. Gupta, M. Agarwal, A. Mishra, When trade finance meets Blockchain technology, Int. J. Innov. Sci. Res. Technol. 4 (10) (2019) 342–346.
- [94] J.N. Dewey, R. Hill, R. Plasencia, Blockchain and 5G-enabled Internet of Things (IoT) will redefine supply chains and trade finance, Proc. Secured Lender (2018) 43–45.
- [95] M. Campbell-Verduyn, F. Giumelli, Enrolling into exclusion: African blockchain and decolonial ambitions in an evolving finance/security infrastructure, J. Cultural Economy (2022) 1–20.
- [96] R. Ashima, A. Haleem, S. Bahl, M. Javaid, S.K. Mahla, S. Singh, Automation and manufacturing of smart materials in Additive Manufacturing technologies using Internet of Things towards the adoption of Industry 4.0, Mater. Today: Proc. 45 (2021) 5081–5088.
- [97] S. Kesharwani, M.P. Sarkar, S. Oberoi, Impact of blockchain technology and 5G/IoT on supply chain management and trade finance, Cybernomics 1 (1) (2019) 18–20.
- [98] E.A. Boakye, H. Zhao, B.N.K. Ahia, Emerging research on blockchain technology in finance; a conveyed evidence of bibliometric-based evaluations, J. High Technol. Manag. Res. 33 (2) (2022) 100437.
- [99] M.K. Al Kemyani, J. Al Raisi, A.R.T. Al Kindi, I.Y. Al Mughairi, C.K. Tiwari, Blockchain applications in accounting and finance: qualitative Evidence from the banking sector, J. Res. Bus. Manag. 10 (4) (2022) 28–39.
- [100] R.E. Spiridonov, Restricted move-a smart contract description language to create and control finance instruments on DFinance blockchain platform, J. Phys.: Conf. Ser. 1864 (1) (2021) 012109.
- [101] J. Thomason, M. Ahmad, P. Bronder, E. Hoyt, S. Pocock, J. Bouteloupe, et al., Blockchain—powering and empowering the poor in developing countries, in: Transforming Climate Finance and Green Investment with Blockchains, Academic Press, 2018, pp. 137–152.

- [102] C. Vijai, S.M. Suriyalakshmi, D. Joyce, The blockchain technology and modern ledgers through blockchain accounting, Adalya J. 8 (12) (2019).
- [103] J. Li, Y. Wang, Y. Li, Q.L. Li, A simple survey for supply chain finance risk management with applications of blockchain, in: International Conference of Celebrating Professor Jinhua Cao's 80th Birthday, Springer, Singapore, 2019, pp. 116–133.
- [104] L. Ning, Y. Yuan, How blockchain impacts the supply chain finance platform business model reconfiguration, Int. J. Logist. Res. Appl. (2021) 1–21.
- [105] W. Yang, W. Ziyang, Z. Xiaohao, Y. Jianming, The optimisation research of Blockchain application in the financial institution-dominated supply chain finance system, Int. J. Prod. Res. (2022) 1–21.
- [106] J. Duchenne, Blockchain and smart contracts: Complementing climate finance, legislative frameworks, and renewable energy projects, in: Transforming Climate Finance and Green Investment with Blockchains, Academic Press, 2018, pp. 303–317.
- [107] J. Kaur, S. Kumar, B.E. Narkhede, M. Dabić, A.P.S. Rathore, R. Joshi, Barriers to blockchain adoption for supply chain finance: the case of Indian SMEs, Electron. Commerce Res. (2022) 1–38.
- [108] V. Natanelov, S. Cao, M. Foth, U. Dulleck, Blockchain smart contracts for supply chain finance: Mapping the innovation potential in Australia-China beef supply chains, J. Ind. Inf. Integr. 30 (2022) 100389.
- [109] B. Scott, How Can Cryptocurrency and Blockchain Technology Play a Role in Building Social and Solidarity Finance? (No. 2016-1). UNRISD Working Paper, 2016.
- [110] Y. Chen, C. Bellavitis, Decentralized Finance: Blockchain Technology and the Quest for an Open Financial System, Stevens Institute of Technology School of Business Research Paper, 2019.
- [111] M. Osmani, R. El-Haddadeh, N. Hindi, M. Janssen, V. Weerakkody, Blockchain for next generation services in banking and finance: cost, benefit, risk and opportunity analysis, J. Enterp. Inf. Manag. (2020).
- [112] S.E. Chang, H.L. Luo, Y. Chen, Blockchain-enabled trade finance innovation: A potential paradigm shift on using letter of credit, Sustainability 12 (1) (2019) 188.
- [113] B. Jessel, A. DiCaprio, Can blockchain make trade finance more inclusive? J. Financ. Transform. 47 (2018) 35–50.
- [114] J. Liu, L. Yan, D. Wang, A hybrid blockchain model for trusted data of supply chain finance, Wirel. Pers. Commun. (2021) 1–25.
- [115] J. Huertas, H. Liu, S. Robinson, Eximchain: Supply Chain Finance Solutions on a Secured Public, Permissioned Blockchain Hybrid, Eximchain White Paper, 2018, 13.
- [116] Y. Wang, D.K. Kim, D. Jeong, A survey of the application of blockchain in multiple fields of financial services, J. Inf. Process. Syst. 16 (4) (2020) 935–958.
- [117] M. Kherbouche, G. Pisoni, B. Molnár, Model to program and blockchain approaches for business processes and workflows in finance, Appl. Syst. Innov. 5 (1) (2022) 10.
- [118] R.K. Nurmukhametov, P.D. Stepanov, T.R. Novikova, Blockchain technology and its application in trade finance, Financ. Anal.: Probl. Solut. 11 (2) (2018) 179–190.
- [119] M. Kumar, N. Nikhil, R. Singh, Decentralising finance using decentralised blockchain oracles, in: 2020 International Conference for Emerging Technology, INCET, IEEE, 2020, pp. 1–4.
- [120] L. Liu, Y. Li, T. Jiang, Optimal strategies for financing a three-level supply chain through blockchain platform finance, Int. J. Prod. Res. (2021) 1–18.
- [121] R. Owen, C. Mac an Bhaird, J. Hussain, T. Botelho, Blockchain and other innovations in entrepreneurial finance: Implications for future policy, Strateg. Change 28 (1) (2019) 5.
- [122] Y. Zhao, Research on personal credit evaluation of internet finance based on blockchain and decision tree algorithm, EURASIP J. Wireless Commun. Networking 2020 (1) (2020) 1–12.
- [123] A. Haleem, M. Javaid, R.P. Singh, S. Rab, R. Suman, Hyperautomation for the enhancement of automation in industries, Sensors Int. 2 (2021) 100124.
- [124] C.I. Bogusz, C. Laurell, C. Sandström, Tracking the digital evolution of entrepreneurial finance: the interplay between crowdfunding, blockchain technologies, cryptocurrencies, and initial coin offerings, IEEE Trans. Eng. Manage. 67 (4) (2020) 1099–1108.
- [125] S.R. Kursh, N.A. Gold, Adding fintech and blockchain to your curriculum, Bus. Ed. Innov. J. 8 (2) (2016) 6–12.
- [126] R. Yaksick, Overcoming supply chain finance challenges via blockchain technology, in: Disruptive Innovation in Business and Finance in the Digital World, Emerald Publishing Limited, 2019.
- [127] S. Zhang, D. Zhang, J. Zheng, W. Aerts, Does policy uncertainty of the blockchain dampen ICO markets? Account. Finance 61 (2021) 1625–1637.
- [128] I. Antoniadis, S. Kontsas, K. Spinthiropoulos, Blockchain applications in marketing, in: The Proceedings of 7th ICCMI, 2019.
- [129] Y. Guo, C. Liang, Blockchain application and outlook in the banking industry, Financ. Innov. 2 (1) (2016) 1–12.
- [130] X. Zhang, Y. Han, Research on the energy distribution and tendency of blockchain in the E-commerce field, Int. J. Front. Eng. Technol. 1 (1) (2019).

- [131] J. Ke, Z. Tang, Innovation of supply chain finance model based on blockchain technology, in: International Conference on Cognitive Based Information Processing and Applications, CIPA 2021, Springer, Singapore, 2022, pp. 680–686.
- [132] M. O'Dair, R. Owen, Financing new creative enterprise through blockchain technology: Opportunities and policy implications, Strateg. Change 28 (1) (2019) 9–17.
- [133] A. Haleem, M. Javaid, R.P. Singh, R. Suman, S. Rab, Blockchain technology applications in healthcare: An overview, Int. J. Intell. Netw. 2 (2021) 130–139.
- [134] S. Alam, M. Shuaib, W.Z. Khan, S. Garg, G. Kaddoum, M.S. Hossain, Y.B. Zikria, Blockchain-based initiatives: Current state and challenges, Comput. Netw. 198 (2021) 108395.
- [135] M. Friedlmaier, A. Tumasjan, I.M. Welpe, Disrupting industries with blockchain: The industry, venture capital funding, and regional distribution of blockchain ventures, in: Venture Capital Funding, and Regional Distribution of Blockchain Ventures (September 22, 2017). Proceedings of the 51st Annual Hawaii International Conference on System Sciences, HICSS, 2018.
- [136] S.M. Idrees, M. Nowostawski, R. Jameel, A.K. Mourya, Security aspects of blockchain technology intended for industrial applications, Electronics 10 (8) (2021) 951.
- [137] D.W. Allen, C. Berg, M. Novak, Blockchain: an entangled political economy approach, J. Public Finance Public Choice 33 (2) (2018) 105–125.
- [138] P. Yu, S. Lu, E. Hanes, Y. Chen, The role of blockchain technology in harnessing the sustainability of Chinese digital finance, in: Blockchain Technologies for Sustainable Development in Smart Cities, IGI Global, 2022, pp. 155–186.
- [139] R. Nair, A. Bhagat, An application of blockchain in stock market, in: Transforming Businesses with Bitcoin Mining and Blockchain Applications, IGI Global, 2020, pp. 103–118.
- [140] S. Kurpjuweit, C.G. Schmidt, M. Klöckner, S.M. Wagner, Blockchain in additive manufacturing and its impact on supply chains, J. Bus. Logist. 42 (1) (2021) 46–70.
- [141] K. Zheng, Z. Zhang, J. Gauthier, Blockchain-based intelligent contract for factoring business in supply chains, Ann. Oper. Res. (2020) 1–21.
- [142] Y. Li, Emerging blockchain-based applications and techniques, Serv. Orient. Comput. Appl. 13 (4) (2019) 279–285.
- [143] P. Grover, A.K. Kar, M. Janssen, Diffusion of blockchain technology: Insights from academic literature and social media analytics, J. Enterp. Inf. Manag. (2019).
- [144] R.A. Memon, J. Li, J. Ahmed, A. Khan, M.I. Nazir, M.I. Mangrio, Modeling of blockchain-based systems using queuing theory simulation, in: 2018 15th International Computer Conference on Wavelet Active Media Technology and Information Processing, ICCWAMTIP, IEEE, 2018, pp. 107–111.
- [145] T. Yu, Z. Lin, Q. Tang, Blockchain: The introduction and its application in financial accounting, J. Corp. Account. Finance 29 (4) (2018) 37–47.
- [146] D. Puthal, N. Malik, S.P. Mohanty, E. Kougianos, C. Yang, The blockchain as a decentralized security framework [future directions], IEEE Consum. Electron. Mag. 7 (2) (2018) 18–21.
- [147] A. Tezel, E. Papadonikolaki, I. Yitmen, P. Hilletofth, Preparing construction supply chains for blockchain technology: An investigation of its potential and future directions, Front. Eng. Manag. 7 (4) (2020) 547–563.
- [148] E. Sujatha, J.S.J. Sundar, C.C. Deboral, D.N. Raju, P.V. Gopirajan, E-connect for agro products using supply chain with micro-finance: a blockchain approach, in: 2022 8th International Conference on Smart Structures and Systems, ICSSS, IEEE, 2022, pp. 1–7.
- [149] J. Moosavi, L.M. Naeni, A.M. Fathollahi-Fard, U. Fiore, Blockchain in supply chain management: a review, bibliometric, and network analysis, Environ. Sci. Pollut. Res. (2021) 1–15.
- [150] Y. Shibuya, V. Babich, Multi-Tier Supply Chain Financing with Blockchain, Georgetown McDonough School of Business Research Paper, 2021, (3787044).
- [151] Y. Wang, The exploration of blockchain technology in supply chain finance, Acad. J. Bus. Manag. 3 (2) (2021).
- [152] Y. El Faqir, J. Arroyo, S. Hassan, An overview of decentralized autonomous organizations on the blockchain, in: Proceedings of the 16th International Symposium on Open Collaboration, 2020, pp. 1–8.
- [153] E. Nikbakht, M. Shahrokhi, A. Corriette, Blockchain & distributed financial data, Manag. Finance (2019).
- [154] R.R. Mukkamala, R. Vatrapu, P.K. Ray, G. Sengupta, S. Halder, Blockchain for social business: principles and applications, IEEE Eng. Manag. Rev. 46 (4) (2018) 94–99.
- [155] J. Wang, P. Wu, X. Wang, W. Shou, The outlook of blockchain technology for construction engineering management, Front. Eng. Manag. (2017) 67–75.
- [156] K.M. Khan, J. Arshad, M.M. Khan, Investigating performance constraints for blockchain-based secure e-voting system, Future Gener. Comput. Syst. 105 (2020) 13–26.
- [157] S. Kumar, W.M. Lim, U. Sivarajah, J. Kaur, Artificial intelligence and blockchain integration in business: trends from a bibliometric-content analysis, Inf. Syst. Front. (2022) 1–26.
- [158] A. Klarin, The decade-long cryptocurrencies and the blockchain rollercoaster: Mapping the intellectual structure and charting future directions, Res. Int. Bus. Finance 51 (2020) 101067.

- [159] M. Sivaram, E.L. Lydia, I.V. Pustokhina, D.A. Pustokhin, M. Elhoseny, G.P. Joshi, K. Shankar, An optimal least square support vector machine-based earnings prediction of blockchain financial products, IEEE Access 8 (2020) 120321-120330.
- [160] H. Kim, U. Sengupta, M. Laskowski, Business in the front, crypto in the back: how the blockchain startup, novera, is innovating upon traditional finance, Int. Rev. Entrepreneurship 19 (1) (2021).
- [161] Z. Li, R.Y. Zhong, Z.G. Tian, H.N. Dai, A.V. Barenji, G.Q. Huang, Industrial blockchain: a state-of-the-art survey, Robot. Comput.-Integr. Manuf. 70 (2021) 102124.
- [162] R. Vinciguerra, F. Cappellieri, M. Pizzo, A disruptive innovation of blockchain in finance and banking: A jump into the future, in: Intellectual Capital, Smart Technologies and Digitalization, Springer, Cham, 2021, pp. 159–170.
- [163] M. Buitenhek, Understanding and applying blockchain technology in banking: Evolution or revolution? J. Digit. Bank. 1 (2) (2016) 111–119.
- [164] Z. Zhou, R. Li, Y. Cao, L. Zheng, H. Xiao, Dynamic performance evaluation of blockchain technologies, IEEE Access 8 (2020) 217762-217772.
- [165] T.M. Choi, Supply chain financing using blockchain: impacts on supply chains selling fashionable products, Ann. Oper. Res. (2020) 1–23.
- [166] Y. Zhang, Z. Wang, J. Deng, Z. Gong, I. Flood, Y. Wang, Framework for a blockchain-based infrastructure project financing system, IEEE Access 9 (2021) 141555-141570.
- [167] Z. Zheng, S. Xie, H.N. Dai, X. Chen, H. Wang, Blockchain challenges and opportunities: A survey, Int. J. Web Grid Serv. 14 (4) (2018) 352–375.
- [168] Y. Ke, D. Xianghui, G. Wenxin, T. Wei, Research on the financing mode and path innovation of small and medium-sized private enterprises under the background of blockchain finance, Acad. J. Bus. Manag. 3 (2) (2021).
- [169] S. Kim, H. Park, J. Lee, Word2vec-based latent semantic analysis (W2V-LSA) for topic modeling: A study on blockchain technology trend analysis, Expert Syst. Appl. 152 (2020) 113401.
- [170] C.W. Cai, Triple-entry accounting with blockchain: How far have we come? Account. Finance 61 (1) (2021) 71–93.
- [171] I. Antoniadis, K. Spinthiropoulos, S. Kontsas, Blockchain applications in tourism and tourism marketing: A short review, Strateg. Innov. Mark. Tourism (2020) 375–384.
- [172] W.A.N.G. Yu, S.H.E.N.G. Yang, X.U.E. Xing-Qun, Blockchain technology and internet financial risk prevention and control route, Stud. Sci. Sci. 40 (2) (2022) 257.
- [173] Y. Luo, Y. Tang, Comparative research on the game behavior of the participants in the traditional supply chain finance and the supply chain finance on the blockchain, Discrete Dyn. Nat. Soc. 2022 (2022).
- [174] G. Maulani, E.W. Musu, Y.J.W. Soetikno, S. Aisa, Education management using blockchain as future application innovation, IAIC Trans. Sustain. Digit. Innov. 3 (1) (2021) 60–65.
- [175] J. Kokina, R. Mancha, D. Pachamanova, Blockchain: Emergent industry adoption and implications for accounting, J. Emerg. Technol. Account. 14 (2) (2017) 91–100.
- [176] K. Ghosh, Blockchain technology: a catalyst for change in finance. Ghosh, K.(2021). Blockchain technology: a catalyst for change in finance, Drishtikon: Manag, J. 12 (1) (2021) 01–22.
- [177] Q. Rao, W. Li, Risk evaluation and forecast behavior analysis of supply chain financing based on blockchain, Wirel. Commun. Mob. Comput. 2022 (2022).
- [178] Z. Feng, H. Da, M. Xiao, C. Zhao, Z. Zuo, J. Qi, Supply chain finance in blockchain-based digital asset platform, in: The 2022 4th International Conference on Blockchain Technology, 2022, pp. 74–80.
- [179] Y. Fu, J. Zhu, Big production enterprise supply chain endogenous risk management based on blockchain, IEEE Access 7 (2019) 15310–15319.
- [180] Y. Zhang, L. Xu, Z. Gong, Y. Wang, Z. Wang, Study blockchain-based supply chain finance system of the construction industrialization, in: International Symposium on Advancement of Construction Management and Real Estate, Springer, Singapore, 2022, pp. 804–816.
- [181] N.R. Mosteanu, A. Faccia, Fintech frontiers in quantum computing, fractals, and blockchain distributed ledger: Paradigm shifts and open innovation, J. Open Innov.: Technol. Mark. Complex. 7 (1) (2021) 19.
- [182] F. Liu, H.Y. Fan, J.Y. Qi, Blockchain technology, cryptocurrency: entropy-based perspective, Entropy 24 (4) (2022) 557.
- [183] A. Faccia, M.Y.K. Al Naqbi, S.A. Lootah, Integrated cloud financial accounting cycle: how artificial intelligence, blockchain, and XBRL will change the accounting, fiscal and auditing practices, in: Proceedings of the 2019 3rd International Conference on Cloud and Big Data Computing, 2019, pp. 31–37.
- [184] K. Korpela, J. Hallikas, T. Dahlberg, Digital supply chain transformation toward blockchain integration, in: Proceedings of the 50th Hawaii International Conference on System Sciences, 2017.
- [185] D. Zhao, Application and development trend of blockchain in the financial field, in: International Conference on Application of Intelligent Systems in Multi-Modal Information Analytics, Springer, Cham, 2020, pp. 558–564.
- [186] W. Sun, Application of blockchain technology in the supply chain finance, in: 2022 7th International Conference on Cloud Computing and Big Data Analytics, ICCCBDA, IEEE, 2022, pp. 205–209.
- [187] I. Weber, M. Staples, Programmable money: next-generation blockchain-based conditional payments, Digit. Finance (2022) 1–17.

- [188] M.J. Bürer, M. de Lapparent, V. Pallotta, M. Capezzali, M. Carpita, Use cases for blockchain in the energy industry opportunities of emerging business models and related risks, Comput. Ind. Eng. 137 (2019) 106002.
- [189] X. Xu, Y. He, Blockchain application in modern logistics information sharing: A review and case study analysis, Prod. Plan. Control (2022) 1–15.
- [190] J. Zhong, H. Xie, D. Zou, D.K. Chui, A blockchain model for word-learning systems, in: 2018 5th International Conference on Behavioral, Economic, and Socio-Cultural Computing, BESC, IEEE, 2018, pp. 130–131.
- [191] A. Pal, C.K. Tiwari, A. Behl, Blockchain technology in financial services: a comprehensive review of the literature, J. Global Oper. Strateg. Sourc. (2021).
- [192] O. Alao, P. Cuffe, Towards a blockchain special purpose vehicle for financing independent renewable electricity projects in sub-Saharan Africa, in: 2020 6th IEEE International Energy Conference, ENERGYCon, IEEE, 2020, pp. 1041–1046.
- [193] Y.H. Chen, L.C. Huang, I.C. Lin, M.S. Hwang, Research on the secure financial surveillance blockchain systems, Int. J. Netw. Secur. 22 (4) (2020) 708–716.
- [194] Y. Hong, New model of food supply chain finance based on the internet of things and blockchain, Mob. Inf. Syst. 2021 (2021).
- [195] A.I. Sanka, M. Irfan, I. Huang, R.C. Cheung, A survey of breakthrough in blockchain technology: Adoptions, applications, challenges and future research, Comput. Commun. 169 (2021) 179–201.
- [196] G. Erefae, M. Nuseir, Blockchain in global finance make-over: Exploring the mediating role of supply chain flexibility, Uncertain Supply Chain Manag. 10 (3) (2022) 983–992.
- [197] S. Ahluwalia, R.V. Mahto, M. Guerrero, Blockchain technology and startup financing: A transaction cost economics perspective, Technol. Forecast. Soc. Change 151 (2020) 119854.
- [198] T.M. Hughes, The global financial services industry and the blockchain, J. Struct. Finance 23 (4) (2018) 36–40.
- [199] L. Yuyan, W. Lan, The role of blockchain technology on supply chain finance, in: Proceedings of the 3rd International Conference on Data Science and Information Technology, 2020, pp. 153–158.
- [200] P.K. Sharma, N. Kumar, J.H. Park, Blockchain technology toward green IoT: Opportunities and challenges, IEEE Netw. 34 (4) (2020) 263–269.
- [201] K.J. Smith, G. Dhillon, Assessing blockchain potential for improving the cybersecurity of financial transactions, Manag. Finance (2019).
- [202] K. Padia, V. Shah, D. Desai, P. Kubde, Trade finance using blockchain, in: Proceedings of the 3rd International Conference on Advances in Science & Technology, ICAST, 2020.
- [203] T. Ahram, A. Sargolzaei, S. Sargolzaei, J. Daniels, B. Amaba, Blockchain technology innovations, in: 2017 IEEE Technology & Engineering Management Conference, TEMSCON, IEEE, 2017, pp. 137–141.
- [204] W. Baiod, J. Light, A. Mahanti, Blockchain technology and its applications across multiple domains: a survey, J. Int. Technol. Inf. Manag. 29 (4) (2021) 78–119.
- [205] N. Salisbury, J. Khvatsky, Using smart algorithms, machine learning, and blockchain technology to streamline and accelerate deal flow in climate finance, in: Transforming Climate Finance and Green Investment with Blockchains, Academic Press, 2018, pp. 179–188.
- [206] V. Janaviciene, V.V. Fomin, Systematic literature mapping on blockchain application in the context of economics, finance and management, in: The 11th International Scientific Conference New Challenges of Economic and Business Development-2019: Incentives for Sustainable Economic Growth, 2019, pp. 310–319.
- [207] S. Zhang, J.A. Tang, Q. Hu, F. Liu, Research on blockchain financial derivatives cluster, in: 2020 the 3rd International Conference on Blockchain Technology and Applications, 2020, pp. 40–46.
- [208] J.M. Woodside, F.K. Augustine Jr., W. Giberson, Blockchain technology adoption status and strategies, J. Int. Technol. Inf. Manag. 26 (2) (2017) 65–93.
- [209] L. Su, H. Wang, Supply chain finance research in digital bulk commodities service platform based on blockchain, in: 2020 International Conference on E-Commerce and Internet Technology, ECIT, IEEE, 2020, pp. 341–344.
- [210] D. Boughaci, A.A. Alkhawaldeh, Enhancing the security of financial transactions in Blockchain by using machine learning techniques: Towards a sophisticated security tool for banking and finance, in: 2020 First International Conference of Smart Systems and Emerging Technologies, SMARTTECH, IEEE, 2020, pp. 110–115.
- [211] W. Gao, P.K. Chen, Analysis of barriers to implement blockchain in supply chain finance, in: 2021 3rd International Conference on E-Business and E-Commerce Engineering, 2021, pp. 46–52.
- [212] A. Batwa, A. Norrman, A framework for exploring blockchain technology in supply chain management, Oper. Supply Chain Manag.: Int. J. 13 (3) (2020) 294–306.
- [213] W. Zhang, The current situation and trend of blockchain technology in the financial field, Mob. Inf. Syst. 2022 (2022).
- [214] S. Dos Santos, J. Singh, R.K. Thulasiram, S. Kamali, L. Sirico, L. Loud, A new era of blockchain-powered decentralized finance (DeFi)-a review, in: 2022 IEEE 46th Annual Computers, Software, and Applications Conference, COMPSAC, IEEE, 2022, pp. 1286–1292.
- [215] A. Gurtu, J. Johny, Potential of blockchain technology in supply chain management: a literature review, Int. J. Phys. Distrib. Logist. Manage. (2019).

- [216] M. Swan, Anticipating the economic benefits of blockchain, Technol. Innov. Manag. Rev. 7 (10) (2017) 6–13.
- [217] P.N. Petratos, N. Ljepava, A. Salman, Blockchain technology, sustainability and business: A literature review and the case of Dubai and UAE, Sustain. Dev. Soc. Responsib. 1 (2020) 87–93.
- [218] A. Hooper, D. Holtbrügge, Blockchain technology in international business: changing the agenda for global governance, Rev. Int. Bus. Strateg. 30 (2) (2020) 183–200.
- [219] A. Nasir, K. Shaukat, K.I. Khan, I.A. Hameed, T.M. Alam, S. Luo, What is core and what future holds for blockchain technologies and cryptocurrencies: A bibliometric analysis, IEEE Access 9 (2020) 989–1004.
- [220] J. Woo, R. Fatima, C.J. Kibert, R.E. Newman, Y. Tian, R.S. Srinivasan, Applying blockchain technology for building energy performance measurement, reporting, and verification (MRV) and the carbon credit market: A review of the literature, Build. Environ. 205 (2021) 108199.
- [221] J. Chiu, T.V. Koeppl, Blockchain-based settlement for asset trading, Rev. Financ. Stud. 32 (5) (2019) 1716–1753.
- [222] H. Hassani, X. Huang, E. Silva, Banking with blockchain-ed big data, J. Manag. Anal. 5 (4) (2018) 256–275.
- [223] C. Zhu, Z. Fu, Research on the development trend of traditional financial industry based on blockchain technology, Asian J. Res. Comput. Sci. 5 (3) (2020) 1–6.
- [224] A. Sheth, H. Subramanian, Blockchain and contract theory: modeling smart contracts using insurance markets, Manag. Finance (2019).
- [225] Z. Fauziah, H. Latifah, X. Omar, A. Khoirunisa, S. Millah, Application of blockchain technology in smart contracts: a systematic literature review, Aptisi Trans. Technopreneurship 2 (2) (2020) 160–166.
- [226] R. Tsepeleva, V. Korkhov, Implementation of the cross-blockchain interacting protocol, in: International Conference on Computational Science and its Applications, Springer, Cham, 2021, pp. 42–55.
- [227] Y. Xue, X. Liang, D. Zhao, A blockchain-based rice supply chain system, in: MATEC Web of Conferences, Vol. 336, EDP Sciences, 2021, p. 09003.
- [228] Y. Xinyi, Z. Yi, Y. He, Technical characteristics and model of blockchain, in: 2018 10th International Conference on Communication Software and Networks, ICCSN, IEEE, 2018, pp. 562–566.
- [229] U. Hacioglu, Digital Business Strategies in Blockchain Ecosystems, Springer International Publishing, 2020, DOI, 10, 978-3.
- [230] M. Zachariadis, G. Hileman, S.V. Scott, Governance and control in distributed ledgers: Understanding the challenges facing blockchain technology in financial services, Inf. Org. 29 (2) (2019) 105–117.
- [231] D. Puthal, N. Malik, S.P. Mohanty, E. Kougianos, G. Das, Everything you wanted to know about the blockchain: Its promise, components, processes, and problems, IEEE Consum. Electron. Mag. 7 (4) (2018) 6–14.
- [232] W. Viriyasitavat, L. Da Xu, Z. Bi, V. Pungpapong, Blockchain and internet of things for modern business process in digital economy—the state of the art, IEEE Trans. Comput. Soc. Syst. 6 (6) (2019) 1420–1432.
- [233] J. Frizzo-Barker, P.A. Chow-White, P.R. Adams, J. Mentanko, D. Ha, S. Green, Blockchain as a disruptive technology for business: A systematic review, Int. J. Inf. Manage. 51 (2020) 102029.
- [234] D. Khan, L.T. Jung, M.A. Hashmani, Systematic literature review of challenges in blockchain scalability, Appl. Sci. 11 (20) (2021) 9372.
- [235] F. Alkurdi, I. Elgendi, K.S. Munasinghe, D. Sharma, A. Jamalipour, Blockchain in IoT Security: A Survey. 2018 28th International Telecommunication Networks and Applications Conference, ITNAC, IEEE, 2018, pp. 1–4.
- [236] C. Zhao, X. Meng, Research on innovation and development of blockchain technology in financial field, in: 2019 International Conference on Pedagogy, Communication and Sociology, ICPCS 2019, Atlantis Press, 2019, pp. 421–424.
- [237] A. Muneeza, N.A. Arshad, A.T. Arifin, The application of blockchain technology in crowdfunding: towards financial inclusion via technology, Int. J. Manag. Appl. Res. 5 (2) (2018) 82–98.
- [238] S. Nanayakkara, S. Perera, S. Senaratne, G.T. Weerasuriya, H.M.N.D. Bandara, Blockchain and smart contracts: A solution for payment issues in construction supply chains, in: Informatics, Vol. 8, No. 2, Multidisciplinary Digital Publishing Institute, 2021, p. 36.
- [239] I. Cioroianu, S. Corbet, C. Larkin, The differential impact of corporate blockchain-development as conditioned by sentiment and financial desperation, J. Corp. Finance 66 (2021) 101814.
- [240] Y. Jiang, K. Zhou, X. Lu, S. Yang, Electricity trading pricing among prosumers with game theory-based model in energy blockchain environment, Appl. Energy 271 (2020) 115239.
- [241] L. Mishra, V. Kaushik, Application of blockchain in dealing with sustainability issues and challenges of financial sector, J. Sustain. Finance Invest. (2021) 1–16.
- [242] U. Ravindran, P.R. Vamsi, A secure blockchain based finance application, in: 2021 Thirteenth International Conference on Contemporary Computing, IC3-2021, 2021, pp. 99–109.
- [243] K.A. Schulz, O.J. Gstrein, A.J. Zwitter, Exploring the governance and implementation of sustainable development initiatives through blockchain technology, Futures 122 (2020) 102611.
- [244] D. Levis, F. Fontana, E. Ughetto, A look into the future of blockchain technology, PLoS One 16 (11) (2021) e0258995.

- [245] S. Wang, X. Qu, Blockchain applications in shipping, transportation, logistics, and supply chain, in: Smart Transportation Systems 2019, Springer, Singapore, 2019, pp. 225–231.
- [246] E. Hofmann, P. Wetzel, D. Röck, Blockchain can solve the pain points of trade and supply chain finance, in: B2B Fintech: Payments, Supply Chain Finance and E-Invoicing Market Guide, Vol. 2018, 2018, pp. 50–51.
- [247] M. Krichen, M. Ammi, A. Mihoub, M. Almutiq, Blockchain for modern applications: A survey, Sensors 22 (14) (2022) 5274.
- [248] Y. Tribis, A. El Bouchti, H. Bouayad, Supply chain management based on blockchain: A systematic mapping study, in: MATEC Web of Conferences, Vol. 200, EDP Sciences, 2018, p. 00020.
- [249] Y. Issaoui, A. Khiat, A. Bahnasse, H. Ouajji, Smart logistics: Study of the application of blockchain technology, Procedia Comput. Sci. 160 (2019) 266–271.
- [250] D. Yermack, A. Fingerhut, Blockchain technology's potential in the financial system, in: Proceedings of the 2019 Financial Market's Conference, 2019.
- [251] D.E. O'Leary, Configuring blockchain architectures for transaction information in blockchain consortiums: The case of accounting and supply chain systems, Intell. Syst. Account. Finance Manag. 24 (4) (2017) 138–147.
- [252] M. Javaid, A. Haleem, R.P. Singh, S. Khan, R. Suman, Blockchain technology applications for Industry 4.0: A literature-based review, Blockchain: Res. Appl. (2021) 100027.
- [253] T. Chen, D. Wang, Combined application of blockchain technology in fractional calculus model of supply chain financial system, Chaos Solitons Fractals 131 (2020) 109461.
- [254] C. Dang, F. Wang, Z. Yang, H. Zhang, Y. Qian, Evaluating and forecasting the risks of small to medium-sized enterprises in the supply chain finance market using blockchain technology and deep learning model, Oper. Manag. Res. (2022) 1–14.
- [255] Q.K. Nguyen, Blockchain-a financial technology for future sustainable development, in: 2016 3rd International Conference on Green Technology and Sustainable Development, GTSD, IEEE, 2016, pp. 51–54.
- [256] N. Ariffin, A.Z. Ismail, The design and implementation of trade finance application based on hyperledger fabric permissioned blockchain platform, in: 2019 International Seminar on Research of Information Technology and Intelligent Systems, ISRITI, IEEE, 2019, pp. 488–493.
- [257] C.M.A. Natalia Maslova, P. CTP, Blockchain: disruption and opportunity, Strateg. Finance 100 (1) (2018) 24–29.
- [258] M. Xu, X. Chen, G. Kou, A systematic review of blockchain, Financ. Innov. 5 (1) (2019) 1–14.
- [259] C. Jiang, C. Ru, Application of blockchain technology in supply chain finance, in: 2020 5th International Conference on Mechanical, Control and Computer Engineering, ICMCCE, IEEE, 2020, pp. 1342–1345.
- [260] I. Ioannou, G. Demirel, Blockchain and supply chain finance: a critical literature review at the intersection of operations, finance and law, J. Bank. Financ. Technol. (2022) 1–25.
- [261] J. Chen, S. Chen, Q. Liu, M.I. Shen, Applying blockchain technology to reshape the service models of supply chain finance for SMEs in China, Singapore Econ. Rev. (2021) 1–18.
- [262] Y.X. Huang, Z.H. Liang, Y.H. Wang, G. Lu, Research on the application of blockchain in supply chain finance, Comput. Sci. Appl. 8 (1) (2018) 78–88.
- [263] E. Hofmann, U.M. Strewe, N. Bosia, Introduction—why to pay attention on blockchain-driven supply chain finance? in: Supply Chain Finance and Blockchain Technology, Springer, Cham, 2018, pp. 1–6.
- [264] A.A. Baev, V.S. Levina, A.V. Reut, A.A. Svidler, I.A. Kharitonov, V.V. Grigor'ev, Blockchain technology in accounting and auditing, Account. Anal. Audit. 7 (1) (2020) 69–79.
- [265] Z. Liu, Literature review of supply chain finance based on blockchain perspective, Open J. Bus. Manag. 9 (1) (2020) 419–429.
- [266] C.W. Cai, Disruption of financial intermediation by FinTech: a review on crowdfunding and blockchain, Account. Finance 58 (4) (2018) 965–992.
- [267] W. Pan, M. Qiu, Application of blockchain in asset-backed securitization, in: 2020 IEEE 6th Intl Conference on Big Data Security on Cloud (BigDataSecurity), IEEE Intl Conference on High Performance and Smart Computing, (HPSC) and IEEE Intl Conference on Intelligent Data and Security, IDS, IEEE, 2020, pp. 71–76.
- [268] L.R. Cohen, L. Samuelson, H. Katz, How securitization can benefit from blockchain technology, J. Struct. Finance 23 (2) (2017) 51–54.

- [269] I. Karagiannis, K. Mavrogiannis, J. Soldatos, D. Drakoulis, E. Troiano, A. Polyviou, Blockchain based sharing of security information for critical infrastructures of the finance sector, in: Computer Security, Springer, Cham, 2019, pp. 226–241.
- [270] Y.J. An, P.M.S. Choi, S.H. Huang, Blockchain, cryptocurrency, and artificial intelligence in finance, in: Fintech with Artificial Intelligence, Big Data, and Blockchain, Springer, Singapore, 2021, pp. 1–34.
- [271] S. Halilbegovic, A. Arapovic, N. Celebic, T. Atovic, Exploratory analysis of blockchain application in trade finance, Eur. J. Econ. Stud. 2 (8) (2019) 110–119.
- [272] V. Babich, G. Hilary, Blockchain and other distributed ledger technologies in operations, Found. Trends Technol. Inf. Oper. Manag. 12 (2-3) (2019) 152–172.
- [273] J.H. Lee, M. Pilkington, How the blockchain revolution will reshape the consumer electronics industry [future directions], IEEE Consum. Electron. Mag. 6 (3) (2017) 19–23.
- [274] Z. Wang, Analysis of risks and regulatory issues in the development of blockchain finance, Acad. J. Bus. Manag. 4 (1) (2022) 61–66.
- [275] A. Collomb, K. Sok, Blockchain/distributed ledger technology (DLT): What impact on the financial sector? Digiworld Econ. J. (103) (2016).
- [276] P.R.D. Cunha, P. Soja, M. Themistocleous, Blockchain for development: a guiding framework, Inf. Technol. Dev. 27 (3) (2021) 417–438.
- [277] D. Li, D. Han, H. Liu, Fabric-chain & chain: a blockchain-based electronic document system for supply chain finance, in: International Conference on Blockchain and Trustworthy Systems, Springer, Singapore, 2020, pp. 601–608.
- [278] M. Chanson, N. Martens, F. Wortmann, The role of user-generated content in blockchain-based decentralized finance, in: ECIS 2020 Proceedings, 2020.
- [279] M. Hamilton, Blockchain distributed ledger technology: An introduction and focus on smart contracts, J. Corp. Account. Finance 31 (2) (2020) 7–12.
- [280] A. Gump, C. Leonard, Blockchain:: regulating the future of finance, Int. Financ. Law Rev. 35 (13) (2016) 1.
- [281] D. Gunasekera, E. Valenzuela, Adoption of blockchain technology in the australian grains trade: An assessment of potential economic effects, Econom. Pap.: J. Appl. Econ. Policy 39 (2) (2020) 152–161.
- [282] A. Pal, C.K. Tiwari, N. Haldar, Blockchain for business management: Applications, challenges and potentials, J. High Technol. Manag. Res. 32 (2) (2021) 100414.
- [283] P. Martino, K.J. Wang, C. Bellavitis, C.M. DaSilva, An introduction to blockchain, cryptocurrency and initial coin offerings, in: New Frontiers in Entrepreneurial Finance Research, 2020, pp. 181–206.
- [284] P. Dutta, T.M. Choi, S. Somani, R. Butala, Blockchain technology in supply chain operations: Applications, challenges and research opportunities, Transp. Res. E 142 (2020) 102067.
- [285] J. Kim, S. Kim, A survey of decentralized finance (DeFi) based on blockchain, J. Korea Soc. Comput. Inf. 26 (3) (2021) 59–67.
- [286] J. Fu, B. Cao, X. Wang, P. Zeng, W. Liang, Y. Liu, BFS: A blockchain-based financing scheme for logistics company in supply chain finance, Connect. Sci. 34 (1) (2022) 1929–1955.
- [287] P. Meena, M. Pal, P. Jain, R. Pamula, 6G communication networks: introduction, vision, challenges, and future directions, Wirel. Pers. Commun. 125 (2) (2022) 1097–1123, http://dx.doi.org/10.1007/s11277-022-09590-5.
- [288] A. Yadav, A. Prakash, Factors influencing sustainable development integration in management education: An Empirical Assessment of management education institutions in India, Int. J. Manag. Ed. 20 (1) (2022) 100604.
- [289] S. Khan, M. Kaushik, R. Kumar, W. Khan, Investigating the barriers of blockchain technology integrated food supply chain: a BWM approach, Benchmarking: Int. J. (2022) http://dx.doi.org/10.1108/bij-08-2021-0489.
- [290] S. Khan, R. Singh, Kirti, Critical factors for blockchain technology implementation: a supply chain perspective, J. Ind. Integr. Manag. (2021) 2150011, http://dx.doi.org/10.1142/s2424862221500111.
- [291] A. Rizwan, D. Karras, J. Kumar, M. Sánchez-Chero, M. Mogollón Taboada, G. Altamirano, An internet of things (IoT) based block chain technology to enhance the quality of supply chain management (SCM), Math. Probl. Eng. 2022 (2022) 1–12, http://dx.doi.org/10.1155/2022/9679050.
- [292] H. Sharma, A. Chaturvedi, Adoption of smart technologies: an Indian perspective, in: 2021 5th International Conference on Information Systems and Computer Networks, ISCON, 2021, http://dx.doi.org/10.1109/iscon52037. 2021.9702390.
- [293] R. Singh, S. Khan, J. Dsilva, A framework for assessment of critical factor for circular economy practice implementation, J. Modell. Manag. (2022) http: //dx.doi.org/10.1108/jm2-06-2021-0145.

TBench Editorial Board

Co-EIC

Prof. Dr. Jianfeng Zhan, ICT, Chinese Academy of Sciences and BenchCouncil Prof. Dr. Tony Hey, Rutherford Appleton Laboratory STFC, UK

Editorial office

Dr. Wanling Gao, ICT, Chinese Academy of Sciences and BenchCouncil Shaopeng Dai, ICT, Chinese Academy of Sciences and BenchCouncil Dr. Chunjie Luo, University of Chinese Academy of Sciences, China

Advisory Board

Prof. Jack Dongarra, University of Tennessee, USA Prof. Geoffrey Fox, Indiana University, USA Prof. D. K. Panda, The Ohio State University, USA

Founding Editor

Prof. H. Peter Hofstee, IBM Systems, USA and Delft University of Technology, Netherlands Dr. Zhen Jia, Amazon, USA Prof. Blesson Varghese, Queen's University Belfast, UK Prof. Raghu Nambiar, AMD, USA Prof. Jidong Zhai, Tsinghua University, China Prof. Francisco Vilar Brasileiro, Federal University of Campina Grande, Brazil Prof. Jianwu Wang, University of Maryland, USA Prof. David Kaeli, Northeastern University, USA Prof. Bingshen He, National University of Singapore, Singapore Dr. Lei Wang, Institute of Computing Technology, Chinese Academy of Sciences, China Prof. Weining Qian, East China Normal University, China Dr. Arne J. Berre, SINTEF, Norway Prof. Ryan Eric Grant, Sandia National Laboratories, USA Prof. Rong Zhang, East China Normal University, China Prof. Cheol-Ho Hong, Chung-Ang University, Korea Prof. Vladimir Getov, University of Westminster, UK Prof. Zhifei Zhang, Capital Medical University Prof. K. Selcuk Candan, Arizona State University, USA Dr. Yunyou Huang, Guangxi Normal University Prof. Woongki Baek, Ulsan National Institute of Science and Technology, Korea Prof. Radu Teodorescu, The Ohio State University, USA Prof. John Murphy, University College Dublin, Ireland Prof. Marco Vieira, The University of Coimbra (UC), Portugal Prof. Jose Merseguer, University of Zaragoza (UZ), Spain Prof. Xiaoyi Lu, University of California, USA Prof. Yanwu Yang, Huazhong University of Science and Technology, China Prof. Jungang Xu, University of Chinese Academy of Sciences, China Prof. Jiaquan Gao, Professor, Nanjing Normal University, China

Associate Editor

Dr. Chen Zheng, Institute of Software, Chinese Academy of Sciences, China Dr. Biwei Xie, Institute of Computing Technology, Chinese Academy of Sciences, China Dr. Mai Zheng, Iowa State University, USA Dr. Wenyao Zhang, Beijing Institute of Technology, China

Dr. Bin Liao, North China Electric Power University, China

More information about this series at https://www.benchcouncil.org/tbench/

TBench Call For Papers

BenchCouncil Transactions on Benchmarks, Standards and Evaluations (TBench) ISSN:2772-4859

Aims and Scopes

BenchCouncil Transactions on Benchmarks, Standards, and Evaluations (TBench) publishes position articles that open new research areas, research articles that address new problems, methodologies, tools, survey articles that build up comprehensive knowledge, and comments articles that argue the published articles. The submissions should deal with the benchmarks, standards, and evaluation research areas. Particular areas of interest include, but are not limited to:

• 1. Generalized benchmark science and engineering (see

https://www.sciencedirect.com/science/article/pii/S2772485921000120), including but not limited to

- measurement standards
- standardized data sets with defined properties
- representative workloads
- ➢ representative data sets
- ➢ best practices
- 2. Benchmark and standard specifications, implementations, and validations of:
 - Big Data
 - ≻ AI
 - ➢ HPC
 - ➢ Machine learning
 - Big scientific data
 - ➢ Datacenter
 - ➤ Cloud
 - Warehouse-scale computing
 - Mobile robotics
 - Edge and fog computing
 - ≻ IoT
 - Chain block
 - Data management and storage
 - Financial domains
 - Education domains
 - Medical domains
 - Other application domains
- 3. Data sets
 - Detailed descriptions of research or industry datasets, including the methods used to collect the data and technical analyses supporting the quality of the measurements.
 - Analyses or meta-analyses of existing data and original articles on systems, technologies, and techniques that advance data sharing and reuse to support reproducible research.
 - Evaluating the rigor and quality of the experiments used to generate the data and the completeness of the data description.
 - > Tools generating large-scale data while preserving their original characteristics.
- 4. Workload characterization, quantitative measurement, design, and evaluation studies of:
 - > Computer and communication networks, protocols, and algorithms
 - ▶ Wireless, mobile, ad-hoc and sensor networks, IoT applications
 - Computer architectures, hardware accelerators, multi-core processors, memory systems, and storage networks
 - High-Performance Computing
 - > Operating systems, file systems, and databases

- > Virtualization, data centers, distributed and cloud computing, fog, and edge computing
- Mobile and personal computing systems
- Energy-efficient computing systems
- Real-time and fault-tolerant systems
- Security and privacy of computing and networked systems
- > Software systems and services, and enterprise applications
- > Social networks, multimedia systems, Web services
- Cyber-physical systems, including the smart grid
- 5. Methodologies, metrics, abstractions, algorithms, and tools for:
 - Analytical modeling techniques and model validation
 - Workload characterization and benchmarking
 - > Performance, scalability, power, and reliability analysis
 - Sustainability analysis and power management
 - > System measurement, performance monitoring, and forecasting
 - > Anomaly detection, problem diagnosis, and troubleshooting
 - > Capacity planning, resource allocation, run time management, and scheduling
 - > Experimental design, statistical analysis, simulation
- 6. Measurement and evaluation
 - Evaluation methodology and metric
 - Testbed methodologies and systems
 - > Instrumentation, sampling, tracing, and profiling of Large-scale real-world applications and systems
 - > Collection and analysis of measurement data that yield new insights
 - Measurement-based modeling (e.g., workloads, scaling behavior, assessment of performance bottlenecks)
 - > Methods and tools to monitor and visualize measurement and evaluation data
 - Systems and algorithms that build on measurement-based findings
 - Advances in data collection, analysis, and storage (e.g., anonymization, querying, sharing)
 - ➢ Reappraisal of previous empirical measurements and measurement-based conclusions
 - > Descriptions of challenges and future directions the measurement and evaluation community should pursue

Bench 2022 Call For Papers

2022 BenchCouncil International Symposium on Benchmarking, Measuring and Optimizing (Bench 2022) Calls For Papers

https://www.benchcouncil.org/bench22/index.html

Full Papers deadline: July 28, 2022, 23:59:59 AoE Notification: September 6, 2022, 23:59:59 AoE Final Papers due: October 11, 2022, 23:59:59 AoE Conference date: Nov. 7th - Nov. 9th, 2022 (Virtual) Submission site: https://bench2022.hotcrp.com/

Introduction

Benchmarks, Data, Standards, Measurements, and Optimizations are fundamental human activities and assets. The Bench conference has two essential duties: promote data or benchmark-based quantitative approaches to tackle multidisciplinary and interdisciplinary challenges; connect architecture, system, data management, algorithm, and application communities to better co-design for the inherent workload characterizations.

The Bench conference provides a high-quality, single-track forum for presenting results and discussing ideas that further the knowledge and understanding of the benchmarks, data, standards, measurements, and optimizations community as a whole. It is a multidisciplinary and interdisciplinary conference. The past meetings attracted researchers and practitioners from the architecture, system, algorithm, and application communities. It includes both invited sessions and contributed sessions.

Regularly, the Bench conference will present the BenchCouncil Achievement Award (\$3000), the BenchCouncil Rising Star Award (\$1000), the BenchCouncil Best Paper Award (\$1000), and the BenchCouncil Distinguished Doctoral Dissertation Awards in Computer Architecture (\$1000) and in other areas (\$1000). This year, the BenchCouncil Distinguished Doctoral Dissertation Award includes two tracks: computer architecture and other areas. Among the submissions of each track, four candidates will be selected as finalists. They will be invited to give a 30-minute presentation at the Bench' 22 Conference and contribute research articles to BenchCouncil Transactions on Benchmarks, Standards and Evaluation. Finally, for each track, one among the four will receive the award for each track, which carries a \$1,000 honorarium.

Organization

General Co-Chairs Emmanuel Jeannot, INRIA, France Peter Mattson, Google, USA Wanling Gao, University of Chinese Academy of Sciences, China

Program Co-Chairs

Chunjie Luo, ICT, Chinese Academy of Sciences, China Ce Zhang, ETH Zurich, Switzerland Ana Gainaru, Oak Ridge National Laboratory, USA

Publicity Co-Chairs

David Kanter, MLCommons Rui Ren, Beijing Institute of Open Source Chip Zhen Jia, Amazon

Web Co-Chairs Jiahui Dai, BenchCouncil Jiahui Dai, BenchCouncil Qian He, Beijing Institute of Open Source Chip

Award Committees

BenchCouncil Distinguished Doctoral Dissertation Award Committee in Other Areas: Jack Dongarra, University of Tennessee Xiaoyi Lu, The University of California, Merced Jeyan Thiyagalingam, STFC-RAL Lei Wang, ICT, Chinese Academy of Sciences Spyros Blanas, The Ohio State University

BenchCouncil Distinguished Doctoral Dissertation Award Committee in Computer Architecture: Peter Mattson, Google Vijay Janapa Reddi, Harvard University Wanling Gao, Chinese Academy of Sciences

Bench Steering Committees

Jack Dongarra, University of Tennessee Geoffrey Fox, Indiana University D. K. Panda, The Ohio State University Felix, Wolf, TU Darmstadt Xiaoyi Lu, University of California, Merced Resit Sendag, University of Rhode Island, USA Wanling Gao, ICT, Chinese Academy of Sciences & UCAS Jianfeng Zhan, ICT, Chinese Academy of Sciences &BenchCouncil

Call For Papers

The Bench conference encompasses a wide range of areas and topics in benchmarking, measurement, evaluation methods and tools. We solicit papers describing original and previously unpublished work. The areas and topics of interest include, but are not limited to the following.

- 1. Areas:
 - > Architecture
 - Data Management
 - ➢ Algorithm
 - Datasets
 - ➢ System
 - ➢ Network
 - Reliability and Security
 - Application
- 2. Topics:
 - > Benchmark and standard specifications, implementations, and validations
 - Dataset Generation and Analysis
 - > Workload characterization, quantitative measurement, design and evaluation studies
 - Methodologies, abstractions, metrics, algorithms and tools
 - Measurement and evaluation

Paper Submission

Papers must be submitted in PDF. For a full paper, the page limit is 15 pages in the LNCS format, not including references. For a short paper, the page limit is 8 pages in the LNCS format, not including references. The submissions will be judged based on the merit of the ideas rather than the length. The reviewing process is double-blind. Upon acceptance, the proceeding will be published by Springer LNCS (Indexed by EI). Please note that the LNCS format is the final one for publishing. Distinguished papers will be recommended to and published by the BenchCouncil Transactions on Benchmarks, Standards and Evaluation (TBench).

At least one author must pre-register for the symposium, and at least one author must attend the symposium to present the paper. Papers for which no author is pre-registered will be removed from the proceedings.

Submission site: https://bench2022.hotcrp.com/

LNCS Latex template: https://www.benchcouncil.org/file/llncs2e.zip

Awards

* BenchCouncil Achievement Award (\$3,000)

- This award recognizes a senior member who has made long-term contributions to benchmarking, measuring, and optimizing. The winner is eligible for the status of a BenchCouncil Fellow.

* BenchCouncil Rising Star Award (\$1,000)

- This award recognizes a junior member who demonstrates outstanding potential for research and practice in benchmarking, measuring, and optimizing.

* BenchCouncil Best Paper Award (\$1,000)

- This award recognizes a paper presented at the Bench conferences, which demonstrates potential impact on research and practice in benchmarking, measuring, and optimizing.

* BenchCouncil Distinguished Doctoral Dissertation Award (\$2000)

- This award recognizes and encourages superior research and writing by doctoral candidates in the broad field of benchmarks, data, standards, evaluations, and optimizations community. This year, the award includes two tracks, including the BenchCouncil Distinguished Doctoral Dissertation Award in Computer Architecture (\$1000) and BenchCouncil Distinguished Doctoral Dissertation Award in other areas (\$1000).

Technical Program Committee

Murali Krishna Emani, ANL Shin-ying Lee, AMD Steve Farrell, NERSC Krishnakumar Nair, Meta Greg Diamos, Landing.AI Fei Sun, Alibaba Narayanan Sundaram, Facebook Zhen Jia, Amazon Shengen Yan, SenseTime Gang Lu, Tencent Rui Ren, Beijing Open-Source IC Academy Bin Hu, ICT, CAS Khaled Ibrahim, Lawrence Berkeley National Laboratory Sascha Hunold, TU Wien Woongki Baek, UNIST Mario Marino, Leeds Beckett University Bin Ren, William & Mary

Gwangsun Kim, POSTECH Vladimir Getov, University of Westminster Guangli Li, ICT, CAS Biwei Xie, ICT, CAS Nicolas Rougier, INRIA