Fifty Years of the International Symposium on Computer Architecture: A Data-Driven Retrospective

Matthew D. Sinclair, *University of Wisconsin–Madison, Madison, WI, 53706, USA*
Parthasarathy Ranganathan and Gaurang Upasani, *Google, Mountain View, CA, 94043, USA*
Adrian Sampson, *Cornell University, Ithaca, NY, 14850, USA*
David Patterson, *University of California, Berkeley and Google, Berkeley, CA, 94720, USA*
Rutwik Jain, *University of Wisconsin–Madison, Madison, WI, 53706, USA*
Nidhi Parthasarathy, *Lynbrook High School, San Jose, CA, 95129, USA*
Shaan Shah, *University of California, San Diego, San Diego, CA, 92093, USA*

2023 marked the fiftieth year of the International Symposium on Computer Architecture (ISCA). As one of the oldest and preeminent computer architecture conferences, ISCA represents a microcosm of the broader community; correspondingly, a 50-year-retrospective offers us a great way to track the impact and evolution of the field. Analyzing the content and impact of all the papers published at ISCA so far, we show how computer architecture research has been at the forefront of advances that have driven the broader computing ecosystem. Decadal trends show a dynamic and rapidly-evolving field, with diverse contributions. Examining how the most highly-cited papers achieve their popularity reveals interesting trends on technology adoption curves and the path to impact. Our data also highlights a growing and thriving community, with interesting insights on diversity and scale. We conclude with a summary of the celebratory panel held at ISCA, with observations on the exciting future ahead.

The International Symposium on Computer Architecture (ISCA) is one of the oldest, longest-running, and most prestigious venues for publishing computer architecture research. Since 1973, ISCA has been organized annually, except for 1975. This year, therefore, marked ISCA's 50th year. These past five decades have seen a total of 2134 papers published at ISCA across a variety of topics pertaining to computer architecture. Ranging from the smallest handhelds and edge devices to the largest data centers, across hardware and software, and addressing varied topics like performance, programmability, power, reliability, security, and sustainability, these advances at ISCA represent a nice microcosm of the broader computer architecture field. Accordingly, as we look back at, and celebrate, the past 50 years of ISCA, we also have a unique opportunity to reflect on the impact and evolution of the area and the community as a whole.

In this article, we provide such a retrospective. Building on a multitude of data sources, from digital archived papers in Association for Computing Machinery (ACM's) Digital Library to open-service bibliographic databases like the DBLP computer science bibliography and Google Scholar and a collection of rich data analysis and visualization tools, we take a *data-driven approach* to draw interesting insights about the past five decades. Using automated scripts, we identify the 50 top-cited papers of all time as well as the highest-cited papers for each year and categorize them across key computer architecture topics. We also study keyword frequencies ("word clouds") across paper abstracts. Our analysis

0272-1732 © 2023 IEEE
Digital Object Identifier 10.1109/MM.2023.3324465
Date of current version 2 November 2023.

highlights key trends in the evolution of computing systems—multiprocessors, very long instruction word (VLIW) and superscalar processors, simultaneous multithreading (SMT), multicore architectures, and, most recently, custom silicon accelerators—and the computer architecture research underpinning these transformational advances. Analyzing the popularity trends over time of individual highly cited papers highlights different paths to technology adoption and impact, as well as the surprising longevity of foundational ideas.

Beyond technical trends, we also track the growth of the computer architecture *community*. Our data show a growing community: a 6× increase in the number of submissions and a 10× increase in program committee (PC) sizes, with increased collaboration (an increasing number of authors per paper) and globalization (a healthy geographical representation). However, although our data on gender diversity shows good progress (15% of authors identified as female last year versus 3% in 1973), there are opportunities to do more, similar to other engineering disciplines.[9,10]

Furthermore, from a more qualitative perspective, we also surveyed the community and present key takeaways from the responses as well as summarize the discussions from a lively celebratory panel held at ISCA. These reinforce the same conclusions as our data-driven study: *ISCA, and more broadly, the computer architecture community, is healthy, vibrant, and impactful, and there is significant excitement for the future.*

COMPUTER ARCHITECTURE OVER FIVE DECADES

Analysis Methodology

To understand how ISCA's community and mechanics have evolved over time, we collected data from multiple sources. First, we collected digital copies of all papers from the ACM Digital Library. For each paper, we collected paper title, author names, affiliation(s), and citation count (including citation count per year). We also used the conference forewords to determine PC and External Review Committee size, paper submissions and acceptances, and conference location. To collect this data, we used a variety of sources: DBLP, Google Scholar, and the ISCA Hall of Fame,[4] which itself uses a modified version of the CSRankings scripting infrastructure.[1] In particular, to automate the process as much as possible and reduce errors, we created a scripting ecosystem that leverages DBLP and Google Scholar to gather most of the information. Using these scripts, we determined the top-cited ISCA paper per year, the top-50-cited ISCA papers of all time, number of authors per paper, and diversity information. However, classifying *topics* automatically was difficult, so we manually taxonomized them (discussed further in the "More Fine-Grained Analysis Across Decadal Trends" section).

We chose to use both DBLP and Google Scholar because each had a subset of the desired information in an easily accessible format. For example, Google Scholar makes it easy to count citations of papers, while DBLP makes it easy to perform more historical searches on ISCA publications in a specific year and/or by an author. Although these corpora are not guaranteed to have the exact same information (e.g., schema/formatting differences and temporal inconsistencies), we tried to keep each experiment self-contained and self-consistent with the data for that experiment. Finally, in some situations, especially for ISCA instances that predated the modern Internet, neither DBLP nor Google Scholar had the necessary information. In these situations, we used paper copies of each year's ISCA proceedings (e.g., the general chair's and program chair's welcome notes) as the final authority on information like number of submissions or accepted papers.

Additionally, when binning data across decades, because ISCA did not occur in 1975, we chose to represent the decades as 1973–1982, 1983–1992, 1993–2002, 2003–2012, and 2013–2022. Note that this means the first decade only contains nine years of publications, while the subsequent decades contain 10 years of publications. Moreover, in several cases, we chose not to include the 2023 publications for consistency. Likewise, in some cases (for example, creating word clouds), we selected 1973, 1983, 1993, 2003, 2013, and 2023 as "sample" years to examine trends; examining the yearly trends would have resulted in both a significant overlap and a significant processing overhead, without significantly affecting the results. Overall, this is the first comprehensive collection of information for a conference like ISCA, and we are releasing it publicly on Kaggle to enable further analysis and archival purposes.[13]

Top-Cited Papers and Decadal Trends

Table 1 shows the top-cited paper (per Google Scholar) for each year that ISCA occurred, with the highest-cited paper per decade highlighted in green.

Computer Architecture Over the Past Five Decades

Examining these top-cited papers per year shows some interesting trends in how ISCA's content and topics have evolved over the past 50 years.

TABLE 1. Top-cited paper for each ISCA from 1973 to 2022. The top-cited paper from each decade is bolded.

Year	Paper	Citations	Type	Topic
Decade 1 (1973–1982)				
1973	"Banyan Networks for Partitioning Multiprocessor Systems"	937	Micro	Interconnect
1974	"A Preliminary Architecture for a Basic Data Flow Processor"	854	Arch	Parallelism
1976	"Improving the Throughput of a Pipeline by Insertion of Delays"	126	Micro	Pipelining
1977	"A Large Scale, Homogeneous, Fully Distributed Parallel Machine, I"	747	Micro	Parallelism
1978	"DIRECT—A Multiprocessor Organization for Supporting Relational Database Management Systems"	364	Arch	Parallelism
1979	"Processor-Memory Interconnections for Multiprocessors"	183	Arch	Interconnect
1980	"A Comparison Connection Assignment for Diagnosis of Multiprocessor Systems"	711	Arch	Parallelism
1981	"**A Study of Branch Prediction Strategies**"	1203	Micro	Parallelism
1982	"Decoupled Access/Execute Computer Architectures"	373	Micro	Pipelining
Decade 2 (1983–1992)				
1983	"Very Long Instruction Word Architectures and the ELI-512"	854	Arch	Parallelism
1984	"A Low-Overhead Coherence Solution for Multiprocessors With Private Cache Memories"	800	Micro	Consistency/ Coherence
1985	"Implementing a Cache Consistency Protocol"	408	Micro	Consistency/ Coherence
1986	"Memory Access Buffering in Multiprocessors"	742	Micro	Consistency/ Coherence
1987	"Checkpoint Repair for Out-of-Order Execution Machines"	342	Arch	Parallelism
1988	"An Evaluation of Directory Schemes for Cache Coherence"	826	Micro	Consistency/ Coherence
1989	"Can Dataflow Subsume von Neumann Computing?"	391	Micro	Parallelism
1990	"Improving Direct-Mapped Cache Performance by the Addition of a Small Fully-Associative Cache and Prefetch Buffers"	2247	Micro	Cache
1991	"IMPACT: An Architectural Framework for Multiple-Instruction-Issue Processors"	503	Micro	Parallelism
1992	"**Active Messages: A Mechanism for Integrated Communication and Computation**"	2489	Micro	Parallelism
Decade 3 (1993–2002)				
1993	"Transactional Memory: Architectural Support for Lock-Free Data Structures"	3390	Micro	Parallelism
1994	"The Stanford FLASH Multiprocessor"	1052	Arch	Parallelism
1995	"**The SPLASH-2 Programs: Characterization and Methodological Considerations**"	5361	Tools	Benchmark
1996	"Exploiting Choice: Instruction Fetch and Issue on an Implementable Simultaneous Multithreading Processor"	1210	Micro	Parallelism
1997	"Complexity-Effective Superscalar Processors"	1285	Micro	Parallelism
1998	"Pipeline Gating: Speculation Control for Energy Reduction"	633	Micro	Power
1999	"PipeRench: A Coprocessor for Streaming Multimedia Acceleration"	707	Tools	Power

TABLE 1. (*Continued*)

Year	Paper	Citations	Type	Topic
2000	"Wattch: A Framework for Architectural-Level Power Analysis and Optimizations"	3840	Tools	Power
2001	"Cache Decay: Exploiting Generational Behavior to Reduce Cache Leakage Power"	1006	Micro	Power
2002	"Drowsy Caches: Simple Techniques for Reducing Leakage Power"	1221	Micro	Power
Decade 4 (2003–2012)				
2003	"Temperature-Aware Microarchitecture"	1644	Micro	Power
2004	"Transactional Memory Coherence and Consistency"	1027	Micro	Consistency/ Coherence
2005	"Interconnections in Multi-Core Architectures: Understanding Mechanisms, Overheads and Scaling"	646	Micro	Interconnect
2006	"Techniques for Multicore Thermal Management: Classification and New Exploration"	678	Arch	Power
2007	"**Power Provisioning for a Warehouse-Sized Computer**"	2625	Micro	Power
2008	"Technology-Driven, Highly-Scalable Dragonfly Topology"	967	Micro	Interconnect
2009	"Architecting Phase Change Memory as a Scalable DRAM Alternative"	1806	Micro	NVRAM
2010	"Debunking the 100X GPU Versus CPU Myth: An Evaluation of Throughput Computing on CPU and GPU"	1167	Tools	Simulator
2011	"Dark Silicon and the End of Multicore Scaling"	2405	Micro	Parallelism
2012	"RAIDR: Retention-Aware Intelligent DRAM Refresh"	617	Micro	DRAM
Decade 5 (2013–2022)				
2013	"GPUWattch: Enabling Energy Optimizations in GPGPUs"	710	Tools	Power
2014	"A Reconfigurable Fabric for Accelerating Large-Scale Datacenter Services"	1406	Micro	Interconnect
2015	"ShiDianNao: Shifting Vision Processing Closer to the Sensor"	1081	Arch	Machine Learning
2016	"EIE: Efficient Inference Engine on Compressed Deep Neural Network"	2727	Arch	Machine Learning
2017	"**In-Datacenter Performance Analysis of a Tensor Processing Unit**"	4307	Arch	Machine Learning
2018	"A Configurable Cloud-Scale DNN Processor for Real-Time AI"	532	Arch	Machine Learning
2019	"FloatPIM: In-Memory Acceleration of Deep Neural Network Training With High Precision"	182	Arch	Machine Learning
2020	"MLPerf Inference Benchmark"	307	Tools	Benchmark
2021	"Ten Lessons From Three Generations Shaped Google's TPUv4i: Industrial Product"	143	Arch	Machine Learning
2022	"BTS: An Accelerator for Bootstrappable Fully Homomorphic Encryption"	31	Arch	Security
2022	"CraterLake: A Hardware Accelerator for Efficient Unbounded Computation on Encrypted Data"	31	Arch	Security

In *decade 1* (*1973–1982*), multiprocessor papers were on the rise in ISCA's opening decade. The papers spanned several topics—networks/interconnects, fault diagnosis, and even relational database systems—all for multiprocessor systems or using multiprocessors. However, the top-cited ISCA paper of the decade was James Smith's "A Study on Branch Prediction Strategies." Although less highly cited, ISCA's first decade also saw papers on other topics such as computer architecture education (e.g., Jonathan Allen's ISCA 1976 paper). Although today these are more likely to appear in specialized conferences, in the 1970s, these were still nascent, and their presence in the main program demonstrated ISCA's focus on innovation across a variety of domains.

In *decade 2* (*1983–1992*), VLIW and multiple-issue processors add on to multiprocessor papers. Thanks to all this parallelism, cache coherence and consistency become important: papers on these topics were the top-cited papers in 1984, 1985, and 1988. Additionally, memory and cache architectures started to appear. For example, the Jouppi paper on victim caches and prefetch buffers is often still required reading for graduate computer architecture courses (perhaps unsurprising, given that it is currently the 10th-most cited ISCA paper of all time[12]).

In *decade 3* (*1993–2002*), the year 1995 introduced the SMT paper (again, often required reading for graduate computer architecture courses). Although the top-cited themes include superscalar processors and transactional memory, the highlight of the decade is the focus on power and thermal management. After decades of technology and power scaling, this flurry of activity suggested a leading indicator of the slowing of Dennard scaling that happened in the early 2000s. For example, the 1998 paper on "Pipeline Gating: Speculation Control for Energy Reduction" and back-to-back top-cited papers from 2000 to 2003 all talk about techniques to reduce power and/or model it.

In *decade 4* (*2003–2012*), ISCA witnessed the boom and, to some extent, the bust in the excitement around multicore architectures. Almost all the top-cited papers are either on multicore architectures or interconnects for multicore architectures. Power and thermal considerations carry over from the previous decade. Toward the end of the decade, the "Dark Silicon and the End of Multicore Scaling" paper highlighted major challenges with multicore scaling, which is power limited.

In *decade 5* (*2013–2022*), the natural next step to tackle Dark Silicon is specialization, and the final decade clearly demonstrates this: ISCA's fifth decade starts off with a top-cited paper on GPU power modeling and reconfigurable architecture, but subsequently, machine learning (ML) dominates the rest of the list. The top-cited papers from 2015 to 2021 are all ML architecture papers, while ISCA 2022's top-cited papers (thus far) are on accelerators for homomorphic encryption and encrypted data.

Notably, these big epochal changes were at the forefront of computer architecture advances that underpinned corresponding innovations in the broader computing ecosystem. Many of the ideas in these top-cited papers have directly influenced pervasively used industry products and systems.

More Fine-Grained Analysis Across Decadal Trends

Examining the top-cited paper each year provides broad trends about the most highly cited papers from each ISCA. To additionally provide more fine-grained analysis of how the field has evolved, we taxonomize these most cited ISCA papers by type and topic.[6] We first divide the papers across three broad types:

1) *Microarchitecture*: Architecture techniques that could be used inside many computers.
2) *Architecture*: A description or proposal of a full computer architecture.
3) *Tools*: Tools that help architects design computers, such as simulators or benchmarks.

We then further break down each of these areas into additional topics within these areas, specifically benchmarks, caches, consistency/coherence, dynamic random-access memory, interconnects, ML, Non-Volatile Random Access Memory (NVRAM), parallelism, power efficiency, pipeline, security, reliability, simulator, and tools. The results of this classification are shown in the "Type" and "Topic" columns in Table 1 as well as in Figure 1(a) and (b).

When evaluating the breakdown of the types of these papers, unsurprisingly, microarchitecture (57%) and architecture (31%) dominate the type of papers published at ISCA [see Figure 1(a)]; after all, these are and/or have been the primary focus areas in the computer architecture community over the past 50 years. However, it is interesting that tools that help architects make up a nontrivial (12%) number of these papers. Although only six of the 50 papers are tools papers, two ("The SPLASH-2 Programs: Characterization and Methodological Considerations" and "GPUWattch: Enabling Energy Optimizations in GPGPUs") are among the five most cited ISCA papers ever, and three of them are in the top-50 most cited papers ever thus far.[12] Given the importance of ML in recent years, we expect that another of them ("MLPerf Inference Benchmark") will

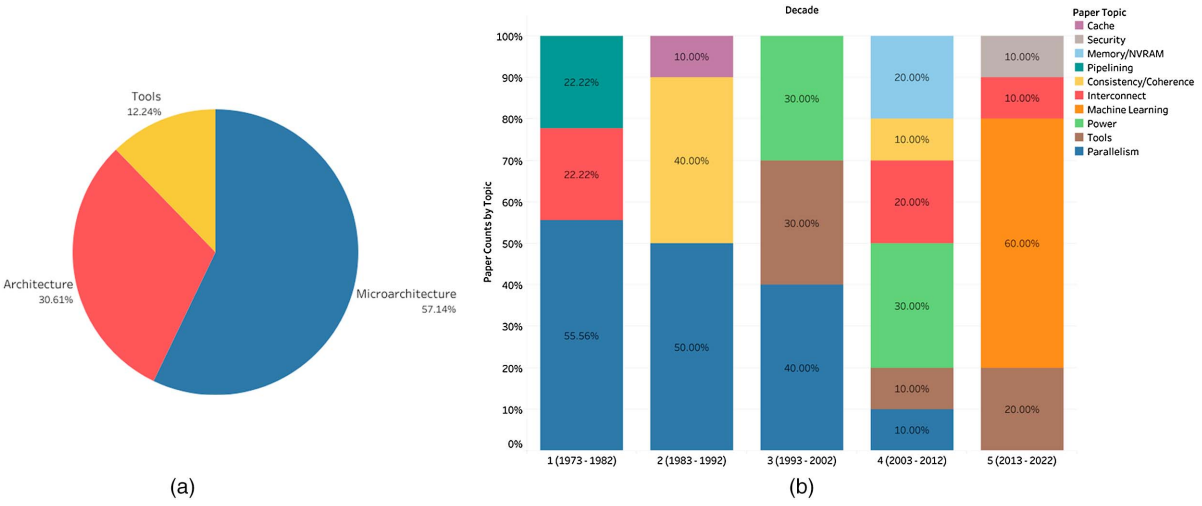

FIGURE 1. (a) Per-year most cited ISCA papers broken down by type. (b) Topic distribution in each decade for per-year most cited ISCA papers.

likely join them in the near future. Collectively, this highlights the importance of well designed, easy-to-use, high-fidelity tools and benchmarks that enable research in other subareas. These tools papers also often have significant staying power. For example, "The SPLASH-2 Programs: Characterization and Methodological Considerations," the highest-cited paper thus far at ISCA, appeared 12 years before the next oldest paper in the top-10 cited papers.[12]

The architecture topics in Figure 1(b) provide a finer-grained classification. Several interesting trends also emerge that complement our analysis in the "Computer Architecture Over the Past Five Decades" section. Similar to our analysis in that section, the top per-year cited papers show the importance, and pervasiveness, of how parallelism-related topics (e.g., parallelism, consistency/coherence, and interconnects) dominated the per-year top-cited ISCA papers in ISCA's early years. This echoes the community's transition from multiprocessors to multicore systems. However, over the first four decades, parallelism appears increasingly less as the top-cited topic from each year's ISCA, especially as Dark Silicon emerged. Likewise, as Dennard's scaling neared its end, papers on power efficiency became highly cited in ISCA's third and fourth decades. Collectively, these factors introduced the need for specialization (e.g., general-purpose GPUs and accelerators), which led to these topics increasingly being the most highly cited per-year ISCA papers in the fourth and especially fifth decades, where ML and security (due to issues like Spectre and Meltdown) dominated.

ML accelerators have an unexpectedly large slice of the *overall top-50 most-cited ISCA papers* as the current excitement about deep neural networks started only a decade ago; the seven ML papers in the top 50 were published at ISCA between 2014 and 2017. But given that ML papers in general are among the most highly cited in all of science and engineering,[2] the huge popularity of ML today likely accelerates citations to ML accelerators.

Pipelining (e.g., for multi-issue and VLIW processors) was highly cited in ISCA's first decade, but subsequently appears less frequently. Finally, other "eternal" topics such as caches, consistency/coherence, interconnects, security/reliability, and tools appear consistently across the decades, albeit sometimes with varying degrees of high citations. This highlights the importance of these topics in multiple different domains, as well as how their designs evolved over time.

A Different Perspective on High-Impact Papers

Thus far we have used the single highest-cited paper per year as a proxy for the most impactful paper for that year. However, citation count is a noisy and imperfect signal for impact. Thus, Figure 2 examines all the papers with citation counts within 15% of the top-cited paper each year. Overall, the trends are similar to Figure 1(a), and many of the topics have similar percentages. For example, parallelism and parallelism-related topics still dominate the first four decades of ISCA, and each decade their dominance decreases. However, there are some important distinctions. For example, examining additional papers show that security and reliability were not just highly cited in ISCA's fifth decade, they were also highly cited in ISCA's second decade. This corresponds to the challenges in designing

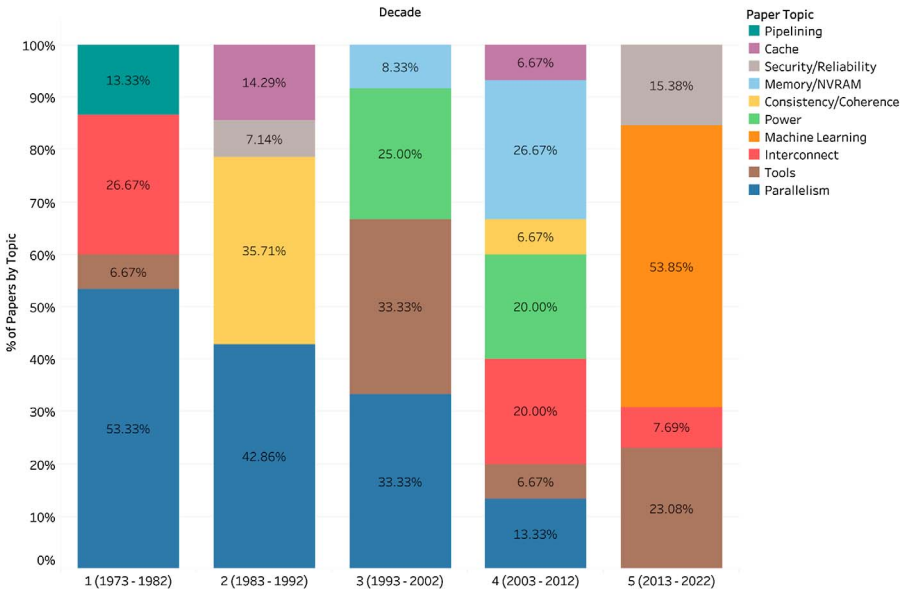

FIGURE 2. Topic distribution in each decade for papers with citation counts within 15% of the top-cited paper each year.

early multiprocessor and multicore CPUs (discussed further in the "Per-Decade Word Clouds" section). Thus, although Spectre and Meltdown fueled a resurgence in security as a highly cited topic at ISCA, it was not the first time that they were highly cited in the community.

Moreover, focusing on the single highest-cited paper per year, or even the papers within 15% of the top-cited paper per year, may potentially miss out on other high-impact papers from a given year. For example, in 1995, SPLASH-2 (thus far the top-cited ISCA paper of all time), the aforementioned SMT paper (thus far the ninth highest-cited ISCA paper of all time), and Multiscalar (thus far the 20th highest-cited ISCA paper of all time) all appeared. Each of these papers are more highly cited than the top-cited paper from many years of ISCA, including the first 16 ISCAs! Thus, only studying the top-cited paper *per year* would miss out on these foundational, important papers. However, SPLASH-2 is so highly cited that the SMT and Multiscalar papers are not within 15% of it. To address this, we also looked at the *top-50 highest-cited papers*, as shown in Figure 3(a) and (b). Comparing the per-year top 50 (Table 1) to the overall top 50 highlights several interesting trends. Certain topics like parallelism, interconnects, and ML have similar proportions across the decades, as shown in Figure 3(b). However, in Figure 3(a), microarchitecture (57%) and tools (12%) papers constitute a significantly different proportion of the per-year top-cited papers than the overall top 50 (70% and 8%, respectively). This potentially highlights how papers, especially those from earlier versions of ISCA, may not be cited as much relative to newer ISCA papers but still made significant impacts at the time. Moreover, it also shows how papers from recent years (e.g., on security), which have not yet had enough time to be cited enough to reach the top 50, are also making significant impact.

Note that focusing on citation counts across years comes with caveats. Over the years, ISCA has also seen significant growth and changes in many facets. One of the most interesting growth vectors has been in how other papers were cited. ISCA predates search engines and digital libraries. For example, ACM's Digital Library (https://dl.acm.org/) started in 1998 and IEEE *Xplore* (https://ieeexplore.ieee.org/Xplore/home.jsp) in 2000. Finding related work in the 1970s and 1980s was a lot of work: you had to go to libraries to read physical papers and to chase references, which reduced the number of references per paper. Moreover, ISCA [and other conferences, like the Architectural Support for Programming Languages and Operating Systems (ASPLOS) conference] changed policies on paper length for references in ~2015, hoping to increase the number of citations to computer architecture papers. The addition of arXiv, newer conferences like IEEE's International Symposium on High-Performance Computer Architecture (HPCA), and existing conferences like ASPLOS becoming annual instead of biannual, also collectively increase the number of computer architecture papers per year. As a result, there are roughly 5× more computer architecture papers published in top venues per year in 2023 than there were in 1990. Consequently,

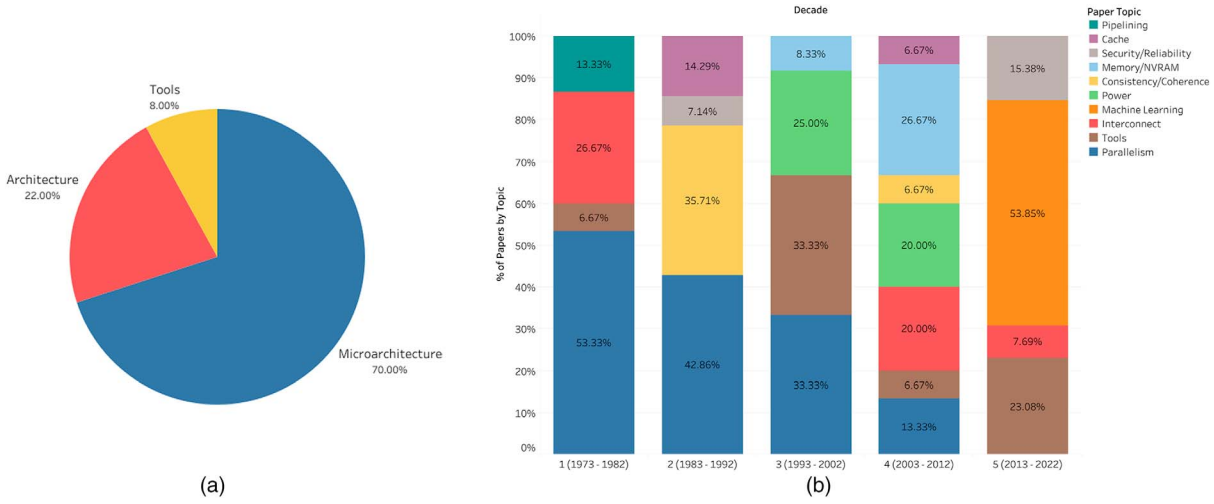

FIGURE 3. (a) Top-50 most cited ISCA papers of all time broken down by type. (b) Topic distribution in each decade for the top-50 most cited ISCA papers of all time.

publishing more papers per year and encouraging those papers to cite more related work has significantly increased the number of citations for more recent ISCA papers compared to ISCA papers from the 1970s and 1980s. For example, the 27 papers of the 1978 ISCA averaged 12 references (316 total), while 40 years later, ISCA had 63 papers that averaged 58 citations or 5× more (3636 total). The ISCA 2018 total was ∼12× more than 1978's. Thus, there might be 25× more citations from refereed computer architecture papers published in 2023 than from those in 1990. We also find that newer papers are more likely to cite other, more recent papers than older papers. Likewise, the number of authors per ISCA paper has been increasing over the years.[5,8] This evolution means that more recent ISCA papers are more likely to be highly cited, and because there are often more authors on those papers it is more likely that they will get cited, which impacts and biases metrics like examining citation count across all ISCA instances. As a result, recent ISCA papers are much more common in the overall top 50.

These hurdles also make it even more impressive that one paper from the first ISCA in 1973 made the top-50 list (number 42).[12] Its authors include Jack Lipovski, who was also the first general chair of ISCA. This is discussed further in an SIGARCH blog post that analyzes these trends in more detail.[6]

Per-Decade Word Clouds

To provide another perspective on how the computer architecture community's focus and content has evolved over years, we analyze keyword frequencies across the decades. Figure 4(a)–(f) shows word clouds of the most common words in abstracts from ISCA 1973, ISCA 1983, ISCA 1993, ISCA 2003, ISCA 2013, and ISCA 2023. The word clouds show the unique words across all abstracts in each year, after filtering for common words like "the." The number of unique words varies by year and ranges from 164 to 286 words. These word clouds highlight several interesting trends, some of which reinforce our findings from the "Computer Architecture Over the Past Five Decades," "More Fine-Grained Analysis Across Decadal Trends," and "A Different Perspective on High-Impact Papers" sections.

ISCA papers in the initial decade focused on designing hardware and systems (often for minicomputers), including specialized designs for avionics and meteorology. To a lesser extent, other topics like fault tolerance also appeared prominently in the first year of ISCA, hinting at the growing importance of this field as transistors became smaller and designs more complex and reinforcing why security/reliability were important in ISCA's early decades. In subsequent decades, parallelism at various levels was a common topic in ISCA 1983, 1993, and 2003. This includes fault-aware computing, multiprocessors, synchronization, memory, and caches, as well as instruction-level parallelism and other techniques to increase parallelism within a given processor. This mirrors the dominance of these topics in the per-year most cited papers and per-decade trends. However, as discussed previously, as Dennard's scaling faded, architects increasingly turned further toward parallelism and virtualization, initially continuing previous work on multiprocessors and multicore CPUs (ISCA 2003 and 2013) and later turning increasingly

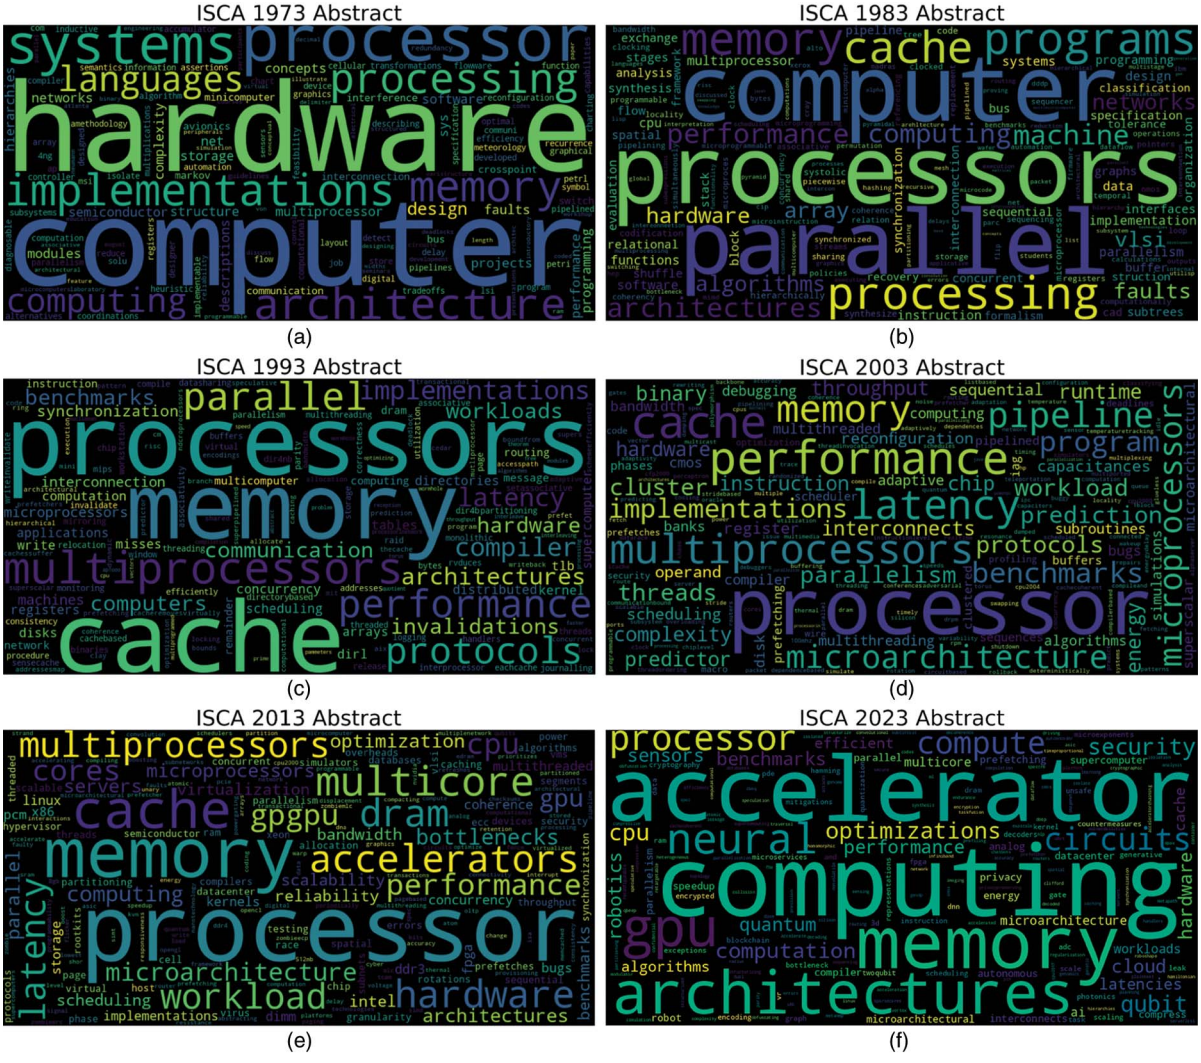

FIGURE 4. (a) Word cloud of the most common words in abstracts for ISCA 1973. (b) Word cloud of most common words in abstracts for ISCA 1983. (c) Word cloud of most common words in abstracts for ISCA 1993. (d) Word cloud of most common words in abstracts for ISCA 2003. (e) Word cloud of most common words in abstracts for ISCA 2013. (f) Word cloud of most common words in abstracts for ISCA 2023.

toward GPUs and other specialized accelerators (ISCA 2013 and 2023). Unsurprisingly, topics like accelerators (e.g., ML and robotics), quantum computing, and security appear much more prominently in ISCA 2023 relative to previous years, just as they dominated ISCA's fifth decade, as shown in Figures 1 and 2. This signifies the importance of these workloads in driving the requirements of future systems as well as the challenges associated with continuing to scale multicore CPU performance due to Dark Silicon, the end of Dennard's scaling, and the slowing of Moore's law.

The Life of a Highly Cited Paper

Figure 5 shows the trends in per-year citations for the top-cited papers per-year in the "Top-Cited Papers and Decadal Trends." The *y*-axis shows the number of citations for each year after publication. The dot closest to the *x*-axis represents the first year after publication, the dot after that the year after, and so on. The size of the circles is proportional to the number of citations that year. Moreover, for a few of the early ISCA papers, we do not have citation history, so the top-cited papers from those years are absent.

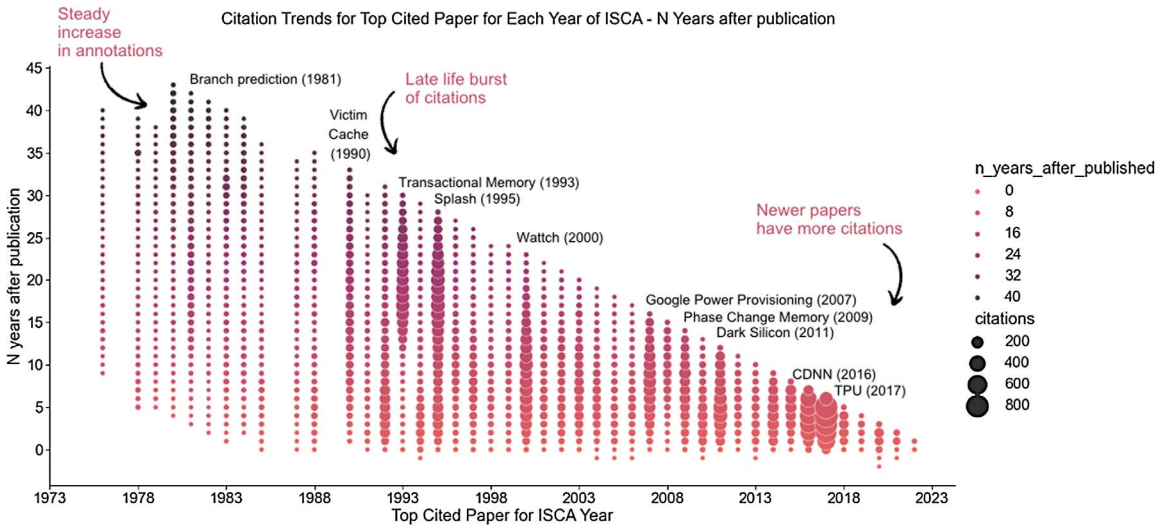

FIGURE 5. Citation trends for the top-cited paper of each year of ISCA. The *y*-axis shows the number of citations for each year after publication (proportional to the size of the circles). The gaps are because we do not have data for citation history for some of the earlier papers. CDNN represents EIE: Efficient Inference Engine on Compressed Deep Neural Network.

Although the citation paths are somewhat different for each of these papers, several common trends exist. We discuss those as well as highlight a few example papers that stand out. First, many of the top papers from the 1970s and 1980s see slow, steady citations, often continuing 35+ years beyond their publication date. However, with the exception of Smith's branch prediction paper in ISCA 1981 (a popular paper even today in modern computer architecture courses), these papers largely do not see the same large bursts of citations that papers from more recent ISCAs see. In comparison, the top papers from ISCAs after 2005, including the ISCA 2009 "Phase Change Memory" paper, ISCA 2011 "Dark Silicon and the End of Multicore Scaling" paper, Google's ISCA 2017 "In-Datacenter Performance Analysis of a Tensor Processing Unit" paper, and others often see very large bursts of citations very early. Some of these papers continue to be cited highly today. As a result, these papers have reached very high citation counts and are some of the highest-cited ISCA papers ever.[12] Collectively, these trends seem to further highlight the changes in citation policies and the increase in computer architecture papers per year. Interestingly, in between these time frames, there are several top-cited papers, including SPLASH-2 (the most cited ISCA paper ever), that exhibit even more interesting trends. These papers see initial bursts of high citations, followed by a period of fewer citations, and are then cited highly again later on. We believe that this highlights the prescience of some of these papers: as Dennard's scaling ended and Moore's law slowed, topics like parallelism (SPLASH-2 and "Transactional Memory: Architectural Support for Lock-Free Data Structures" at ISCA 1993) and power consumption (Wattch at ISCA 2000) became even more popular, and thus, these papers became highly relevant again later in life. However, as discussed in the "A Different Perspective on High-Impact Papers" section, changes in conferences and citation approaches also bias these results toward more recent papers.

COMMUNITY TRENDS OVER THE PAST FIVE DECADES

Scale and Growth Trends

Figure 6 shows the number of papers submitted to ISCA, number of papers accepted at ISCA, and ISCA's acceptance rate across the first 50 years of ISCA. The green, red, and blue lines show the rough trends for each of these metrics, respectively.

ISCA Paper Submissions

The growing number of submissions across the years reflect the growth of the computer architecture community: ISCA 1973 had roughly 70 paper submissions, while ISCA 2023 had roughly 400, a 6× increase in paper submissions. However, this growth has not always been linear. For example, ISCA 1988 had a significant spike in the number of paper submissions relative to the previous and subsequent years. More broadly,

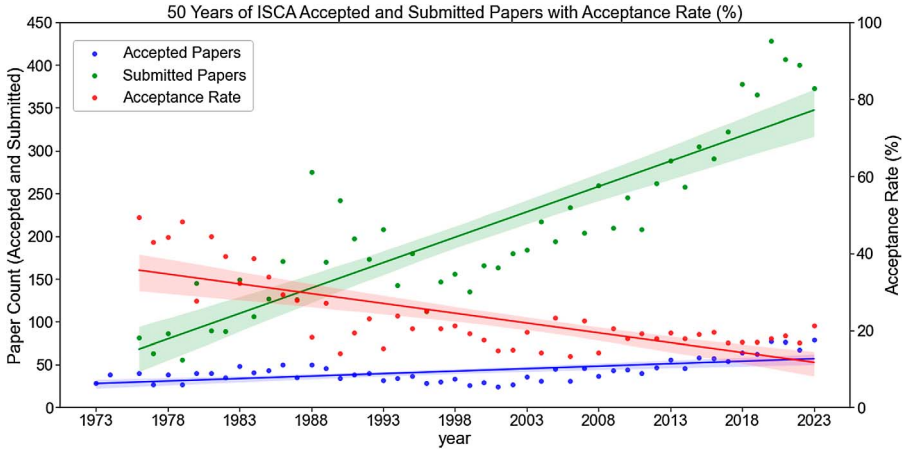

FIGURE 6. How the number of submissions, accepted papers, and acceptance rate varied for the first 50 years of ISCA.

ISCA in the late 1980s and early 1990s sometimes saw the number of paper submissions decrease year over year. In the case of the big increase in submissions for ISCA 1988, we believe that this can be explained by ISCA occurring in Hawaii for the first time. However, it is less clear why other years in this time frame saw fluctuations relative to the previous year. One possible explanation is that ISCA was attempting to keep to a single track. Although this seems like it would impact the acceptance rate (discussed further in the next section) more than paper submissions, perhaps this also had a chilling effect on the number of submissions as well. This was also the time period when a few other architecture conferences were ramping up, so there might have been a split in where papers were submitted, leading to some fluctuations. Nevertheless, an additional analysis as to why paper submissions decreased in this time frame is warranted.

Accepted ISCA Papers and Acceptance Rate

Although paper submissions have significantly increased in the last 50 years, the number of accepted papers (the blue dots in Figure 6) has grown much more slowly over the first 50 years of ISCA: from ISCA 1973 to ISCA 2023 the number of accepted papers has only grown by 2×, despite the 6× increase in paper submissions. Accordingly, the ISCA acceptance rate (the red dots in Figure 6) has steadily declined from ∼50% (ISCA 1973) to ∼20% (ISCA 2023). We believe that this highlights the increasing selectiveness of ISCA as the prominence of the conference has grown.

Number of Authors Per Paper

Figure 7 shows how the mean number of authors per accepted ISCA paper has varied across the first 50 years. Previous work in the late 1990s showed that the average number of authors per paper was steadily

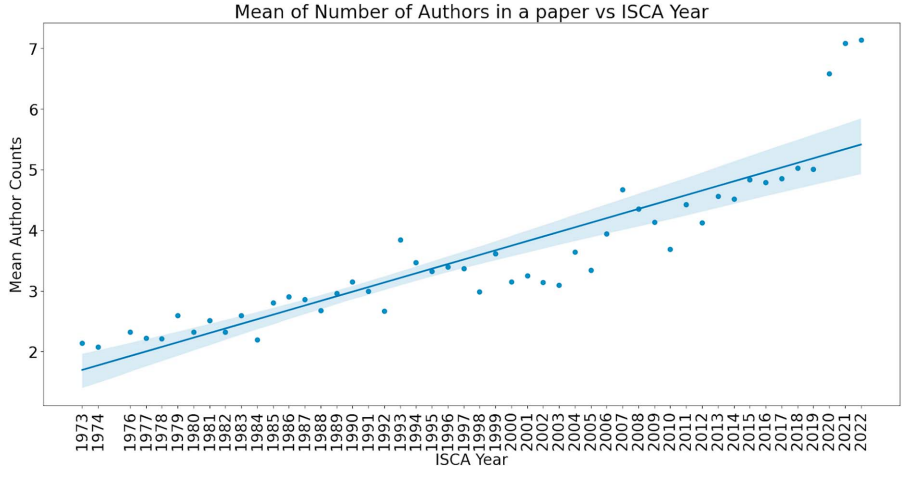

FIGURE 7. A depiction of how the mean number of authors per accepted paper has varied for the first 50 years of ISCA.

increasing at ISCA.[8] Figure 7 shows that this trend has largely continued to the modern day. Although the early years of ISCA had a number of papers with single or very few authors (e.g., a mean of 2.1 authors in 1973), the data show that these papers are becoming increasingly rare. As highlighted in previous work,[8,9] we believe that this highlights the increasing difficulty and challenges in building and researching increasingly large, more complex systems and the accompanying methodology [e.g., simulators, field-programmable gate-array (FPGA) platforms, or real systems] required to perform the research. However, interestingly, the trend is not completely linear; for example, 2007–2010 saw the mean number of authors decrease from 4.7 to 3.7. It is unclear why this drop happened, however, as there were papers such as the Anton and Corona papers, with large author counts in this time frame. Thus, we suspect that there is perhaps some noise in the cadence of paper acceptances here. There is also a large spike in the mean number of authors between 2019 (5) and 2020 (6.6). We believe that this spike correlates with the introduction of an industry track[7] as well as several industry papers in recent years, such as AMD's Exascale paper, Google's Tensor Processing Unit (TPU) papers, and Groq's Tensor Streaming Processor (TSP) papers. Each of these papers had at least 22 authors, and in general, papers on building large-scale projects such as these tend to have many authors, causing the mean to increase significantly as these papers became accepted more frequently at ISCA. However, even in other papers besides these, the number of authors appears to be steadily increasing.

PC Size

Given the large increase in paper submissions, the ISCA's main PC (i.e., not including external reviewers or referees) has needed to grow. As shown in Figure 8, although ISCA 1973 only had a PC of 11 members, ISCA 2023 had a PC of 113 members, a 10× increase in the number of PC members across the last 50 years! In fact, at last year's ISCA PC meeting, the number of people in the "conflict room" was sometimes larger than the size of the original ISCA PC from 1973! More generally, the number of people in the PC initially tracked the exponential growth in paper submissions. However, in recent years, PC growth rate has started to exceed the growth rate of paper submissions.

Diversity Trends

Geographical Diversity

Figure 9 presents the geographic locations ISCA has been held at in the first 50 years. Perhaps unsurprisingly given that ACM's Special Interest Group on Computer Architecture (SIGARCH) and IEEE's Technical Committee on Computer Architecture (TCCA) had a rule about hosting three out of every four ISCAs in North America, the majority of ISCAs have been held in the United States. Within the United States, ISCA has been distributed fairly evenly across the east and west coasts with a strong central presence as well.

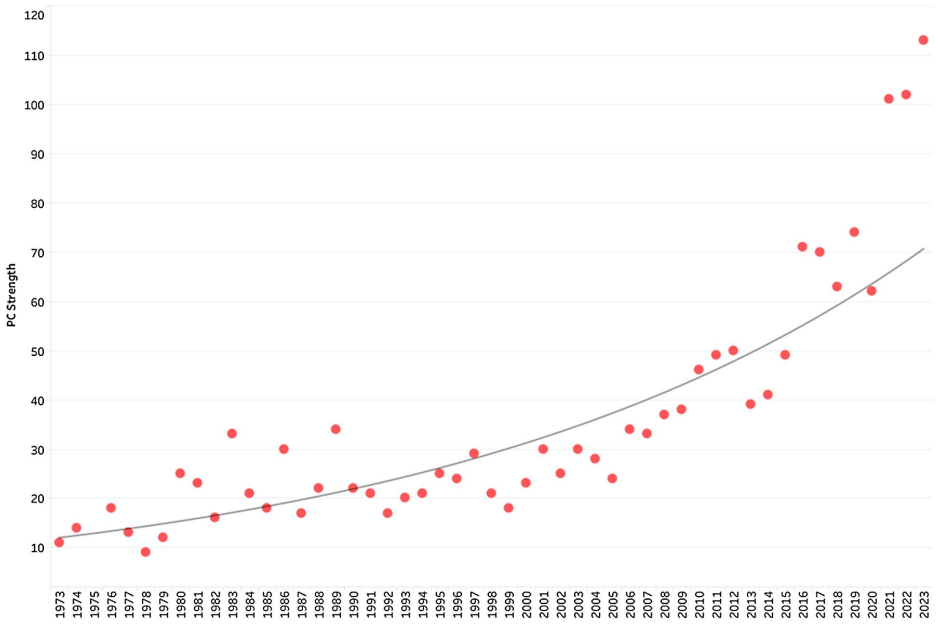

FIGURE 8. A depiction of how ISCA's PC size varied for the first 50 years of ISCA.

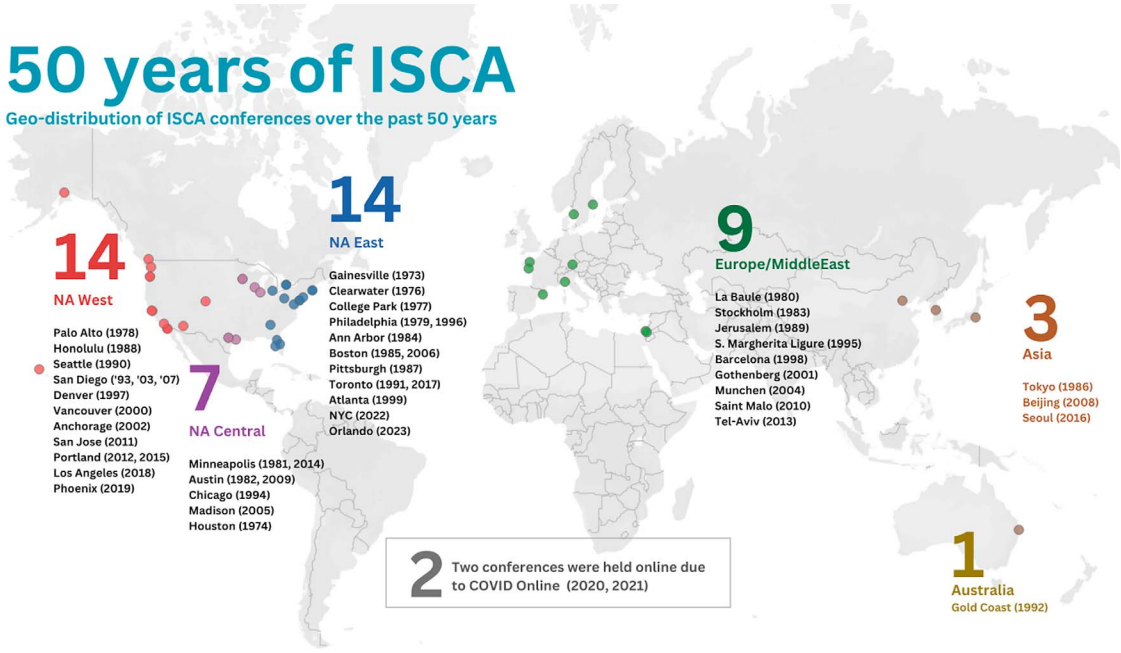

FIGURE 9. Geographical distribution of where ISCA conferences have been held in the past 50 years. NYC: New York City; S. Margherita Ligure: Santa Margherita Ligure.

Internationally, especially in recent years, ISCA has increasingly been hosted in a number of countries in Asia, Europe, and North America, including Canada, China, France, Israel, Japan, and South Korea. This continued geodiversity of where ISCA is held will continue next year as ISCA will be held in Buenos Aires, Argentina, the first time that ISCA will be held in South America. More broadly, the increased geodiversity is both a reflection of the increased internationalization of our community and the growing diversity in authors and paper submissions.

Gender Diversity

We also examined diversity in both ISCA's Hall of Fame[4] as well as for all authors who have had at least one ISCA paper in the past 50 years. We chose to study both because the ISCA Hall of Fame represents some of the most prolific and distinguished architects. Overall, both numbers highlight the need to continue working on improving diversity in our community: although we don't have self-reported data, we used Gender API to obtain an estimate.[11] Based on this, we found that only 9% of Hall of Fame authors are female, confirming trends in previous work that the field skews heavily male.[10] Figure 10 further confirms this. It plots the number of female authors each year as a percentage of all authors, with the standard deviation representing how confident Gender API was in its classification. The percentage of female authors has increased over the years, from 3% (1973) to 15% (2022). However, this fraction is still much lower compared to male authors, suggesting that we need to continue our outreach efforts in this area.

SOME QUALITATIVE PERSPECTIVES

ISCA Memories

To add a qualitative dimension to our retrospective, we also surveyed the community to get their memories of ISCAs over the decades. We e-mailed all 123 living members of the ISCA Hall of Fame and asked each of them three questions: 1) their favorite ISCA-related memory, 2) any highlights from the decade when they first attended ISCA, and 3) their favorite ISCA paper of all time. Although many declined to answer the third question, we collected a treasure trove of memories from attending the first 50 ISCAs.

The overwhelming impression that these stories yielded is that, in addition to the technical component, the respondents really valued the community and social aspect of ISCA. The most storied years were destinations in Europe and Asia: we heard many times about the banquet and its afterparty in Barcelona in 1998, the exquisite setting of ISCA 2010 in the French port town of Saint-Malo, and, more recently, the food

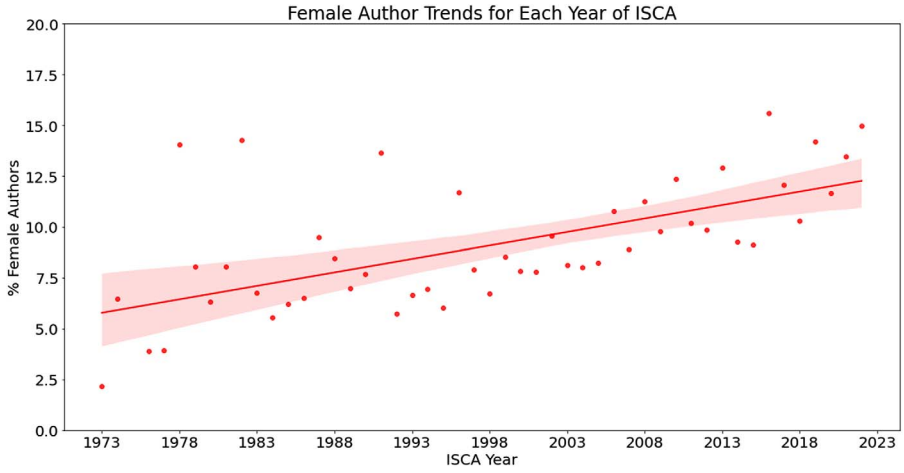

FIGURE 10. Author trends that are identified as female per Gender API for every year of ISCA. Gender classification was done using Gender API (a more detailed figure with error bars is in Upsani et al.[12]).

and entertainment in Seoul in 2016. ISCA researchers remembered when they met each other for the first time, when the banquet featured particularly good food, and when the conference gifts were especially useful. One attendee of ISCA 1996 in San Diego, for example, reports that she still uses the flip-flops from that year, which leave "ISCA-34" footprints on the sand behind you when you walk on the beach. In aggregate, the stories depict a long history of adding an in-person, interpersonal dimension to the research that makes up our field. And although it was only one, we did even hear from one ISCA attendee who has positive memories from a "virtual ISCA" in 2020 during the height of the COVID-19 pandemic.

THE STORIES HIGHLIGHT A KEY FUNCTION OF ISCA THAT HAS REMAINED CONSTANT THROUGH ITS FIVE DECADES: CONNECTING NEW ENTRANTS TO THE FIELD WITH ESTABLISHED RESEARCHERS TO COALESCE A COMMUNITY.

Another common theme recounted researchers' first-ever talk at ISCA. Some of these memories were purely fond ones, but more often they were memories of the nervous moments before the talk. That intimidating moment when you see famous and respected senior researchers in the audience seems to be a constant across the decades—even for ISCA denizens who now, decades later, play that same role for new authors at ISCA in 2023. The set of stories we collected contained at least one "I-first-met chain": when researcher *A* fondly remembered meeting their research hero, researcher *B*, at their first ISCA while, separately, researcher *B* submitted a nearly identical story about their excitement upon meeting researcher *C* at their first. The stories highlight a key function of ISCA that has remained constant through its five decades: connecting new entrants to the field with established researchers to coalesce a community. For more specific examples, we refer the reader to the special series of social media posts we launched as part of the #ISCA50 celebratory campaign.

ISCA50 Celebratory Panel

To celebrate ISCA's 50th anniversary, the conference hosted a panel discussion led by Partha Ranganathan and featuring David Patterson, Margaret Martonosi, Todd Austin, Onur Mutlu, Luis Ceze, and Thierry Tambe. In a lively and informative discussion, panelists shared insights about the past, present, and future of the field as well as insights into the community and their personal journeys.

Challenged to "channel their inner Gordon Moore" and make bold predictions, the panelists came up with some fascinating answers: a planet spackled with FPGA fabric enabling anyone to be a hardware designer, the blurring of lines between hardware and software and artificial intelligence (AI) automation in computer architecture, new technologies beyond CMOS, a new Moore's law where the number of chiplets doubled every generation, a perpetual AI assistant on everyone's wrist, nuclear fusion for energy sustainability, and even

AI-generated papers in the 75th ISCA! The panel also paid tribute to Moore, highlighting the boldness of his predictions and his charisma in getting the entire industry to coalesce around those. It is instructive that Moore's original prediction was published in a trade magazine, not at a top-tier venue, further highlighting the boldness of Moore's prediction at the time. Sharing advice for young researchers, the panelists encouraged working on new underserved areas; collaborating across disciplines; taking a full system's view, including software; and highlighted emerging areas like the impact of computing on global climate change, and quantum and biological computing.

Beyond technical topics, the panel also discussed the culture and community at ISCA. The wide-ranging discussion reiterated several themes discussed in earlier sections of this article. Memorably, a student question asking the panelists to "list their favorite instruction" stole the show: the resulting discussion and banter ended up being a live demonstration of the mix of technical passion and bonhomie defining the community. Panelists also offered valuable insights on fostering diversity (e.g., CARES, the Committee to Aid REporting on discrimination and haraSsment, policy violations) and the importance of being adaptive to change (e.g., taking a broader view of computer architecture or adding a demo session to ISCA). They also shared valuable advice on the nontechnical aspects of life, such as work–life balance and mental health. From pursuing "authentic happiness," getting enough sleep, and embracing analog activities like cooking or swimming, to lifelong learning and even suggestions for good books to read, the discussion featured a treasure trove of great ideas.

Overall, the panel reinforced the same findings from the rest of the analysis in this article: ISCA and more broadly computer architecture, is healthy, vibrant, and impactful, and there is significant excitement for the future. Perhaps this is best summarized by what one of the panelists (David Patterson) memorably said:

> "If the current AI/ML revolution is akin to searching for new worlds, and our colleagues in the machine learning sciences are the astronauts driving that mission, the computer architects are the ones building the rockets!"

50 YEARS OF ISCA: LOOKING BACK AND LOOKING AHEAD

When we set out to write this retrospective, much like the data-driven approach to computer architecture that is common at ISCA, we decided to take a data-driven approach to understanding and analyzing the past 50 years. To the best of our knowledge, our work is unique in its rich archival data sets and analysis tools (collated and curated across numerous data sources), and our subsequent analysis to tease out insights about the community and its impact. We have been surprised at the richness of insights from such a data-driven exercise and the power of the visual narratives from our analysis. For example, inspired by one of the pictures from this study (see Figure 9), researchers are organizing the first-ever computer architecture conference in Africa. We wish that effort well. To enable further research, we are also making the data available publicly, along with a companion arXiv version of this article with extended additional analysis.[12] At the same time, traditional qualitative approaches to reflecting on the past have also continued to be effective. For example, discussions at the ISCA50 celebratory panel have led to serious discussions about having a new demo session at future ISCA conferences.

RESEARCH AT ISCA HAS PARALLELED AND PRESAGED THE COMPUTING ADVANCES WE HAVE SEEN IN THE LAST HALF CENTURY.

Overall, ISCA has had a phenomenal 50 years of impact. The topics have changed and the community has grown, but a constant across five decades has been the growing influence and impact to the broader computer architecture community and to the broader computing ecosystems. Research at ISCA has paralleled and presaged the computing advances we have seen in the last half century, including minicomputers, general-purpose uniprocessor CPUs, multiprocessor and multicore CPUs, general-purpose GPUs, and accelerators. Looking ahead, the future looks exciting. The slowing of traditional technology scaling at a time when demand for computing is exploding due to workloads like AI/ML provide rich impetus and opportunities for new, foundational computing advances, from underlying substrates, to new architectures, to novel hardware–software co-design. Topics like accelerators, biological-inspired computing, carbon-aware computing/environmental sustainability, edge computing, ethical computing, quantum computing, and security/privacy are all likely to play important roles in the next few decades. New approaches to computer systems, such as the use of ML in designing, verifying, and operating systems, are likely to change how we think about research. Beyond innovating around new technologies,

there are rich opportunities to rethink abstraction layers (across compilers, operating systems, VLSI, and others). We can't wait to see what the 75th year retrospective will look like!

ACKNOWLEDGMENT

We would like to acknowledge Lizy John, Hyesoon Kim, and Tushar Krishna for their support and input to this article. We also thank Mark Hill for his hand-me-down collection of ISCA proceedings, which proved invaluable.

REFERENCES

1. E. D. Berger. "CSRankings: Computer Science Rankings." CSRankings. Accessed: Jun. 1, 2023. [Online]. Available: https://csrankings.org
2. B. Crew, "Google Scholar reveals its most influential papers for 2020," *Nature*, Jul. 2020. [Online]. Available: https://www.nature.com/nature-index/news/google-scholar-reveals-most-influential-papers-research-citations-twenty-twenty
3. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, 2016, pp. 770–778.
4. M. D. Hill, M. D. Sinclair, and G. S. Sohi, "An International Symposium on Computer Architecture (ISCA) Hall of Fame." UW-Madison. Accessed: Jun. 1, 2023. [Online]. Available: https://pages.cs.wisc.edu/~arch/www/iscabibhall.html
5. P. M. Kogge, "Long term trends of computer architecture research funding as seen through ISCA," Univ. of Notre Dame, Dame, IN, USA, TR #2005-16, 2005.
6. D. Patterson. "What are the most cited ISCA papers?" Computer Architecture Today Blog. [Online]. Available: https://www.sigarch.org/what-are-the-most-cited-isca-papers/
7. D. Patterson. "Genesis and reflections on the return of industry products to ISCA 2020." SIGARCH Blog. [Online]. Available: https://www.sigarch.org/genesis-and-reflections-on-the-return-of-industry-products-to-isca-2020/
8. S. S. Mukherjee, "What should graduate students know before joining a large computer architecture project?" *ACM SIGARCH Comput. Archit. News*, vol. 25, no. 1, pp. 23–26, Mar. 1997, doi: 10.1145/250015.250019.
9. S. M. Blackburn, K. S. McKinley, and L. Xie, "Author growth outstrips publication growth in computer science and publication quality correlates with collaboration," Sep. 2019, *arXiv:1909.02212*.
10. L. L. Wang, G. Stanovsky, L. Weihs, and O. Etzioni, "Gender trends in computer science authorship," *Commun. ACM*, vol. 64, no. 3, pp. 78–84, Mar. 2021, doi: 10.1145/3430803.
11. Gender API. [Online]. Available: https://gender-api.com/
12. G. Upasani et al., "Fifty years of ISCA: A data-driven retrospective on key trends," Jun. 2023, *arXiV:2306.03964*.
13. G. Upasani et al., ISCA@50 Datasets. Accessed: Oct. 16, 2023. [Online]. Available: https://www.kaggle.com/datasets/isca50/50-years-of-isca

MATTHEW D. SINCLAIR is an assistant professor with the Computer Sciences Department, with a joint appointment with the Department of Electrical and Computer Engineering, University of Wisconsin–Madison, Madison, WI, 53706, USA. Contact him at sinclair@cs.wisc.edu.

PARTHASARATHY RANGANATHAN is an engineering fellow/vice president at Google, Mountain View, CA, 94043, USA, where he is the area technical lead for systems hardware and data centers. Contact him at partha.ranganathan@google.com.

GAURANG UPASANI is a software engineer at Google, Mountain View, CA, 94043, USA. Contact him at gupasani@google.com.

ADRIAN SAMPSON is an associate professor at Cornell University, Ithaca, NY, 14850, USA. Contact him at asampson@cs.cornell.edu.

DAVID PATTERSON is a distinguished engineer at Google and a professor emeritus at the University of California, Berkeley and Google, Berkeley, CA, 94720, USA. Contact him at pattrsn@berkeley.edu.

RUTWIK JAIN is a Ph.D. candidate with the Computer Sciences Department, University of Wisconsin–Madison, Madison, WI, 53706, USA. Contact him at rnjain@wisc.edu.

NIDHI PARTHASARATHY is a high school student at Lynbrook High School, San Jose, CA, 95129, USA. Contact her at nidhi.parthasarathy@gmail.com.

SHAAN SHAH is a Ph.D. candidate with the Department of Electrical and Computer Engineering, University of California, San Diego, San Diego, CA, 92093, USA. Contact him at shaan.shah365@gmail.com.