

Using High-Performance Computing Artifacts as a Learning Intervention: A Systematic Literature Review

Nkundwe Moses Mwasaga[†]
School of Computing
University of Eastern Finland
Joensuu, Finland
nkundwe.mwasaga@uef.ac.fi

Mike Joy
Computer Science
Warwick University
Coventry, UK
m.s.joy@warwick.ac.uk

ABSTRACT

High-Performance Computing (HPC) artifacts provide opportunities for students to improve their understanding of parallel computing, which is important for students who study computer science. In line with that, many computing departments are integrating HPC systems into their curricula. However, there is a need to investigate HPC artifacts that have been used as learning interventions. This study has employed a systematic literature review to investigate published papers on HPC education from 1988 to 2018. The findings of our investigation of a stratified sample of 211 papers reveal the state of the practice of application of HPC artifacts in computing education in terms of the contexts, themes, nature and topics of the publications. The study revealed that a majority of publications reported the usage of Beowulf and other clusters as the pedagogical tools. Furthermore, the study discovered gaps in research on the application of HPC artifacts in ability and aptitude, teaching and learning, teaching and learning techniques, curriculum, parallel programming, and parallel processing. This study contributes to our understanding of what HPC artifacts are used in computer science education.

CCS CONCEPTS

• Computer systems organization • Applied computing • Applied computing-Education • Applied computing-Education-Computer-managed instruction

KEYWORDS

High-Performance Computing artifact, Systematic literature review, HPC education, Computer science education, Beowulf clusters

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

ICONIC, September 24–25, 2020, Plaine Magnien, Mauritius

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7558-0/20/09...\$15.00

<https://doi.org/10.1145/3415088.3415130>

ACM Reference format:

Nkundwe Mwasaga and Mike Joy. 2020. Using High-Performance Computing Artifacts as a Learning Intervention: A Systematic Literature Review. In *Proceedings of 2nd International Conference on Intelligent and Innovative Computing Applications (ICONIC'20)*. ACM, Plaine Magnien, Mauritius, 10 pages. <https://doi.org/10.1145/3415088.3415130>

1 Introduction

The increased usage of multi-core systems and computing clusters has created a demand for parallel computing skills, economies demand a workforce with parallel computing knowledge and skills, and a prior survey study has shown the use of HPC in government and industry [26]. As a result, supercomputers, cloud computing systems, and High-Performance Computing (HPC) systems have been integrated to support academic curricula [1], scientific computing [75] and business [25], and a survey study of cloud computing systems and education has shown improvements in classroom teaching [40]. Other HPC systems that have been developed and integrated in academic institutions to improve students' HPC skills and access to affordable HPC include virtual clusters [51][58], a bootable cluster CD [64], multi-platform for distributed computing [76], exascale systems [21][53][62], embedded and mobile system-on-chips [52]; and a portable computational cluster [13].

The rapid implementation of HPC systems has been facilitated by the construction of the first Beowulf Linux commodity cluster at NASA's Goddard Space Flight Center in 1994 [33]. Empirical research projects that use Beowulf cluster for learning and teaching [1][27]; programming[4]; exascale computing[16]; education[28]; massive parallel computing[30]; and education in HPC[71] have demonstrated evidence that since then, Beowulf clusters have been widely used to support computer science teaching and research [48]. Beowulf clusters have been used in the areas of large-scale visualization, simulation of complex neural network models, and evolutionary computation [30] [28]. The support provided by HPC systems can be classified in order to gain an understanding of how different implementations of clusters support computer science curricula.

In this study, publications were investigated to establish evidence of the usage of HPC artifacts using a systematic literature review method under the design science paradigm. In investigating the HPC education publications, we adopted Simon's classification [67][5] as an ontological classification, which classifies papers within the domain of computer science education into four broad dimensions, and the topics, context, nature, and scope of each are then subsequently broken down into categories. Moreover, the study [5] demonstrated that the context identifies the generic thrust of the published paper, the topics list the themes the paper addresses, the scope is the involvement of actors in the research, and nature shows the distinction between research and practice papers.

This study aims to investigate whether HPC artifacts have been used as learning interventions to improve parallel computing knowledge and skills in the context of computer science education over time. The investigation helps gain a more in-depth insight into the trends of using HPC artifacts in education and how HPC artifacts have been used as learning intervention tools.

This paper is organized to answer five research questions commencing with the description of the protocol that is used, followed by a synthesis of the findings which address the research questions, and a summary concludes.

2 Methodology

We use the Systematic Literature Review (SLR) methodology under the paradigm of Design Science Research [2][41][37] [66] to guide the investigation of HPC artifacts that have been used as learning intervention and reported in online scientific publication depositories. According to Kitchenham [10], SLR, as the critical step in conducting scientific research, is the method that analyses all relevant primary research publications by identifying, mapping, evaluating, aggregating, and interpreting based on a particular research question, phenomenon of interest or topic area. SLR also is used to identify gaps in particular topics to be filled. This is in line with the aim of SLR, which is to construct a broader view of the research question by summarizing the literature with minimum bias [3]. Hence this section describes an overview of the research methodology that has been followed in this study, explains the protocol used, articulates the research questions pursued in the study, describes the search strategy, and articulates the exclusion and inclusion criteria used.

2.1 Protocol for SLR

In an SLR, the protocol that guides the study to identify publications should be described and justified to guarantee the reproducibility of the work [14] and minimize subjectivity, during data collection, by the researchers. The protocol used in this SLR study details the planning method on how the research was conducted. In this study, we reviewed existing protocols that guide how to perform SRL[9]; SRL in software engineering[11]; a

guideline for performing SRL in software engineering [69]; SRL for mobile cloud computing [20]; guidelines on how to perform SRL[17][44]; and procedures on how to perform SRL[10] and agreed to follow the modified SLR protocol [17] that includes the following steps:

- Research questions and conceptual framework;
- Selection of work team;
- Research Strategy;
- Search, eligibility and coding;
- Quality assessment;
- Synthesis of results; and
- Study presentation.

The SLR protocol that we adopted follows the steps that are described in the following sections.

2.1.1 Research Questions. The research questions for this study aimed to investigate primary studies that are related to interventions of HPC artifacts in the context of computer science education, which is useful for gaining a broader view of the HPC education field. In categorizing HPC education, we have considered publications on HPC if they did refer explicitly to any form of HPC artifacts in the education context. The decision to focus on the labeling of HPC artifacts guaranteed the appropriate level of focus needed in the SLR, though such a decision may have led to the exclusion of publications that *indirectly* referred to HPC without using education [61]. The retrieved studies relevant to our research questions were diverse, heterogeneous, and difficult to classify. To elucidate the process of classification, we drew on Simon's classification [67] to sort 211 studies in categories. Using Simon's classification, the investigation categorized the interventions into topics, nature, scope, and contexts. For the sake of this study, we have included parallel computing artifacts as a separate category. The purpose of this category is to identify the types of HPC artifacts that have been used as learning intervention tools in computer science education. The investigation provides answers to the following research questions.

RQ1: What are the reported topic areas in computer science education publications where HPC artifacts have been used as learning interventions?

RQ2: What are the contexts in which computer science education uses HPC artifacts?

RQ3: What is the scope in the usage of HPC artifacts as learning interventions supporting different curricula?

RQ4: What is the nature of the publications that have used HPC artifacts as learning interventions supporting different curricula?

RQ5: What parallel computing artifacts have been used as learning interventions in computer science education?

We extracted initial keywords from research questions [9], as indicated in Table 1. Based on the research questions, the HPC education publications were used in this study. To be specific, we avoided restricting our review study, and we decided to cover any

type of HPC system used as an intervention in computer science education, learning, teaching, or training. The intervention is any use of an HPC system as an educational tool that addresses an educational issue in computer science and has been implemented at one or more academic institutions. The academic environment is the context that is reviewed in the published HPC studies, and the HPC system artifacts must be implemented or used within such an academic environment.

Based on that, initial keywords were used to derive compound keywords that were used to search for a variety of HPC publications relevant to the research questions of this study.

Table 1. Research Questions and Keywords

Research Question	Keywords
RQ1	Computer, Science, Education, HPC, Artifacts, Learning, Interventions
RQ2	Computer, Science, Education, HPC, Artifacts.
RQ3	HPC, Artifacts, Learning, Intervention, Curricula.
RQ4	HPC, Artifacts, Learning, Intervention, Curricula.
RQ5	Parallel, Computing, Computer, Science, Education, Artifacts, Learning, Interventions.

In order to answer the research questions, relevant publications from online scientific publications depositories were reviewed.

2.1.2 Selection of Work Team. We created a work team that was led by the first author, who did almost all the processes of this SLR, although it is rare for a single individual to have all necessary methodological and technical skills and knowledge to conduct an SLR. The other reason for the creation of the team is the inherent time taken for a single person to conduct an SLR study. Having a team for an SLR increases the quality of the review when eligibility and search of the publications and the coding of results are independently performed by two members of the team[42].

Based on the review questions, the technical knowledge of HPC and methodological knowledge were applied in order to conduct the SLR. In order to increase the quality of the review, the processes of searching sources and criteria, coding, eligibility of publications, and the synthesis process were performed independently by each member of the team.

2.1.3 Research Strategy. We used the SLR method to investigate the state of HPC artifacts used as interventions in computing education and training [46]. The investigation began in July 2016 and continued until December 2018. The methods used to retrieve publications were searching from scientific publication depositories and chaining from known publications. We reviewed studies to find the database repositories that are

useful for our study [47][36][49][57][46] and identified Scopus, SciFinder, Web of Science, ACM, Google Scholar, and IEEE as suiting our research questions. The study reviewed HPC publications from 1988 to 2018, which related to education. The investigation used three online scientific publication repositories as the data sources (ACM Digital Library, IEEE Xplore, and Google Scholar). The pilot search showed they were sufficient to retrieve education-related publications HPC.

We used the online scientific publications repositories to search for keywords in order to conduct the SLR. In order to do extensive searches, the repositories were searched using full text, abstract, and title of desired study types that were relevant to answering the questions posed in the context of this study. The review process included publications that were published from 1988 to 2018, where 1988 was the years of publication of the three papers which commenced the corpus that was available for review [63][42][10]. The publications reported usage of a supercomputer in the contexts of research, learning, and training, to enhance the management and delivery of innovative education models [59][34]. The first publication focused on the integration of supercomputers to support research in Boeing, and the lessons learned [63], and the thrust of the second one was the study of the watershed phenomenon in economics that measured the demand and supply of supercomputers in research institutions [70]. The third one reported on the supercomputing programme [21], whose aim was to create adaptable consumers of supercomputing technology and to bridge a gap between computer science education and computer-specific training [22]. The use of Simon's classification lets us discover the trends in the use of HPC artifacts as pedagogical tools in computer science and other disciplines that use HPC artifacts.

When there was any potential ambiguity about the inclusion of a paper or challenges associated with suppositions and assumptions raised during the review, the paper was included if all members of the team agreed by consensus.

2.1.4 Search, Eligibility and Coding. The process of searching papers in online scientific publication depositories requires search strings to enable retrieval of relevant publications. The search strings have been derived from the research questions in this paper. We composed the initial search strings from the search terms and main keywords identified in these research questions and the relevant synonyms. To reduce the likelihood of bias, we calibrated our search strings in a pilot search in a single scientific publication repository.

The research questions concern education topics where HPC artifacts have been used as educational tools, hence the keywords for RQ1 are *computer, science, education, HPC, artifacts, learning, and interventions*. We used keywords derived from Table 1 to construct search strings for this study [9]. Similar terms and synonyms for HPC artifacts that were used to create search strings in this research are *parallel distributed systems, high-*

performance computing, parallel computing, Beowulf cluster, supercomputing, parallel computing system, and high-performance computing system. Based on what was mentioned earlier, the synonyms for education used in this research are *training, learning, and teaching.* The combination of these was used to formulate search strings. Therefore the keywords derived from RQ1 were used with the synonyms and related terms to create search strings to retrieve relevant data to answer RQ2, RQ3, RQ4, RQ5, and RQ6.

As a consequence of this, the keywords which were used to cover a variety of HPC publications in this study were:

- high-performance computing education;
- parallel computing education;
- Beowulf cluster education;
- supercomputing education;
- teaching parallel distributed computing;
- parallel computing systems education;
- high-performance computing system education;
- parallel distributed computing training;
- high-performance computing training;
- parallel computing training;
- Beowulf cluster training;
- Beowulf cluster teaching;
- supercomputing training;
- parallel computing systems training;
- Beowulf cluster learning;
- learning parallel distributed computing;
- high-performance computing learning;
- learning supercomputing;
- high-performance computing system training.

To improve the study's validity, we assessed the risk of bias in the included publications [3][43]. In doing so, we involved two evaluators who independently performed SLR and ensured the process involved searching multiple publication databases to minimize publication, extraction, selection, and inclusion bias.

2.1.5 Quality Assessment. In an SLR, the confidence in the study's conclusion and usefulness to the consumers of the study depends on relevance and quality of the primary studies and the process of review [6] [72]. We included 211 studies for further quantitative synthesis after passing assessment criteria from the modified dimensions of quality and relevance of the publications [6]. So the assessment criteria we used to ascertain the quality and relevance of the publications are the following.

1. *Quality of the study performance:* The artifact used in the publication meets the required standard for HPC artifact used in education. The required standard is an HPC artifact used as an intervention to support teaching, learning, and research.

2. *Relevance to the research question:* The publication addresses exactly the target subject of the research question of the SLR;
3. *Relevance to the review focus:* The study covered in the publication is in a similar education context to the one defined in the SLR.

In this sense, all 211 studies passed quality and relevance criteria and proceeded to further quantitative analysis using aggregative synthesis.

2.1.6 Description and Analysis of the Studies

2.1.6.1 Criteria of Exclusion and Inclusion Criteria. The search used the stipulated keywords, terms, and synonyms, as listed in the previous section, and the entire text of each publication within each repository was searched in order to retrieve relevant publications.

The search used the entire publication texts, and the scope of the applied searches was restricted to the following inclusion criteria:

- i. Publications that have been written in English;
- ii. The year of publication was from 1988 to 2018 (inclusive); as this SLR was conducted in 2018 no future publications have been included in the study;
- iii. Conference and journal publications;
- iv. Title, keywords and abstract based on the search string and keywords, and the content of the research question matches the content of the abstract;
- v. Introduction and conclusion contain content that deals with HPC education;
- vi. The full content of the publication deals with HPC education and usage of the HPC artifact in the education context;
- vii. The focus of the publication is teaching, learning, training, and research in education.

The exclusion criteria are:

- i. Publications that do not include relevant keywords;
- ii. Publications that are not in the area of teaching, learning, training, and research in education.

2.1.6.2 Analysis of the Studies. It was observed that many publications in the ACM Digital Library could also be searched and found in Google Scholar. Hence we considered IEEE Xplore and Google Scholar for further review. The keywords, terms, synonyms, and phrases were used to create query strings. The created query strings were combined using the OR search operator among related terms. The OR operator has been used in order to retrieve any publication with any of these search terms.

The combined search strings that were joined by the OR operator to retrieved publications were:

- A = "High-Performance Computing education" OR "Parallel computing education" OR "Beowulf cluster education" OR "Supercomputing education."
- B = "Parallel computing systems education" OR "Parallel distributed computing education" OR "High-Performance Computing System education" OR "High-Performance Computing training"
- C = "Parallel computing training" OR "Beowulf cluster training" OR "Supercomputing training" OR "Parallel computing systems training"
- D = "Parallel distributed computing training" OR "High-Performance Computing System training" OR "Teaching Beowulf Cluster" OR "Teaching High-Performance Computing"
- E = "Teaching Parallel distributed computing" OR "Teaching High Performance Computing Cluster" OR "Teaching Supercomputing" OR "Teaching High-Performance Computing System"
- F = "High-performance computing learning" OR "Beowulf cluster learning" OR "learning parallel distributed computing" OR "learning supercomputing."

The number of hits for each final combined search strings and the total number of papers identified was 831, as shown in Table 2.

Table 2. Final Keyword Search Results with OR Operator

No.	Keywords	IEEE Xplore, Hits	Google Scholar, Hits
1.	A	69	409
2.	B	72	43
3.	C	147	57
4.	D	24	104
5.	E	36	22
6.	F	59	42
TOTAL		401	430

The publications were further reviewed and evaluated for inclusion by scrutinizing the titles, abstracts, and full text in a reliable and robust manner. The filtering process used the search strings based on the relevance of the contents of the research questions. After scrutinizing the search results from three online scientific publications depositories for duplicates, we restricted our detailed search to 641 publications. The manual process of searching potentially relevant publications using titles of publications was conducted. It yielded 216 results because the duplicate, non-educational, non-English language, and non-conference or journal publications were removed, as shown in Figure 1, representing the PRISMA diagram that systematic reviewers typically use to represent the flow of studies through the SLR [50].

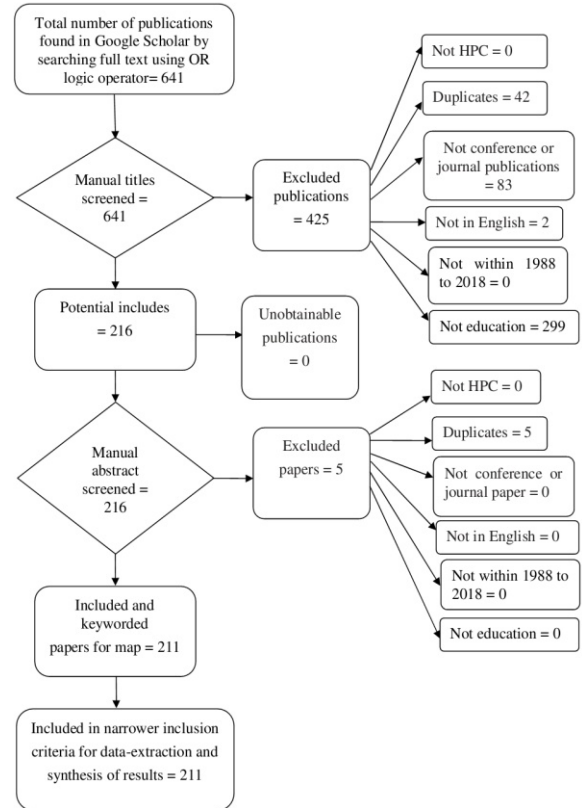


Figure 1: Inclusion and Exclusion Results (adapted from [35])

Furthermore, we manually read the abstracts of each publication, yielding 211 potentially relevant publications. Thus, consistent with the methods used in [7][20][44] and the suggestion of including all [6] as applicable for quantitative synthesis in SLR. Then the inclusion and exclusion criteria enabled the retrieval of 211 eligible publications.

3 Results

A total of 211 HPC education publications written in the English language were identified using the search strategy. The analysis of the publications that cover the period from 1988 to 2018 has shown that the focus in the period from 1988 to 1993 has been the usage of shared memory multiprocessors (supercomputing) systems that support training and research.

From 1994 the usage of distributed memory supercomputing systems (HPC clusters) in education started. This trend coincides with the introduction of affordable Beowulf clusters in the family of HPC devices in 1994, as indicated in [33]. The number of HPC publications increased steadily from 1992 – 2006 to a total of 102

publications. From 2007 – 2012, the number decreased to 46 HPC publications, as shown in Figure 2.

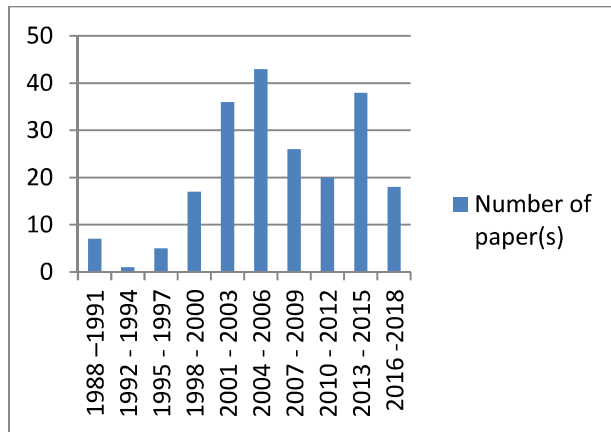


Figure 2: The Number of HPC Education Papers from Range of Years of Publications

The chosen starting point of the reviewed HPC education publications was from 1988 to 2018. In the years 2004-2006, the focus of the publications was in virtual clusters, Beowulf clusters, grid computing, high throughput computing systems, cloud HPC, and multiprocessor systems, which constitute the largest number of publications in this study. The family of Beowulf clusters forms the dominating systems that have been used in HPC education from 1994 to 2018.

As good abstracts enable the classification of the publications [67][5], we used the abstracts of the 211 publications to categorize them in terms of topics, contexts, nature, and scope (using Simon's classification). The systematic descriptive mapping has been employed to study the publications that will answer the research questions. We consider the research questions in turn, and for each question, the data analysis we have performed suggests specific topics that are under-researched.

RQ1: What are the reported topic areas in computer science education publications where HPC artifacts have been used as learning interventions?

The numbers of HPC artifacts categorized by topic were 7 in Ability/aptitude, 10 in Teaching/learning, 6 in Teaching/learning techniques, 72 in Teaching/learning tools, 103 in About research, and 13 in the curriculum. The major use of HPC artifacts is in research, closely followed by teaching as a pedagogical tool.

RQ2: What are the contexts in which computer science education uses HPC artifacts?

The numbers of HPC artifacts categorized by the context of publication are 59 in Parallel Hardware/architecture, 115 in Parallel Computing, 29 in Parallel Programming, and 8 in Parallel Processing. According to [67][5], the context dimension represents the thrust or sort of subject the publication is focusing

on. In [5], the subject has been defined as the unit of teaching where a student can achieve a formal result. In line with the rest of this study, we added the subject of parallel computing to the list of subjects of context dimension instead of using only the category of hardware/architecture in Simon's list. We found that the major use of HPC artifacts is in the context of parallel computing, which is followed by parallel architectures and then parallel programming.

RQ3: What is the scope in the usage of HPC artifacts as learning interventions supporting different curricula?

The numbers of HPC artifacts categorized by scope have been reported as follows: 1 in programme or department, 191 in Institution, 16 in Many institutions (and 3 in Not applicable). Most publications indicate that the greater usage of HPC artifacts is within a single institution (90.5%).

RQ4: What is the nature of the publications that have used HPC artifacts as learning interventions supporting different curricula?

The numbers of HPC artifacts categorized by nature of publication are 5 in position, 163 in the report, 33 in analysis, and 10 in the experiment.

Most of the publications (77.3%) indicate that the major usage of HPC artifacts has been presented as a pedagogical report, but very few publications report experimental results relating to the usefulness of HPC artifacts as pedagogical tools in computer science curricula.

RQ5: What parallel computing artifacts have been used as learning interventions in computer science education?

The analysis 211 publications have shown that 130 publications mention the usage of Beowulf clusters and other families of the cluster as teaching/learning tools in computer science and interventions to support research, compared to other types of HPC systems. However, we discovered from the literature that the others reported types of HPC that have received little attention include: cloud computing systems, cyberinfrastructure, grid computing, high throughput computing systems, massively parallel computers, multi-core systems, multiprocessor systems, a parallel computing system, supercomputing systems, web servers, and servlets.

We also found that from 1988 to 2018, there were 110 reviewed publications that focused on using specifically Beowulf clusters, 32 on using generic HPC artifacts and 22 using specifically HPC clusters. The other types of HPC artifacts formed a minority of those reported.

4 Discussion

The findings from this study show that HPC artifacts have been used as learning interventions to improve parallel computing knowledge and skills in the context of computer science education, as learning/teaching tools, and to support research [65] in the

context of computer science education. In addition, the findings show those families of HPC clusters [55][31][32] have been mentioned to be used extensively as learning interventions to improve parallel computing knowledge and skills.

The topic dimension classifies the papers in terms of ability/aptitude, teaching/learning, teaching/learning techniques, teaching/learning tools, about research, and the curriculum [67]. The significant proportions of publications in the dimension of topics have shown major usage of HPC artifacts as teaching/learning tools [45] and in supporting research. We also observed the usage of HPC artifacts to support teaching/learning parallelism techniques for freshmen [18] and the support of the development of a parallel computing curriculum framework [27].

The context dimension represents the generic thrust [67] of the published publications, which focuses on parallel hardware/architecture, parallel computing, parallel programming, and parallel processing categories. The literature analysis has shown that majority usage of the HPC artifacts has been reported in parallel hardware/architecture and parallel computing papers. We observed the usage of HPC artifacts for training students in parallel programming [73][19][23]. However, less has been reported on the usage of HPC artifacts for parallel processing in an education context, although Beowulf clusters have been predominantly used in the parallel hardware/architecture, parallel computing, and parallel programming categories. We have discovered that cloud computing systems [12][54] that use an HPC system have been reported to be used in the categories of hardware/architecture and parallel programming. Data also suggested that there should be further research on:

- parallel programming using different programming languages using HPC artifacts;
- parallel algorithms and processing using HPC artifacts.

The scope dimension represents the focus of the papers in terms of engagement with programme or department, and single or multiple institutions [67]. The literature has reported that a majority of the usage of HPC artifacts has been in the scope of single or multiple institutions, where it is reported that the predominant use of HPC artifacts is in experiments, analyses, and report publications. However, only one publication has been reported in the category of programme or department in the scope dimension. In addition, data suggested that there should be further research in order to:

- promote inter-institution engagement in HPC education research;
- encourage departments to engage in using HPC artifacts in programmes.

The nature dimension represents the distinction between practice and research paper [67] and focuses on position, analysis, report, and experiment papers [67]. The majority of the literature has been in reports and analysis papers. However, little has been

reported on the usage of HPC *machines* in reports and experiment papers. On top of that, data suggested further research in:

- designing position paper proposals that address ideas about the implementation of HPC artifacts in education;
- addressing research questions by analyzing data collected from researches that use HPC artifacts in education;
- conducting experiments and surveys about HPC artifacts in education that are used as tools to address research questions.

However, the analysis of the usefulness of HPC artifacts from the 211 publications based on topics has shown that ability/aptitude, teaching/learning, teaching/learning techniques, and curriculum categories have lower numbers of publications than those that have reported the usage of HPC artifacts. According to MacDonald and Atwood, the usefulness of a system can be defined as *“the extent to which a system’s functions allow users to complete a set of tasks and fulfill specific goals in a particular context of use”* [15]. Based on that definition, the tasks are training processes, and the goal is to impart HPC knowledge and skills. We discovered that a majority of publications that have been categorized using Simon’s classification reported the usage of Beowulf clusters [56] and other families of clusters as the pedagogical tools that support computer science education and research. We also observed some publications reporting usage of cloud [8] and virtual HPC clusters [60][24] as learning tools in the research and teaching/learning tool categories. The extensive use of Beowulf clusters as parallel computing artifacts in the context of HPC education research in parallel programming has been reported by Gardner [74]. Furthermore, Beowulf clusters have been used to support research experiments on job scheduling experiments [29]. We have also observed the usage of affordable and portable micro clusters to address the challenge of teaching in the context of parallel and distributed computing education [39]. Micro clusters have also been used in education to support low power and data-intensive applications [68][38].

As with regards to the usefulness of HPC artifacts in the 211 publications, the data suggests that there should be further research on:

- creating novel teaching or learning techniques using HPC artifacts;
- designing HPC courses that will increase the ability of students and the abilities that define what makes an HPC student good;
- the processes involved in teaching and learning HPC;
- HPC course contents and integration of HPC into curricula.

5 Conclusion

In this study, the usage of HPC artifacts as learning interventions to improve parallel computing knowledge and skills in the context of computer science education was investigated. We used SLR methodology to review 211 computer science publications in order to understand the trends of usage of HPC artifacts stratified in terms of types, topics, contexts, nature, and scope of Simon's classifications.

From the publications reviewed regarding whether HPC artifacts were used as the learning interventions, we observed that HPC artifacts were viable computing solutions used in the context of computer science education. The viability of the artifact was supported by the use of a large number of different names of types of HPC artifacts in the reviewed publications: Beowulf cluster, cloud computing system, cluster computing system, cyberinfrastructure, grid computing, HPC cluster, HPC system, HPC visualization cluster, Hadoop cluster, a high-throughput computing system, Linux cluster, massively parallel computer, multi-core system, multiprocessor system, parallel computers, parallel cluster, a parallel computing system, parallel machines, supercomputing system, virtual cluster and web servers with servlets.

However, we gained insights on the inadequacy of research in some categories of dimensions of topics, contexts, nature, and scope of Simon's classification. We discovered little attention to the application of HPC artifacts in ability/aptitude, teaching/learning, teaching/learning techniques, curriculum, parallel programming, parallel processing, programme or department, report, and experiment categories.

Our SLR consulted publications databases that indexed HPC education publications from the known journal and conference, but the limitations of this approach are that it did not include non-academic and non-peer-reviewed journal. We used different search strings in order to retrieve relevant publications that explicitly address HPC education; however, it is possible that other relevant publications could be retrieved using other search strings. In addition, we searched only English publication sources that limit our study from including non-English publications.

Further education research needs to be pursued to report on the usage of cloud, and virtual HPC artifacts as these two artifacts are both easy to deploy and affordable.

In the future, we are going to evaluate the deployment of affordable HPC clusters and evaluate their effectiveness in supporting the curriculum.

REFERENCES

- [1] Ahmad A. Datti, Hadiza A. Umar, and Jamil Galadanci. 2015. A Beowulf Cluster for Teaching and Learning. *Procedia Computer Science*; 70:62-8.
- [2] Ahmad Alturki, Guy G. Gable, and Wasana Bandara. 2011. A Design Science Research Roadmap. DESRIST. Milwaukee: Springer.
- [3] Alison Nightingale. 2009. A guide to systematic literature reviews. *Surgery (Oxford)*;27(9):381-4.
- [4] Amit Sabne, Putt Sakdhnagool, and Rudolf Eigenmann. 2015. HeteroDooP: A MapReduce programming system for accelerator clusters. In *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing*, pp. 235-246, ACM.
- [5] Angela Carbone, Michael de Raadt, Raymond Lister, Margaret Hamilton, and Judy Sheard. 2008. Classifying computing education papers: process and results. In *Proceedings of the Fourth international Workshop on Computing Education Research*, pp. 161-172. ACM.
- [6] Angela Harden and David Gough. 2012. Quality and Relevance Appraisal. In: D. Gough, S. Oliver, and J. Thomas (Eds.), *An Introduction to Systematic Reviews* (pp. 153-178). London: SAGE Publications, Inc.
- [7] Arlene Fink. 2013. *Conducting research literature reviews: From the Internet to paper*. Sage Publications.
- [8] Atanas Radenski. 2012. Integrating data-intensive cloud computing with multi-cores and clusters in an HPC course. In *Proceedings of the 17th ACM annual conference on Innovation and technology in computer science education*, pp. 69-74, ACM.
- [9] Barbara Kitchenham, and Stuart Charters. 2007. Guidelines for performing systematic literature reviews in software engineering. In *Technical report, Ver. 2.3 EBSE Technical Report*, Keele University, UK
- [10] Barbara Kitchenham. 2004. Procedures for performing systematic reviews. Keele, UK, Keele University, 33, 1-26.
- [11] Barbara O. Kitchenham, Pearl Brereton, David Budgen, Mark Turner, John Bailey, and Stephen Linkman. 2009. Systematic literature reviews in software engineering—a systematic literature review. *Information and software technology*;51(1):7-15.
- [12] Ceraj Ivica, Justin T. Riley, and Charles Shubert. 2009. StarHPC—Teaching parallel programming within elastic compute cloud. In *Information Technology Interfaces. ITI'09. Proceedings of the ITI 2009 31st International Conference on Jun 22 (2009)* pp. 353-356, IEEE.
- [13] Charles Peck. 2010. LittleFe: parallel and distributed computing education on the move. *Journal of Computing Sciences in Colleges*;26(1):16-22.
- [14] Claes Wohlin and Rafael Prikladnik. 2013. Systematic literature reviews in software engineering. *Information and Software Technology*;55(6):919-20.
- [15] Craig M. MacDonald and Michael E. Atwood. 2014. What does it mean for a system to be useful?: An exploratory study of usefulness. In *Proceedings of the Conference on Designing interactive systems Jun 21 (2014)* .pp. 885-894, ACM.
- [16] Daniel A Reed. 2014. Beowulf Clusters: From Research Curiosity to Exascale. In *Proceedings of the 20 Years of Beowulf Workshop on Honor of Thomas Sterling's 65th Birthday*.pp. 28-33, ACM.
- [17] David Gough and James Thomas. 2012. Commonality and diversity in reviews. In: Gough, D., & Thomas, J. (Eds.), *An introduction to systematic reviews* (pp. 35-65). London:Sage Publications Ltd.
- [18] Donald Johnson, David Kotz, and Fillia Makedon. 1994. Teaching parallel computing to freshmen. In *Conference on Parallel Computing for Undergraduates*. Colgate University.
- [19] Eduard Ayguadé and Daniel Jiménez-González. 2018. An approach to task-based parallel programming for undergraduate students. *Journal of Parallel and Distributed Computing*, 118, 140-156.
- [20] Eweoya Ibukun, and Olawande Daramola.2015. A Systematic Literature Review of Mobile Cloud Computing. *International Journal of Multimedia and Ubiquitous Engineering*;10(12):135-52.
- [21] Floyd B Hanson, 1990. A real introduction to supercomputing: A user training course. In *Proceedings of the 1990 ACM/IEEE conference on supercomputing*, .pp. 376-385, IEEE Computer Society Press.
- [22] Floyd Hanson, Thomas Moher, Nora Sabelli, and Ann Solem. 1988. A training program for scientific supercomputing users. In *Proceedings*

- of the ACM/IEEE conference on Supercomputing 1988 Nov 1 (pp. 342-349). IEEE Computer Society Press.
- [23] Francesco Gregoretti and Zary Segall. 1986. Programming for observability support in a parallel programming environment. In *Proceedings of the 1986 ACM fourteenth annual conference on Computer science* (pp. 189-197). ACM.
- [24] Frank Doelitzscher, Markus Held, Christoph Reich, and Anthony Sulistio. 2011. Viteraas: Virtual cluster as a service. In *Cloud Computing Technology and Science (CloudCom)*, IEEE Third International Conference on Nov 29 (2011) .pp. 652-657., IEEE.
- [25] Fred L. Kitchens, and Sushil K. Sharma. 2003. Affordable Supercomputing Solutions: Cluster Computers: In *Business Applications*. InPACIS, p. 126.
- [26] Geoffrey C. Fox. 1992. Parallel computing in industry {an initial survey. Technical Report SCCS-302b, Northeast Parallel Architectures Center.
- [27] Greg Wolffe and Christian Trefftz. 2009. Teaching parallel computing: new possibilities. *Journal of Computing Sciences in Colleges*; 25(1):21-8.
- [28] Han Xiaoying, Chen Feng, and Chen Wei. 2011. A rocks based visualization cluster platform design and application for virtual reality education. In *IT in Medicine and Education (ITME)*, International Symposium on Dec 9 (2011) ,Vol. 1, pp. 735-739, IEEE.
- [29] Hélio Silva, Anders Lyhne Christensen, and Sancho Oliveira. 2011. Performance study of Conillon: a platform for distributed computing. In *Proceedings of the 2011 Workshop on Open Source and Design of Communication* , (pp. 13-18). ACM.
- [30] Hong Lin and Khoi Nguyen. 2005. Classroom simulation of massive parallel computing. In *Proc. The 16th IASTED International Conference on Modeling and Simulation*, pp. 18-20.
- [31] Hong Lin, Ongard Sirisaengtaksin, and Ping Chen. 2010. A cluster computing environment at a small institution to support faculty/student projects. *Journal of Computing Sciences in Colleges*; 25(4):30-6.
- [32] Hong Lin. 2013. Teaching parallel and distributed computing using a cluster computing portal. In *Parallel and Distributed Processing Symposium Workshops & PhD Forum (IPDPSW)*, 2013 IEEE 27th International, pp. 1312-1317, IEEE.
- [33] James R. Fischer. 2014. The Roots of Beowulf. In *Proceedings of the 20 Years of Beowulf Workshop on Honor of Thomas Sterling's 65th Birthday*, pp. 1-6, ACM.
- [34] Jarmila Skrinarova, Adam Dudas, and Eduard Vesel. 2017. Model of education and training strategy for the management of HPC systems. In *Informatics, IEEE 14th International Scientific Conference on 2017 Nov 14* (pp. 400-405). IEEE.
- [35] Jeff Brunton and James Thomas. 2012. Information Management in Reviews. In: D. Gough, S. Oliver, & J. Thomas (Eds.), *An Introduction to Systematic Reviews* , pp. 83 – 106, London: Sage Publications Ltd
- [36] Jie Li, Judy F. Burnham, Trey Lemley, and Robert M. Britton. 2010. Citation analysis: Comparison of web of science®, scopus™, SciFinder®, and google scholar. *Journal of electronic resources in medical libraries*; 7(3):196-217.
- [37] Joan Ernst Van Aken and Georges Romme. 2009. Reinventing the Future: Adding Design Science to the Repertoire of Organization and Management Studies. *Organization Management Journal*, 6, 5-12.
- [38] Joel C. Adams and Tim H. Brom. 2008. Microwulf: a beowulf cluster for every desk. *ACM SIGCSE Bulletin*; 40(1):121-5.
- [39] Joel C. Adams, Jacob Caswell, Suzanne J. Matthews, Charles Peck, Elizabeth Shoop, David Toth, and James Wolfer. 2016. The Micro-Cluster Showcase: 7 Inexpensive Beowulf Clusters for Teaching PDC. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education*, pp. 82-83, ACM.
- [40] José A. González-Martínez, Miguel L. Bote-Lorenzo, Eduardo Gómez-Sánchez, and Rafael Cano-Parra. 2015. Cloud computing and education: A state-of-the-art survey. *Computers & Education*; 80:132-51.
- [41] Joseph G. Walls, George R. Widmeyer, and Omar A. El Sawy. 1992. Building an Information System Design Theory for Vigilant EIS. *Information System Research*, 3,36-60.
- [42] Julian P. Higgins and Sally Green. 2006. Cochrane handbook for systematic reviews of interventions 4.2.6. In: Higgins, J. P., & Green, S.(Eds). *The cochrane library*. Chichester, UK:Wiley.
- [43] Julian PT Higgins and Douglas G. Altman. 2008. Assessing risk of bias in included studies. *Cochrane handbook for systematic reviews of interventions: Cochrane book series*, 187-241.
- [44] Keke Gai, and Saier Li. 2012. Towards cloud computing: a literature review on cloud computing and its development trends. In *Multimedia Information Networking and Security (MINES)*, Fourth International Conference on Nov 2 (2012) .pp. 142-146, IEEE.
- [45] Kian L Pokorny. 2005. Constructing a Beowulf cluster as a pedagogical tool. *Journal of Computing Sciences in Colleges*; 20(4):104-5.
- [46] Kiduk Yang and Lokman I. Meho. 2006. Citation analysis: a comparison of Google Scholar, Scopus, and Web of Science. *Proceedings of the American Society for information science and technology*; 43(1):1-5.
- [47] Kurt D. Bollacker, Steve Lawrence, and C. Lee Giles. 2000. Discovering relevant scientific literature on the web. *IEEE Intelligent Systems and their Applications*; 15(2):42-7.
- [48] Lluc Alvarez, Eduard Ayguade, and Filippo Mantovani. 2018. Teaching HPC Systems and Parallel Programming with Small-Scale Clusters. In *2018 IEEE/ACM Workshop on Education for High-Performance Computing (EduHPC)* (pp. 1-10). IEEE.
- [49] Michael Levine-Clark and Esther L. Gil. 2008. A comparative citation analysis of Web of Science, Scopus, and Google Scholar. *Journal of Business & Finance Librarianship*; 14(1):32-46.
- [50] Moher, David, Alessandro Liberati, Jennifer Tetzlaff, Douglas G. Altman and Doug Altman. 2009. Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. *Annals of internal medicine*, 151(4), 264-269.
- [51] Nicholas Aaron Robison and Thomas J. Hacker. 2012. Comparison of vm deployment methods for hpc education. In *Proceedings of the 1st Annual conference on research in information technology*, (pp. 43-48). ACM.
- [52] Nikola Rajovic, Alejandro Rico, Filippo Mantovani, Daniel Ruiz, Josep Oriol Vilarrubi, Constantino Gomez, Luna Backes, Diego Nieto, Harald Servat, Xavier Martorell and Jesus Labarta. 2016. The Mont-Blanc prototype: an alternative approach for HPC systems. In *High Performance Computing, Networking, Storage and Analysis, SC16: International Conference for 2016 Nov 13* (pp. 444-455). IEEE.
- [53] Norbert Attig, Paul Gibbon, and Th Lippert. 2011. Trends in supercomputing: The European path to exascale. *Computer Physics Communications*, 182(9), 2041-2046.
- [54] Pekka Abrahamsson, Sven Helmer, Nattakarn Phaphoom, Lorenzo Nicolodi, Nick Preda, Lorenzo Miori, Matteo Angriman, Juha Rikkila, Xiaofeng Wang, Karim Hamily, and Sara Bugoloni. 2013. Affordable and energy-efficient cloud computing clusters: the bolzano raspberry pi cloud cluster experiment. In *Cloud Computing Technology and Science (CloudCom)*, IEEE 5th International Conference on Dec 2 (2013) ,Vol. 2, pp. 170-175, IEEE.
- [55] Per Andersen. 1999. The Texas Tech tornado cluster: a Linux/MPI cluster for parallel programming education and research. *Crossroads*; 6(1):28-32.
- [56] Philip R Prins. 2004. Teaching parallel computing using Beowulf clusters: a laboratory approach. *Journal of Computing Sciences in Colleges*; 20(2):55-61.
- [57] Philipp Mayr and Anne-Kathrin Walter. 2007. An exploratory study of Google Scholar. *Online information review*; 31(6):814-30.
- [58] R. L. Warrender, John Tindle, and David Nelson. 2013. Development of a virtual cluster. In *High Performance Computing and Simulation (HPCS)*, International Conference on Jul 1 (2013) .pp. 545-551, IEEE.
- [59] Rashid Mehmood, Furqan Alam, Nasser N. Albogami, Iyad Katib, Aiiad Albesbri, and Saleh M. Altowajiri. 2017. UTiLearn: A

- personalised ubiquitous teaching and learning system for smart societies. *IEEE Access*;5:2615-35.
- [60] Richard Brown, and Elizabeth Shoop. 2013. Teaching undergraduates using local virtual clusters. In *IEEE International Conference on Cluster Computing (CLUSTER)*, pages 1--8. IEEE.
- [61] Richard Thorpe, Robin Holt, Allan Macpherson, and Luke Pittaway. 2005. Using knowledge within small and medium-sized firms: a systematic review of the evidence. *International Journal of Management Reviews*, 7(4), 257-281.
- [62] Ricoux P, Berthou JY, Bidot T. 2013. European exascale software initiative EESI2—towards exascale roadmap implementation. In *2nd IS-ENES workshop on high-performance computing for climate models*.
- [63] Ronald B. Griffin, Janusz S. Kowalik, and Melvin R. Scott. 1988. Supercomputer integration program at Boeing. In *Supercomputing'88*, [Vol. 1], Proceedings, pp. 442-449, IEEE.
- [64] Sarah M. Diesburg, Paul A. Gray, and David Joiner. 2005. High performance computing environments without the fuss: the bootable cluster CD. In *Parallel and Distributed Processing Symposium*, 2005. Proceedings. 19th IEEE International , (pp. 8-pp). IEEE.
- [65] Semir Sarajlic, Naranjan Edirisinghe, Yuriy Lukinov, Michael Walters, Brock Davis, and Gregori Faroux. 2016. Orion: Discovery Environment for HPC Research and Bridging XSEDE Resources. In *Proceedings of the XSEDE16 Conference on Diversity, Big Data, and Science at Scale*, p. 54, ACM.
- [66] Shirley Gregor and David Jones. 2007. The Anatomy of a Design Theory. *Journal of the Association for Information Systems*, 8(5), 312-335.
- [67] Simon. 2007. A Classification of Recent Australasian Computing Education Publications', *Computer Science Education*, 17:3, 155 – 169
- [68] Simon J. Cox, James T. Cox, Richard P. Boardman, Steven J. Johnston, Mark Scott, and Neil S. O'brien. 2014. Iridis-pi: a low-cost, compact demonstration cluster. *Cluster Computing*;17(2):349-58.
- [69] Staffs Keele. 2007. Guidelines for performing systematic literature reviews in software engineering. In *Technical report*, Ver. 2.3 EBSE Technical Report. EBSE, sn.
- [70] Stephanie S. Goodwin and John C. Vukelich. 1988. From user to client services; making the transition for supercomputing. In *Proceedings of the 16th annual ACM SIGUCCS Conference on User Services* , pp. 111-115, ACM.
- [71] Susan Mehringer. 1996. Education in high performance computing via the WWW: Designing and using technical materials effectively. In *Proceedings of the 1996 ACM/IEEE conference on supercomputing*, p. 9, IEEE Computer Society.
- [72] Valerie Smith, Declan Devane, Cecily M. Begley, and Mike Clarke. 2011. Methodology in conducting a systematic review of systematic reviews of healthcare interventions. *BMC medical research methodology*, 11(1), 15.
- [73] William B. Gardner and John D. Carter. 2014. Using the pilot library to teach message-passing programming. In *Education for High Performance Computing (EduHPC), Workshop on Nov 16 (2014)* ,pp. 1-8, IEEE.
- [74] William B. Gardner. 2017. Should We Be Teaching Parallel Programming?. In *Proceedings of the 22nd Western Canadian Conference on Computing Education* ,p. 3, ACM.
- [75] Zhe Fan, Feng Qiu, Arie Kaufman, and Suzanne Yoakum-Stover. 2004. GPU cluster for high performance computing. In *Supercomputing*, 2004. Proceedings of the ACM/IEEE SC2004 Conference, pp. 47-47, IEEE.
- [76] Zoran Constantinescu, and Pavel Petrovic. 2002. Q 2 ADPZ* an open source, multi-platform system for distributed computing. *Crossroads*,9(2):13-20.