

**CLOUD COMPUTING AND EDA (ELECTRONIC
DESIGN AUTOMATION): COST OPTIMIZATION FOR
THE EDA JOBS WITH OBJECT STORAGE**

Ramesh Fernando

219333G

Degree of Master of Science

Department of Computer Science and Engineering

University of Moratuwa

Sri Lanka

March 2024

**CLOUD COMPUTING AND EDA (ELECTRONIC
DESIGN AUTOMATION): COST OPTIMIZATION FOR
THE EDA JOBS WITH OBJECT STORAGE**

Ramesh Fernando

219333G

MSc thesis submitted in partial fulfillment of the requirement for the degree Master
of Science

Department of Computer Science and Engineering

University of Moratuwa

Sri Lanka

March 2024

Declaration

I declare that this is my own work and this report does not incorporate without acknowledgement any material previously submitted for a Degree or Diploma in any other University or institute of higher learning and to the best of my knowledge and belief it does not contain any material previously published or written by another person except where the acknowledgment is made in the text. Also, I hereby grant to University of Moratuwa the non-exclusive right to reproduce and distribute my thesis/dissertation, in whole or in part in print, electronic or other medium. I retain the right to use this content in whole or part in future works.

Signature:

Date: 2024.04.27

The above candidate has carried out research for the MSc research thesis under my supervision. I confirm that the declaration made above by the student is true and correct.

Name of the supervisor: Prof. Indika Perera

Signature of the supervisor:

Date:

Name of the supervisor: Dr. Chathuranga Hettiarachchi

Signature of the supervisor:

Date:

Acknowledgment

I would like to express my profound gratitude to my advisors Prof. Indika Perera, Dr. Chathuranga Hettiarachchi, Dr. Gayashan Amarasinghe and Dr. Ming-Ying Chung (Synopsys Inc.) for their invaluable support, encouragement, supervision and guidance throughout the research journey. The continuous motivation of Dr. Chathuranga Hettiarachchi encouraged me to complete my thesis successfully.

I sincerely express my gratitude to my former colleagues at Synopsys Sri Lanka for their support and guidance in making this work a reality. I'm thankful to Dr. Ranga Rodrigo and Mr. Damith (technical officer) from the Electronics and Telecommunications Department, University of Moratuwa who enabled access to the licensed software setup to run my experiments and helped me to troubleshoot the connectivity issues.

I want to extend my heartfelt gratitude to the JuiceFS community. Through their extensive resources and active participation, I have gained a comprehensive understanding of distributed file storage.

Also, I'm truly grateful to my parents and my wife who always encouraged me to complete the task I started.

Abstract

Cloud computing with its ability to provide on-demand and optimized computing resources is emerging to disrupt the High-Performance Computing (HPC) markets. On the semiconductor domain, the Electronic Design Automation (EDA) has been an essential prerequisite, enabling the design of complex integrated circuits. With the increasing complexity of System on the Chip (SoC) designs, coupled with the need for faster time-to-market has led to a growing demand for scalable cost-effective EDA solutions. The emergence of cloud computing has provided a potential solution to these challenges, offering the ability to scale EDA workflows on-demand and pay-per-usage concept. The thesis explores the avenues in using JuiceFS, the open source, high-performance, POSIX compatible, distributed file system designed for the native cloud applications to deploy EDA workloads in the cloud with the aid of object storage by minimizing the cost for storage. Furthermore, characteristics of EDA workloads are identified. The cost for storage (eg. High-performance Network File System (NFS), AWS FSx for Lustre and Azure NetApp) is a major contributor to the overall cost of the cloud and has been a limiting factor for most of the semiconductor companies, especially for the startups. The storage cost component can be minimized by enabling the file system with cheaper object storage. The tradeoff is the performance downgrade (Run time of the EDA tool) and it is observed that high-performance POSIX-compatible file system is required to minimize the tradeoff. An experiment on VCS* compile time is conducted to identify the tradeoffs using JuiceFS configured with the metadata server and object server.

*Collaborating with Synopsys Inc., the VCS slowdown in the JuiceFS setup is analyzed with the Proof of Concept (POC) developed by Synopsys Inc. on top of JuiceFS.

Keywords: EDA, Cloud Computing, JuiceFS, POSIX, Metadata, Object Storage, HPC, Distributed Computing.

Table of Contents

Declaration	i
Abstract	iii
List of Figures	vi
List of Tables.....	viii
List of Abbreviations.....	ix
Chapter 1. Introduction	1
1.1 Background of Study	1
1.2 Research Problem.....	3
1.3 Research Objectives	4
1.4 Structure of the Thesis	5
Chapter 2. Literature Review	6
2.1 Evolution of Semiconductor Chips	6
2.2 Evolution of the EDA industry with the increase of chip complexity	7
2.3 Cloud Computing and EDA	8
2.4 Barriers to a paradigm shift (To move the EDA workloads to the cloud) .	10
2.4.1 Cost of the Cloud	10
2.4.2 Security of the Cloud	12
2.4.3 Licensing model by the EDA vendors	13
2.5 Expected changes necessary to migrate the EDA workloads from on- premise to the cloud	14
2.6 Object Storage, architecture, pros, and cons	15
2.6.1 In-memory object storage	16
2.6.2 File-based object storage.....	17
2.7 Network File Storage (NFS)	17
2.8 FUSE – File System in Userspace	18
2.9 JuiceFS File System	21

2.11	Related work	32
2.11.1	Synopsys and IBM Research	32
2.11.2	Ceph	32
Chapter 3.	Methodology	35
3.1	Setting up the file system with the aid of JuiceFS and the object storage	35
3.2	Deploying VC Formal demo cases.	38
3.3	Cost optimization with the object storage vs performance	39
3.4	VCS® Compilation.....	40
3.5	Experimental Setup	42
3.6	VCS® Compilation with Customer Designs	43
Chapter 4.	Observations, Results and Analysis	45
4.1	The runtime and the tracing results for the VC Formal® setup (JuiceFS distributed mode).....	45
4.2	VCS® on NFS (Tier 1 storage) and the JuiceFS setup.....	45
4.3	Cost comparison.....	47
4.4	VCS® Compilation slowdown analysis in the JuiceFS setup	51
Chapter 5.	Discussion and Conclusion	56
5.1	Discussion	56
5.2	Study Limitations	57
5.3	Future Work	59
5.4	Conclusion	59
Reference List	61

List of Figures

Figure 1.1: Architecture of the JuiceFS file system [7]	2
Figure 1.2: Top-level architecture of POC developed on top of JuiceFS	3
Figure 2.1: Simplified view of a modern Integrated Circuit(IC) design flow [9].....	8
Figure 2.2: The graph of cost vs beign late to the market [11]	10
Figure 2.3: Overview of block storage, file storage and object storage.....	16
Figure 2.4: Comparison of the traditional object storage model vs OSD mo. [25] ...	17
Figure 2.5: The client-server architecture of the Network File System.....	18
Figure 2.6: High-level architecture of FUSE [28]	19
Figure 2.7: Profiled report of VCS Compilation with JuiceFS – MinSoC design [30]	21
Figure 2.8: Overview of file storage mechanism in JuiceFS [7].....	23
Figure 2.9: Object chunks are mapped to the metadata in JuiceFS [7].....	24
Figure 2.10: Data flow of JuiceFS write request [7].....	25
Figure 2.11: Data flow of JuiceFS read request [7]	26
Figure 2.12: Modes of caching at different stages in JuiceFS [7].....	28
Figure 2.13: Reference EDA workflow in the cloud [29].....	29
Figure 2.14: The Ceph System Architecture.....	33
Figure 2.15: A Software Defined Infrastructure for EDA jobs by IBM [38]	33
Figure 3.1: The experimental setup for the JuiceFS file system	36
Figure 3.2: The no. of CPUs available in the aedc2 server machine	36
Figure 3.3: The MinIO object server	37
Figure 3.4: A mount point of the JuiceFS file system	37
Figure 3.5: Debugging an assertion failure using VC Formal GUI.....	38
Figure 3.6: Representation of VCS parsing and analysis stage	40
Figure 3.7: Representation of VCS elaboration	41
Figure 4.1: The comparison of data in the table 7.....	46

Figure 4.2: Price quotation of AWS object storage (S3 bucket).....	48
Figure 4.3: The AWS bill for 50TB of object storage for 1 month [35].....	48
Figure 4.4: Price quotation of Amazon FSx for Lustre.....	49
Figure 4.5: The AWS bill for the file storage in FSx for Lustre – Scratch.....	49
Figure 4.6: The AWS bill for the file storage in FSx for Lustre – Persistent	50
Figure 4.5: Call Tree of the collect report running on NFS(L*) and JuiceFS file system(R*)	51
Figure 4.6: Statistics of the collect report running on NFS(L*) and JuiceFS file system(R*)	51
Figure 4.7: Timeline of the collect report running on NFS(L*) and JuiceFS file system(R*)	52
Figure 4.8: JuiceFS profiling data for the customer4_perform_uvm BM with ‘-- writeback’ option enabled	53
Figure 4.9: Snapshot of the .accesslog during the compilation of customer4_perform_uvm BM in the JuiceFS mount point.....	54
Figure 4.10: The I/O latency mapped for the benchmark customer4_perform_uvm during the period of compilation.....	54
Figure 4.11: The overall analysis and summary of the experiment	55
Figure 5.1: The control flow and the data flow of an EDA application.....	58

List of Tables

TABLE 1: Nature of the EDA workloads in different stages of the design flow	9
TABLE 2: The cost comparison for Microsoft Azure Blob vs Azure NetApp Files. 11	
TABLE 3: FUSE request types by group [28]	19
TABLE 4: AMAZON Web services resources vs Synopsys on-premise server	31
TABLE 5: Microsoft AZURE machine resources vs Synopsys on-premise server ..	31
TABLE 6: Cloud Storage solutions that support POSIX operations.	39
TABLE 7: Detailed information of the customer benchmarks	44
TABLE 8: CPU Time and the memory usage during the VC Formal run.....	45
TABLE 9: VCS compile time and the simulation time on NFS and JuiceFS setup ..	46
TABLE 10: Cost comparison for object storage (S3 bucket) vs Amazon FSx for Lustre.....	50
TABLE 11: Voluntary and Involuntary context switching of NFS vs JuiceFS file system.....	52

List of Abbreviations

ASIC	Application Specific Integrated Circuits
AWS	Amazon Web Services
BM	Benchmarks
EDA	Electronic Design Automation
FUSE	Filesystem in Userspace
HPC	High Performance Computing
HTTP	Hypertext Transfer Protocol
NFS	Network File Storage
NDA	Non-Disclosure Agreement
OSD	Object-based Storage Device
POC	Proof of Concept
RTL	Register Transfer Level
SoC	System on Chip
TLS	Transport Layer Security
VCS	Verilog Compiler and Simulator
VLSI	Very Large-Scale Integrations