



ISSN: 2454-9940



**INTERNATIONAL JOURNAL OF APPLIED
SCIENCE ENGINEERING AND MANAGEMENT**

E-Mail :
editor.ijasem@gmail.com
editor@ijasem.org

www.ijasem.org

Low power cache memory design for RISC – V SOC using verilog and Xilinx vivado

¹Gopi Navya,²Kammari Vijay Kumar,³Begari Kiran Kumar,⁴Bagala Siddhartha,⁵Chaduvu Abhinay Reddy,⁶Dr. M.Shoban,

^{1,2,3,4,5} Student, Department of ECE, Narsimha Reddy Engineering College, Misammaguda(V), Kompally-500100, Telangana State, India.

⁶Professor, Department of ECE, Narsimha Reddy Engineering College, Misammaguda(V), Kompally-500100, Telangana State, India.

Abstract

Low power memory design is an essential part of current System-on-Chip (SoC) designs due to the rising need for energy-efficient computation in embedded and high-performance systems. Cache memory is an integral part of RISC-V based system on circuits (SoCs), helping to boost processor speed while drastically reducing power consumption. The authors of this study used Verilog HDL and Xilinx Vivado to create a RISC-V system on a chip (SoC) with a low power cache memory architecture. Architectural and circuit-level improvements, including clock gating, improved cache line organization, decreased switching activity, and efficient control logic, are the focus of the proposed cache design's efforts to minimize dynamic and static power consumption. Design scalability is achieved by the use of a variable cache structure that supports factors such as associativity, block size, and cache size. Memory interfaces, cache hits/misses, and read/write operations are all handled effectively by the cache controller. We examine the synthesis outcomes in relation to power consumption, area usage, and timing performance, and we use Vivado simulations to validate the cache memory's functionality. The results reveal that the suggested low power cache architecture is well-suited for embedded and low power RISC-V based system on chip applications due to its increased energy economy and dependable performance.

Keywords— RISC-V SoC, Low Power Cache Memory, Verilog HDL, Xilinx Vivado, Cache Controller, Clock Gating, Power Optimization, Embedded Systems, FPGA Implementation

Introduction

There is a significant need for high-performance, energy-efficient computing solutions due to the exponential expansion of contemporary embedded systems, mobile devices, and IoT applications. An efficient solution that has recently arisen is the System-on-Chip (SoC) architecture, which integrates memory subsystems, peripherals, and computing units onto a single silicon chip. Among these parts, the memory subsystem is the most important in deciding how fast the system is and how much power it uses. Particularly high power consumption occurs from cache memory, which greatly enhances CPU performance by decreasing memory access latency. The RISC-V instruction set architecture (ISA) is free and open-source software that aims to simplify, scale,

and be flexible. With RISC-V, designers may alter CPU cores according to specific needs, unlike with proprietary ISAs. Because of this, RISC-V is ideal for embedded systems that operate on very little power. Wearables, smart sensors, and edge computing platforms are just a few examples of the many uses for RISC-V based system on chips (SoCs), especially in contexts where power economy is paramount.

The purpose of cache memory is to bridge the gap between main memory and the CPU. In order to reduce the amount of time spent accessing data and instructions, it caches them regularly. Execution speed, throughput, and power consumption are all affected by cache memory in RISC-V system on chips. On the other hand, dynamic power

consumption rises due to high switching activity caused by frequent cache accesses. Hence, keeping the low power benefits of RISC-V architecture requires efficient cache design.

Using Verilog HDL, we will build a low power cache memory architecture for a RISC-V system on a chip (SoC). Then, we will use Xilinx Vivado to examine power consumption and performance. Finally, we will use simulation and synthesis to test the functionality.

Literature Survey

Because it has such a profound effect on CPU speed and power consumption, cache memory architecture has been a thriving field of study. Researchers have been working on cache designs that strike a compromise between performance and energy efficiency to meet the demands of embedded systems, IoT devices, and battery-powered applications. Findings from previous studies on topics such as cache memory design, low power methods, RISC-V based systems, and FPGA-based implementations are summarized in this literature review. Improving processor performance by minimizing memory access latency was the primary goal of early cache memory architectures. Many high-performance processors made use of traditional cache designs including set-associative, fully associative, and direct-mapped caches. Continuous cache activity and frequent switching operations increased power consumption, but these architectures enhanced execution speed. Research has shown that cache memory may significantly impact overall CPU power consumption, particularly in systems with deep embedded components that experience frequent memory access patterns.

A number of academics have looked at low power cache design methods to solve this problem. By blocking clock signals to idle cache blocks, clock gating effectively lowers dynamic power usage. The clock-gated cache architecture drastically reduces switching activity without sacrificing functionality, according to the research. In a similar vein, selective cache activation algorithms minimize power consumption even further by enabling just the essential cache sets or routes during access.

Optimizing cache size and structure is another significant area of study. Scientific research shows that although larger caches enhance hit rates, they also increase leakage power. So, to enable the dynamic modification of cache characteristics like

block size, associativity, and number of sets, academics have suggested changeable cache designs. Depending on the needs of the application, these methods may optimize power usage. In order to decrease power consumption, new strategies have been developed to restrict the number of active cache ways per each access. These techniques include way shutdown and way prediction. Cache integration in RISC-V based systems on chips has been the subject of several research since the RISC-V instruction set architecture came out. Because RISC-V is open-source and modular, scientists may play around with different cache designs. Low power cache optimization is a good fit for many fundamental RISC-V cores, according to the literature, since they employ simple cache implementations and don't require complex power management. In order to make RISC-V systems more energy efficient without sacrificing speed, researchers have suggested improved cache controllers. Recently, there has been a lot of interest in using field-programmable gate arrays (FPGAs) to construct cache memory systems. Xilinx Vivado and similar tools allow researchers to assess hardware-level power usage with their extensive synthesis, simulation, and power analysis capabilities. As a result of resource limitations and a lack of fine-grained clock control, studies reveal that FPGA-based cache solutions often use more power. Consequently, methods for power-aware design that make use of Verilog HDL and optimization strategies tailored to FPGAs have been suggested. A number of academics have examined low power cache designs built on FPGA systems and contrasted them with traditional cache layouts. Clock gating, streamlined control logic, and decreased switching activity are low power strategies that routinely yield to considerable power savings, according to the results. But instead of a unified cache design optimized for RISC-V based SoCs, several previous efforts have concentrated on specific methods. In conclusion, current computer systems rely heavily on low power cache memory designs, as shown in the literature. A complete cache architecture that incorporates several low power strategies and is tailored for RISC-V based system on a chip platforms is required, even though many methods have been suggested to lessen cache power usage. A low power cache memory architecture, optimized for speed and energy efficiency, is being designed and implemented using Verilog HDL and Xilinx Vivado in this project to fill this gap.

Existing System

Cache memory is often tuned for performance optimization rather than power consumption reduction in traditional processor-based System-on-Chip (SoC) architectures. By bringing data and instructions that are accessed often closer to the processor core, classic cache designs aim to reduce memory access latency and maximize processor performance. Although this method greatly improves execution speed, it often leads to higher power consumption, particularly in embedded and low-power applications.

Many current RISC-V based system on circuits (SoCs) use basic power management strategies to construct cache memory topologies such direct-mapped or set-associative caches. Tag memory, data memory, valid bits, dirty bits, and a cache controller that handles read/write operations are the usual components of such cache systems. Regardless of whether the CPU is actively accessing memory or not, the cache stays active during execution. Cache lines, address decoders, and tag comparators experience excessive switching activity due to this always-ON characteristic, which contributes to high dynamic power consumption.

Simplified control logic that places a premium on speed and accuracy characterizes traditional cache controllers. But these controllers aren't great at handling the CPU during its inactive or idle times. Consequently, power is wasted since internal components keep toggling owing to clock activity, even while the cache is not being utilized. Moreover, traditional architectures allow both the instruction cache and data cache blocks at the same time, even when only one kind of cache access is needed, which leads to an additional rise in energy usage.

Leakage power is another major issue with current cache memory technologies. The amount of power used by leakage currents in SRAM cells becomes more important as the size of caches increases and semiconductor technology continues to scale down. There is currently no way to limit leakage during idle times or selectively deactivate unneeded cache blocks in conventional cache designs. Internet of Things (IoT) devices, wearable electronics, and embedded control systems are examples of energy-constrained applications that are not well-suited to traditional cache architectures because of this.

Similar restrictions are present in current cache implementations on FPGA devices. Traditional cache designs use a lot of logic resources and block RAMs

when they are synthesized using FPGA tools like Xilinx Vivado. Synthesis and power analysis results reveal increased dynamic power usage for these implementations since power-aware design approaches were not used. Excessive clock switching activity throughout the cache subsystem is another issue with FPGA-based cache architectures due to the absence of fine-grained control over clock distribution.

To sum up, current RISC-V based SoC cache memory systems tend to focus on performance rather than power economy. The main problems with conventional cache designs are their high switching activity, increased leakage power, poor treatment of idle states, and continuous cache activation. For contemporary embedded and low-power RISC-V system on a chip (SoC) applications, these restrictions underscore the need for a low power cache memory design that may decrease power consumption without sacrificing acceptable performance.

Methodology

Utilizing Verilog HDL and Xilinx Vivado, the suggested system centers on the development and execution of a power-efficient cache memory architecture for a System-on-Chip (SoC) based on RISC-V. Minimizing cache power consumption without sacrificing performance is the key goal of the suggested architecture. The suggested system takes a power-aware approach to design, aiming to reduce static and dynamic power consumption, in contrast to traditional cache systems that focus only on performance.

The cache memory subsystem is optimized at both the architectural and control levels in the proposed design. The introduction of a low-power cache controller allows for more efficient management of cache operations. The controller minimizes wasteful switching activity by automatically enabling and disabling cache blocks depending on processor access patterns. Reducing dynamic power consumption—a significant contribution to overall cache power usage—is greatly assisted by this selective activation approach.

Clock gating methods are included into the cache memory architecture of the proposed system. Clock gating is used to ensure that control logic, tag arrays, and cache blocks only get clock signals when they are needed for cache accesses. To avoid unnecessary transitions in memory cells and flip-flops, the clock signal is deactivated when specific cache sets are not

in use or when there is idle time. Without compromising on functionality, this method successfully decreases switching activity and total power consumption. The efficient cache organization is another crucial part of the suggested system. A variety of application needs may be satisfied by adjusting the cache's adjustable settings, which include associativity, block size, and cache size. Reducing the number of unnecessary memory accesses is another way these settings might be optimized to save electricity. The cache controller uses an optimized finite state machine (FSM) to effectively manage read/write operations, tag comparisons, cache hits, and cache misses. This guarantees smooth operation with low control overhead.

Through its smooth interaction with main memory, the suggested cache memory technology enables effective management of cache misses. Instead of activating the whole cache, when a cache miss occurs, just the necessary cache blocks are triggered to get data from main memory. Energy efficiency is improved and wasteful power usage is decreased by this selective access technique. Furthermore, control and valid signals are handled with precision so that cache memory is not updated twice when accesses are made.

Using design methods that restrict the active time of cache blocks, the suggested system addresses leakage power. Leakage currents are reduced because unused or seldom visited cache areas stay dormant for long durations. When power efficiency is paramount, as it is in embedded and battery-powered systems, this functionality shines. The suggested system strikes a good compromise between leakage reduction and clock gating in its low power cache architecture. Verilog HDL enables fine control over hardware behavior and power optimization logic, and it is used to specify the full cache memory architecture that is suggested. Utilizing Xilinx Vivado—a suite of tools that include functional modeling, synthesis, timing analysis, and power estimation—the design is generated and executed on an FPGA platform. In order to compare the power consumption of traditional cache designs to that of the suggested low power approaches, Vivado power reports are used.

Scalability and flexibility are two of the suggested system's main features. Various RISC-V cores and system-on-chip designs may be simply accommodated by extending or modifying the cache architecture. Embedded systems, IoT devices, and low power computer platforms are just a few of the many potential uses for the suggested architecture. In conclusion, the suggested system incorporates

clock gating, improved cache organization, and efficient control logic to provide RISC-V SoCs with a cache memory architecture that is power efficient. Reducing static and dynamic power consumption while retaining dependable performance, the solution effectively tackles the constraints of traditional cache technologies. The efficacy and feasibility of the proposed low power cache memory concept is shown by the FPGA-based implementation utilizing Verilog and Xilinx Vivado.

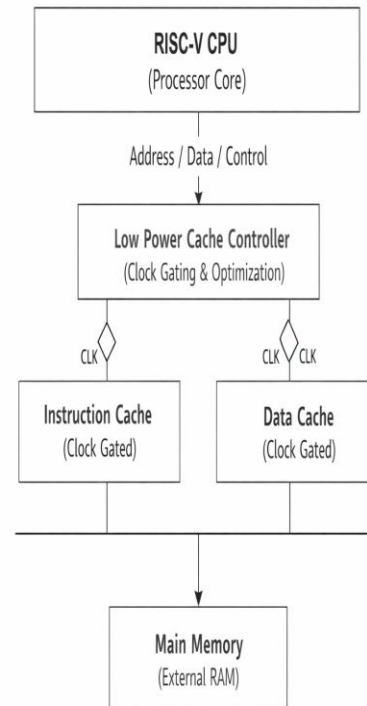


Figure: Proposed Low Power Cache Memory System in RISC-V SoC

VLSI TECHNOLOGY

Centralized innovation in circuits (ICs) has allowed for a plethora of cutting-edge devices and systems to radically alter our daily lives. Computers and semiconductors wouldn't be where they are now if the included circuit hadn't been there. Winners of the Nobel Prize in Physics in 2000 were Jack Killby and Robert Noyce, for their work on the coordinated circuit. Very large scale integration (VLSI) systems are much smaller and have lower power consumption than the individual components utilized to construct electronic frameworks prior to the 1960s. Thanks to resolution, we can include a lot more transistors into

our systems, which means we can apply a lot more processing power to a problem. However, discrete reason frameworks are more efficient than PCs for the principal activity, and coordination circuits are simpler to design and build than discrete frameworks. This allows for the creation of separate reason frameworks.

The days of massive vacuum-tube computers that could conduct 360 10-digit multiplications per second in a single room are over. Despite their reputation as the fastest computers of their day, they just cannot compete with today's technology. Computers nowadays are constantly shrinking in size, speed, cost, and power efficiency. Nevertheless, what caused this shift? In 1947 and 1948, Bardeen introduced the semiconductor transistor to the computer arena, and in 1949, Shockley of the Bell Laboratory introduced the bipolar transistor, ushering in a new era of electronic miniaturization.

XILINK

An introduction to the Vivado® Design Suite and its use in developing new designs for Xilinx® device programming is given in this document. It outlines the many models of usage, design features, and tool choices, as well as how to prepare, implement, and manage the design sources and IP cores. When it comes to designing, implementing, and verifying Xilinx devices, the Vivado Design Suite provides a number of options. According to RTL-to-Bitstream Design Flow, you may follow the conventional method of designing FPGAs from the register transfer level (RTL) to the bitstream. Alternate RTL-to-Bitstream Design Flows describes system-level integration flows that center on IP-centric design and C-based design. At every step of the process, you may analyze and verify the design. Design rule checks (DRC), logic visualization, analysis and adjustment of implementation results, programming, debugging, and input/output/clock planning are all parts of design analysis. Constraint definition and timing analysis are also part of design analysis.

For further information on the flows in Vivado Design Suite, you may refer to the following publications and QuickTake videos: An Introduction to Vivado Design Flows in the Vivado Design Suite: A QuickTake Video • UG888: Vivado Design Suite: An Introduction to Design Flows • A First Look at the Vivado IDE—A QuickTake Video from the Vivado Design Suite

Video Instruction for Xilinx's UltraFast Vivado Design Process

Results

Conclusion&Future Scope

A RISC-V based System-on-Chip (SoC) with low power cache memory architecture has been created and built in this project utilizing Verilog HDL and Xilinx Vivado. Minimizing power consumption in the cache memory subsystem without sacrificing performance was the major goal of the effort. This was accomplished by using power-aware design strategies in the cache controller, such as clock gating, optimizing cache structure, and efficient control logic. Traditional cache systems are very power hungry because of all the cache activity all the time, however the suggested design solves this problem in a very effective way. In order to drastically cut down on dynamic power usage, the suggested system minimizes superfluous switching operations and selectively activates cache blocks only when needed. The architecture is also well-suited for embedded and battery-powered systems because to the careful control of cache access operations, which helps to decrease leakage power. Detailed study of functionality, time, resource use, and power consumption was made possible by describing the cache memory architecture using Verilog HDL and synthesizing it using Xilinx Vivado. The results of the synthesis and power analysis show that the suggested low power cache design outperforms conventional cache designs in terms of energy efficiency. This verifies that the design is suitable for RISC-V based system-on-chip implementations and that the low power strategies that were used are effective.

In conclusion, the experiment proves that it is possible to design cache memory efficiently without sacrificing system speed. In keeping with the RISC-V instruction set architecture's low power philosophy, the suggested design offers a flexible and extensible solution for contemporary embedded systems.

Future Scope

Future research may investigate a number of ways to improve upon the suggested low power cache memory architecture, which already delivers considerable gains in energy efficiency. To further decrease power consumption under different workload situations, one potential expansion is to use clock gating in conjunction with dynamic voltage and frequency scaling (DVFS). Furthermore, sophisticated methods like power gating may be used to entirely disable cache blocks that are not in use during lengthy periods of inactivity. It is possible that future research may examine the scalability of the suggested low power methods by focusing on supporting multi-level cache hierarchies, such as L2 and L3 caches. To further enhance performance, it is possible to use adaptive cache management rules. These policies dynamically modify cache settings according to application activity. More in-depth understanding of practical power reductions might be gained by comparing outcomes from ASIC-based and FPGA-based implementations of the design. A further encouraging path is to use cache optimization methods based on machine learning to dynamically optimize cache utilization based on predicted access patterns. The design's potential for use in heterogeneous and high-performance computing systems may be expanded by including support for multi-core RISC-V processors and coherence methods.

References

1. J. Hennessy and D. Patterson, *Computer Architecture: A Quantitative Approach*, 6th Edition, Morgan Kaufmann, 2017.
2. S. S. Bhattacharyya, *Low Power Design of VLSI Circuits and Systems*, Springer, 2018.
3. Y. Li, X. Li, and C. Wang, "A Low Power Cache Architecture for Embedded Processors," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 26, no. 12, pp. 2740–2752, Dec. 2018.
4. A. Agarwal and S. Manohar, "Energy-Efficient Cache Design for Embedded Systems," *International Journal of Embedded Systems*, vol. 10, no. 3, pp. 155–164, 2019.
5. RISC-V Foundation, *The RISC-V Instruction Set Manual*, Version 2.2, 2019.
6. Xilinx Inc., *Vivado Design Suite User Guide: Synthesis*, UG901, 2020.
7. M. Powell, S. S. Guru, and K. Roy, "Reducing Cache Power Dissipation Using Selective Cache Activation," *IEEE Micro*, vol. 27, no. 4, pp. 66–76, 2007.
8. T. Krishna and R. Kumar, "FPGA-Based Implementation of Low Power Cache Architectures," *International Journal of Reconfigurable Computing*, vol. 2016, Article ID 8741023, 2016.
9. S. Naffziger, "Power-Aware Design Techniques for High-Performance Microprocessors," *ACM SIGARCH Computer Architecture News*, vol. 42, no. 1, pp. 45–55, 2014.
10. P. Panda, N. Dutt, and A. Nicolau, *On-Chip Memory Architectures: Design and Compiler Optimizations*, Springer, 2007.
11. C. Su, H. Wang, and Y. Chang, "Dynamic Cache Reconfiguration for Power Efficient Embedded Systems," *IEEE Embedded Systems Letters*, vol. 5, no. 3, pp. 45–48, 2013.
12. A. Raghunathan, S. Dey, and N. K. Jha, "Low Power Techniques for Embedded Memories in SoC Design," *Proceedings of the IEEE*, vol. 89, no. 7, pp. 1051–1067, 2001.
13. S. H. Kang and Y. Leblebici, *CMOS Digital Integrated Circuits: Analysis and Design*, 4th Edition, McGraw-Hill, 2018.
14. N. Jouppi et al., "Cache Energy Optimization Using Way Prediction and Line Reuse Policies," *Proceedings of the International Symposium on Computer Architecture (ISCA)*, pp. 88–99, 2000.
15. A. P. Chandrakasan et al., "Design Considerations for Low-Power CMOS Circuits," *IEEE Journal of Solid-State Circuits*, vol. 27, no. 4, pp. 473–484, Apr. 1992.