



# Minimizing Memory Contention in an APNG Encoder Using a Grid of Processing Cells

Vivek Govindasamy<sup>(✉)</sup>, Emad Arasteh, and Rainer Dömer

CECS, University of California, Irvine, USA  
{vbgovind,emalekza,doemer}@uci.edu  
<https://www.cecs.uci.edu/>

**Abstract.** Modern processors experience memory contention when the speed of their computational units exceeds the rate at which data can be accessed in memory. This phenomenon is well known as the memory bottleneck and is a great challenge in computer engineering. In order to mitigate the memory bottleneck in classic multi-core architectures, a scalable parallel computing platform called Grid of Processing Cells (GPC) has been proposed. To evaluate its effectiveness, we model the GPC using SystemC TLM-2.0, with a focus on memory contention. As an example, we parallelize an APNG encoder application and map it to the GPC and compare its performance to traditional shared memory processors. Our experimental results show improved execution times on the GPC due to a large decrease in memory contention.

**Keywords:** Memory Bottleneck · Grid of Processing Cells · SystemC TLM-2.0

## 1 Introduction

The increase in processor speeds over the past years has led to increased time spent in accessing the main memory to retrieve data. As many cores try to access the shared memory, this leads to contention and delays each core. The cores suffer from contention and their computations are halted due to sharing of the same main memory. This memory bottleneck applies to most modern CPUs which are usually shared memory processors (SMP).

To deal with slow memory access speeds, various solutions are being researched. The development of hierarchical caches is the main method to address this issue [1]. Another solution is the Berkeley RISC project, in which many complex instructions were removed because they were rarely used [2], and instead replaced with more CPU registers which are much faster to access than main memory [3, 4].

In this paper, we model and evaluate SMPs and a scalable alternative called *Grid of Processing Cells (GPC)* where processors paired with local memories are

arranged in a 2D array [5]. As a specific configuration, the GPC checkerboard architecture (Fig. 1) is aimed at addressing the memory bottleneck. The cores and memories are placed one after another, and each core has access to its own and three neighbour memories, thereby increasing data availability. The checkerboard contains several variations of a main logical component which is termed as a cell. Each cell is designed with the idea that it represents a core and components that are local to that particular core.

In this work we model the GPC architecture in SystemC TLM-2.0 [6] and map an application to it [7].

### 1.1 Problem Definition

Our main contribution in this paper is modeling and demonstrating the improvement of the GPC against the classic SMP architectures in terms of execution time and time spent in main memory access contention when running an APNG encoder on the different architectures.

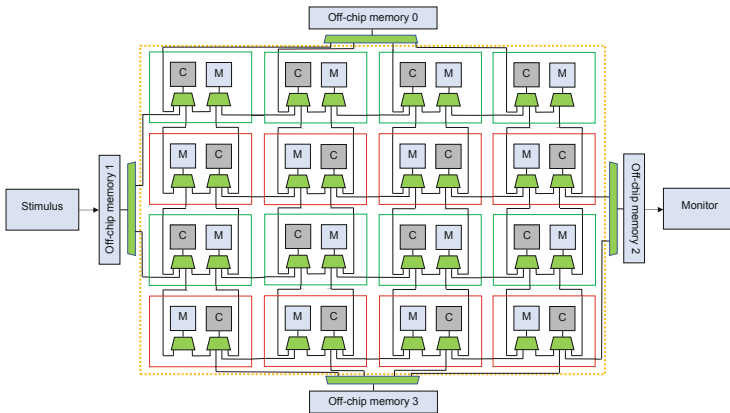


Fig. 1. 4-by-4 checkerboard GPC [5].

### 1.2 Background and Related Work

Over many years the general trend is that processors become faster as designers increase the clock rate but there is little increase in memory access speed [8]. This is known as the memory wall. To mitigate this problem, there has been a lot of focus on improving caches [1]. However, even the most advanced caches suffer from high miss rates if the cache size is too small or if associativity is increased too much [9]. Caches must also implement cache coherence protocols in the case of multi-core processors, and they consume significant space on the chip as well as power. Caches have given rise to Non-Uniform Memory Access (NUMA) where each core can access near memory faster than distant memory. With NUMA the time to maintain cache coherence is usually quite high [10]

and leads to contention as the interconnect is shared for every core. There have been other works related to addressing the memory wall, for example the Illusion system [11], which is similar to the GPC architecture. Both have a mapping step, but in the GPC the memories surrounding each core can be accessed only by neighbours and there is no NOC to facilitate communication between distant cores. This limitation makes the GPC truly scalable because there is no NOC complexity to grow, but puts a burden on application mapping which is restricted to only local communication.

Modification to the architecture itself is one method of reducing the contention between cores. In the checkerboard GPC contention is reduced drastically as separate buses are used between the cores. Architectures similar to the GPC have been proposed in the past, such as the Epiphany-V [12]. While similarities are present, such as the fact that both use a cache-less memory model, the GPC varies in a few aspects: i) The checkerboard architecture uses a different addressing map where each memory has a different address space, ii) The GPC has no operating system running on all of the cores, and iii) GPC uses simple multiplexers-based buses instead of a complex NOC.

Our main objective in this work is to reduce memory contention through software methods, which primarily lies in the mapping operation. We also confirm that memory contention is indeed a major reason for the memory wall, and provide experimental results showing that the GPC minimizes that.

## 2 Modeling of the Checkerboard GPC

The checkerboard model [5] consists of cores with local memory which can be accessed also by their neighbours. The memory size is small and the cores themselves only perform computation on small amounts of data at a time. The small memories are referred to as on-chip memories and are expected to be as fast as caches in a multi-core computer made of static random-access memory (SRAM). The off-chip memories are larger but slower, similar to dynamic random-access memory (DRAM).

**Core Module** - The core module is the computation component of a cell (Fig. 2) and represents a complete processor core with in-order or out-of-order execution but without (or only 1<sup>st</sup>-level) cache. It contains a single socket connected to the core multiplexer. It is a SystemC module containing general arithmetic functions. The primary communication is the SystemC blocking transport interface (*b.transport*) [13] and event synchronization to prevent possible race conditions when interacting with other cores. Each core contains a main thread which performs the actual computation.

**Memory Module** - Each core has its own on-chip memory to store its data. This memory is assumed to be fast as SRAM, but small (up to 128 megabytes). The memory can only be accessed by the four neighbouring cores similar to a local scratchpad memory with explicitly managed address space. The off-chip memories on the edges of the checkerboard are larger (512 megabytes) but have slower access (DRAM).

**Core Demultiplexer Module** - In order to communicate with the neighbours, an addressing scheme for the checkerboard is required. The global address space refers to the four off-chip memories on the outside and the local address space refers to the small memories near the cores. The core demultiplexer routes addresses to the individual memories. It contains one socket to communicate with the core and four sockets to connect to the adjacent memory multiplexers. The core demultiplexer forwards *b\_transport* calls from the core after performing address translation.

**Memory Multiplexer Module** - The memory multiplexer is connected to a memory and to adjacent core demultiplexers. Its purpose is to forward *b\_transport* calls from neighboring cores to its memory. The memory multiplexer permits only one access at a time and performs arbitration. This allows to observe memory contention. Algorithm 1 provides the algorithm of how the time spent waiting for memory access is computed [14].

---

**Algorithm 1:** Maintaining busy state in *b\_transport* inside the core multiplexer (bus access arbitration with FCFS policy)

---

```

initialization: busy_until = 0;
busy = busy_until - current timestamp;
if busy < 0 then
    | busy_until = current timestamp;
    | busy = 0;
end
delay = multiplexer delay + busy;
d1 = delay;
socket->b_transport(transaction, delay);
d2 = delay;
memory delay = d2 - d1;
busy_until += memory delay ;

```

---

## 2.1 Modeling Interconnect Contention

To observe the contention in the traffic of memory transactions, we utilize timing delays in the SystemC TLM-2.0 blocking transport interface inside our interconnect. To ensure that only one transaction at a time uses each memory, the memory multiplexer stores its busy status in a state variable and delays competing transactions accordingly.

As listed in Algorithm 1, we store a timestamp marking the end of memory occupation in a variable `busy_until` which we initialize to zero. When a transaction arrives, we calculate the remaining time left until the memory becomes available again. If `busy` is negative, the transaction arrived at an idle time and `busy_until` is reset.

Before forwarding the transaction to the memory (`b_transport`), we update the `delay` with the sum of the multiplexer latency and the busy delay. The transaction then processes in the memory. We observe the `memory_delay` by taking the time before and after the transaction and update the `busy_until` state variable accordingly.

In summary, our interconnect modeling accurately performs arbitration with first-come-first-serve policy and tracks the busy state of memory transactions. Aware of contention, our model enables accurate observation of any congestion in memory traffic.




Fig. 2. Checkerboard GPC Cell.

### 3 Parallelized APNG Encoder Application

To test the performance of the checkerboard GPC, we need a suitable application which can be run in parallel. We choose an Animated Portable Network Graphics (APNG) encoder [15] which basically is a PNG encoder that concatenates generated PNG images with additional information such as the frame rate. PNG encoders have two main components which perform the actual image compression, the filters and the DEFLATE algorithm [16]. Filters are of five types (None, Sub, Up, Avg, Paeth) and are used to reduce pixel values. The reduced pixel values require less number of bits to transmit, providing some compression. The filtered values are sent in to the DEFLATE algorithm, which uses a combination of Lempel-Ziv-Storer-Szymanski (LZSS) and Huffman encoding to perform lossless compression. DEFLATE works better on values which are highly correlated to each other, which filtering provides [17].

Our SystemC TLM-2.0 APNG encoder consists of eight modules, namely the Color Splitter, Subtract Filter, Up Filter, Average Filter, Paeth Filter, Comparator, Compressor and APNG Encoder. The Color Splitter separates input data into individual color streams coming from the Stimulus of the encoder. The filters perform different mathematical computations which correlate pixel data, improving the compression provided by DEFLATE. The Comparator chooses the best filtered output to send to the Compressor. The Compressor uses DEFLATE to output a compressed row which is sent to the APNG Encoder module and the Monitor which writes the compressed data to a file. The APNG Encoder module generates additional information needed to create an APNG file which

is written in the Monitor module. Our model performs the encoding row wise, with parallel filters.

### 3.1 Backannotation of Delays

In order to evaluate performance, we need to reflect timing in the model. We estimate the computation delay of the major APNG functions by measuring their execution time on a reference platform. Since we are mainly interested in the relative timing of major blocks in the application, we simply run the APNG encoder on a computer (2.4 GHz CPU i5-1135G7) and measure the delays with the `gprof` Linux profiler. The observed delays are listed in Table 1. We note that the filtering operations are most time-consuming in the encoder. Thus, we parallelize the filters in our model. We back-annotate the measured computation delays into the APNG SystemC model and scale them proportionally to the image size.

Every memory access by the cores results in a communication delay. In reality, not every memory access takes the same amount of time as some accesses will be to the cache and others to the main memory. However, we have not modeled caches in our SMP and single core models. Therefore, for fairness purposes we consider every memory access and multiplexer switch to be 10ns uniformly, regardless of on-chip or off-chip memory. This is still an effective measure of performance because the SMP models must perform main memory accesses frequently so that they can communicate the filtered rows to the other cores which use it, and caches are not of much use here.

Another important metric is the contention time, which is how much time each core spends waiting for access to the main memory, as described with Algorithm 1.

## 4 Mapping APNG on the Checkerboard

To evaluate the effectiveness of the checkerboard architecture we map the APNG encoder application on it. Since not all of the memories are available to each core, it is sometimes necessary to forward data through cores. Forwarding increases the communication time for each cell. To reduce the lost time incurred by forwarding, it is better to use data by an adjacent core.

### 4.1 Checkerboard Mappings

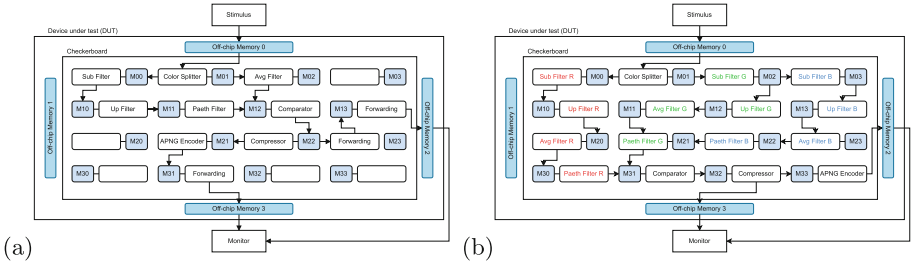
Two checkerboard mappings have been implemented, a simpler initial mapping where individual filtering on the three colors is performed on the same core, and

**Table 1.** APNG computation delays

Module Name	Total time	Time per frame	Time per pixel
Color Splitter	4 s	0.133 s	11 ns
Subtract Filter	30 s	1.000 s	82 ns
Up Filter	33 s	1.100 s	88 ns
Average Filter	50 s	1.667 s	137 ns
Paeth Filter	102 s	3.400 s	274 ns
Comparator	8 s	0.267 s	21 ns
Compressor	14 s	0.467s	38 ns
APNG Encoder	1 s	0.033s	3 ns

an improved mapping which splits the colors between cores to filter. In the initial mapping nine cores are involved in encoding with three cores forwarding data. The improved mapping has every core utilized in APNG encoding.

**Initial Checkerboard Mapping.** The main idea behind the initial mapping is to avoid the use of excessive forwarding (Fig. 3 (a)). For instance, the Subtract filter receives the red, green, and blue values from the Color Splitter. Next, it performs the computation and sends both the filtered values and the unfiltered values to the Up filter. The Up filter performs its own computation and sends the filtered Sub, Up and unfiltered values to the Paeth filter. The Paeth filter then computes the Paeth filtered output, and sends it to the Comparator while also forwarding both outputs from the Up filter and Sub filter. The same process is performed in the Average filter route. At the Comparator the least sum filtered row is chosen and sent to the Compressor. The compressed output is forwarded through neighboring cells to the right and written to the output file using the monitor. It is also sent to the APNG Encoder module for APNG encoding and forwards it to the second monitor thread.



**Fig. 3.** (a) An initial checkerboard mapping which attempts to minimize communication between modules. (b) An improved checkerboard mapping which splits the computational load more evenly.

**Improved Checkerboard Mapping.** While the initial mapping works, it could be further improved by splitting the filtering work of each core to three cores with each core filtering a different color component (Fig. 3 (b)). This is because some filters, such as the Paeth filter, take too much time to compute. Since there are sixteen cores available, it is possible to map every core to one module. This mapping is expected to be faster.

## 5 Experimental Results

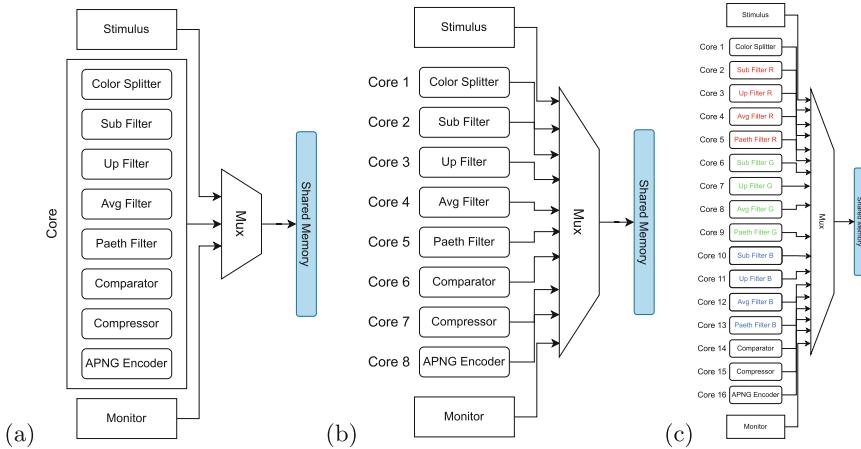
We now compare five different SystemC models. The models are 1) Initial checkerboard (Fig. 3 (a)), 2) Improved checkerboard (Fig. 3 (b)), 3) Single core

(Fig. 4 (a)), 4) 8 core SMP (Fig. 4 (b)), and 5) 16 core SMP (Fig. 4 (c)). These models are evaluated on the basis of their execution times and amount of contention.

## 5.1 Models for Comparison

For comparison, we implement a single core and also two SMP models in SystemC. These models use a single memory with a memory multiplexer connecting the cores to the memory. The single core model performs all computations in the same core (Fig. 4 (a)). The shared memory models work similar to the checkerboard architecture, using the same communication functions to transfer pixel data between cores, but have a greater amount of contention.

The SMP with 8 cores performs the filtering operation on the different color channels in the same core (Fig. 4 (b)). Doubling the number of cores provides a 16 core model in which the cores perform less work but communication is increased per unit time leading to higher contention (Fig. 4 (c)).



**Fig. 4.** Models for comparison, (a) single core architecture, (b) 8 core SMP where threads run in parallel, (c) 16 core SMP with shared memory.

## 5.2 Measurement of Delays

To measure delays, we create global variables of type *sc\_time* called *Computation.Time*, *Read.Time*, *Write.Time* and *Contention.Time*. These variables provide an accurate value of how the time in each SystemC model is spent, and are present in every model. The *Computation.Time* (*Comp.Time*) is the summation of time spent by every core on the times listed in Table 1.

*Read.Time* and *Write.Time* keep track of the total time spent by the model accessing the memory. Added together they reflect the *Memory-Access-Time* or the *MA.Time*.



Lastly, the *Contention-Time* (*Cont.Time*) is the total time spent by every core waiting to access the memory. This time can easily exceed the execution time of the model if there are a lot of cores attempting to access the memory at the same time, as each core will be waiting to access the memory and all of these times add up to the *Contention-Time*.

### 5.3 Simulated Time Results

With the three types of delays back-annotated, we obtain measurements on the five SystemC models. Table 2 shows the variation in model timings as the clock rate is gradually increased, thereby decreasing computation time per module and increasing the frequency of memory access requests. The five SystemC models

**Table 2.** Table for simulated timing results

Model Name	Exec Time (Speedup)	Comp Time	MA Time	Cont Time
<b>Clock rate of 0.25 GHz</b>				
Single Core	714.617 s (1x)	707.231 s	4.619 s	6.791 s
SMP with 8 Cores	326.703 s (1x)		18.518 s	27.130 s
SMP with 16 Cores	131.946 s (1x)		18.525 s	29.992 s
Initial Checkerboard	334.866 s (1x)		27.291 s	1.388 s
Improved Checkerboard	122.943 s (1x)		28.296 s	1.130 s
<b>Clock rate of 0.5 GHz</b>				
Single Core	361.818 s (1.98x)	353.615 s	4.619 s	6.791 s
SMP with 8 Cores	166.704 s (1.96x)		18.518 s	61.124 s
SMP with 16 Cores	81.454 s (1.62x)		18.525 s	64.398 s
Initial Checkerboard	174.859 s (1.92x)		27.291 s	0.699 s
Improved Checkerboard	70.008 s (1.76x)		28.296 s	2.676 s
<b>Clock rate of 1 GHz</b>				
Single Core	185.419 s (1.96x)	176.808 s	4.619 s	6.791 s
SMP with 8 Cores	86.703 s (1.92x)		18.518 s	78.675 s
SMP with 16 Cores	59.332 s (1.38x)		18.525 s	93.821 s
Initial Checkerboard	94.855 s (1.84x)		27.291 s	0.354 s
Improved Checkerboard	43.336 s (1.63x)		28.296 s	2.564 s
<b>Clock rate of 2 GHz</b>				
Single Core	97.219 s (1.91x)	88.404 s	4.619 s	6.791 s
sMP with 8 Cores	48.423 s (1.81x)		18.518 s	88.403 s
SMP with 16 Cores	45.426 s (1.31x)		18.525 s	101.670 s
Initial Checkerboard	54.853 s (1.73x)		27.291 s	0.182 s
Improved Checkerboard	30.001 s (1.43x)		28.296 s	2.712 s
<b>Clock rate of 4 GHz</b>				
Single Core	53.119 s (1.83x)	44.202 s	4.619 s	6.791 s
SMP with 8 Cores	38.048 s (1.26x)		18.518 s	90.867 s
SMP with 16 Cores	39.677 s (1.13x)		18.525 s	116.412 s
Initial Checkerboard	34.852 (1.57x)		27.291 s	0.096 s
Improved Checkerboard	23.332 (1.31x)		28.296 s	2.785 s
<b>Clock rate of 8 GHz</b>				
Single Core	31.069 s (1.71x)	22.101 s	4.619 s	6.791 s
SMP with 8 Cores	38.047 s (1.00x)		18.518 s	91.792 s
SMP with 16 Cores	38.308 s (1.03x)		18.528 s	119.683 s
Initial Checkerboard	24.852 s (1.41x)		27.291s	0.053 s
Improved Checkerboard	19.999 s (1.17x)		28.296 s	2.619 s

are compared at different assumed processor clock rates. The clock rates have been chosen to start at 0.25 GHz and are doubled up to 8 GHz in our simulation.

An important assumption made intentionally is that no cache memory exists in any model. The reasoning for this is that we want to observe memory contention directly, undisturbed by caching behaviour. In other words, for a fair comparison we model all memories as fast and thus do not need caches.

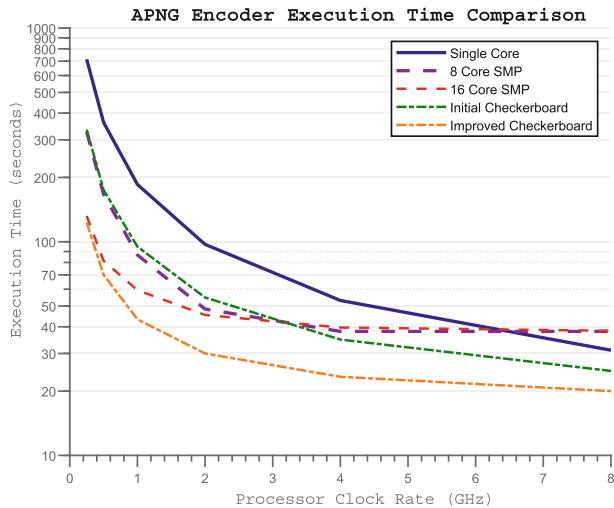
Table 2 shows the comparison of the three types of delays and the overall execution time. Time spent in computation linearly reduces, but has limited effect on the total execution time as the contention time goes up in every model, to varying degrees. The increase in contention time is the reason why execution times start to show less improvement.

Plotting the values from Table 2 provides insight on the change in execution time as the clock rate is doubled (Fig. 5). For the single core model, it is seen that the decrease in computation speed leads to great increase in speedup, until a certain point where diminishing returns are observed. The shared memory processors start off with low execution time, but they start to stagnate at around 4 GHz.

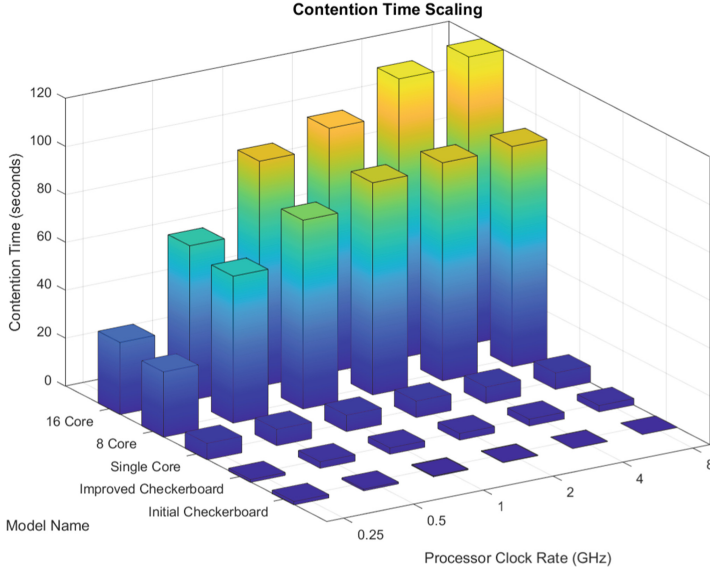
The *Memory.Access.Time* or *MA.Time* varies for each of these models even though the amount of pixels they process is the same. This is because simpler models, like the single core model, need to access the memory only two times (from the stimulus and to the monitor) when processing a row of pixel data. Other models, like the checkerboard, need to pass on the data between adjacent cores, which involves a lot more reading and writing from and to memories. The SMP models need to access the main memory frequently, as almost every core needs to access new data to continue data processing.

## 5.4 Observations and Comparison

The initial checkerboard mapping starts with an execution time similar to the 8 core SMP model, but quickly accelerates as the clock rate is increased. The improved checkerboard is as fast as the 16 core SMP at low clock rates, but is twice as fast when the clock rate reaches 8 GHz. At low clock rates, the limiting factor is the number of cores, and not the memory access contention, whereas at higher clock rates this trend is reversed.



**Fig. 5.** APNG Encoder execution time scaling. Note that the graph y axis is in log scale.



**Fig. 6.** Contention time scaling with increase in clock rate for the five SystemC models.

Figure 6 shows the increase in contention as the clock rate increases. This increases the rate at which memory is accessed by the cores so that they can process more data, which leads to a rise in contention time. For the single core model minimal contention exists because the cores are accessing the memory along with the stimulus and monitor. The initial checkerboard mapping has a linear decrease in contention, but the improved checkerboard suffers from a slight increase. The reasoning for the reduction in contention for the initial mapping is that when the clock rate increases the shorter memory accesses appear to come first, and this leads to a decrease in contention as shortest job first (SJF) reduces wait time. This does not seem to be the case for the SMP models however, as they have a noticeable increase in contention as the clock rate increases and the execution time starts to stagnate. Therefore the checkerboard architecture is a good alternative to shared memory processors as clock rate increases, its contention is much less.

## 6 Conclusion

The increase in processor speeds over the years has resulted in much faster computers but this trend has been hampered due to slower memory speed increases. The newly proposed checkerboard architecture is one possible way to mitigate the effects of slower memory access speeds, as shown by our experimental results.

In future work, we aim to provide more accurate comparisons by also including caches for our SMP models. Further, we would like to continue our modeling

by lowering the level of abstraction of our cores to instruction set simulators [18] and mapping more applications to the GPC.

In the longer term, we plan to address the programmability of the GPC architecture so that applications can be mapped to it automatically by an advanced compiler.

## References

1. Smith, A.J.: Cache memories. *ACM Comput. Surv. (CSUR)* **14**(3), 473–530 (1982)
2. Tanenbaum, A.S.: Implications of structured programming for machine architecture. *Commun. ACM* **21**(3), 237–246 (1978)
3. Patterson, D.A.: Reduced instruction set computers. *Commun. ACM* **28**(1), 8–21 (1985)
4. Cocke, J., Markstein, V.: The evolution of RISC technology at IBM. *IBM J. Res. Dev.* **34**(1), 4–11 (1990)
5. Dömer, R.: A Grid of Processing Cells (GPC) with Local Memories, Tech. Rep. CECS-TR-22-01 (Apr. 2022)
6. Grötter, T., Liao, S., Martin, G., Swan, S.: *System Design with SystemCTM*. Springer Science & Business Media (2007)
7. Govindasamy, V.: Mapping of an APNG Encoder to the Grid of Processing Cells Architecture, Tech. Rep. CECS-TR-22-02 (2022)
8. McKee, S.A.: Reflections on the memory wall. In: *Proceedings of the 1st conference on Computing frontiers*, p. 162 (2004)
9. Patterson, D.A., Hennessy, J.L.: *Computer organization and design ARM edition: the hardware software interface*. Morgan kaufmann (2016)
10. Blagodurov, S., Zhuravlev, S., Dashti, M., Fedorova, A.: A Case for NUMA-aware Contention Management on Multicore Systems. In: *2011 USENIX Annual Technical Conference (USENIX ATC 2011)* (2011)
11. Radway, R.M.: Illusion of large on-chip memory by networked computing chips for neural network inference. *Nat. Electron.* **4**(1), 71–80 (2021)
12. Olofsson, A.: Epiphany-v: A 1024 processor 64-bit risc system-on-chip, arXiv preprint [arXiv:1610.01832](https://arxiv.org/abs/1610.01832) (2016)
13. IEEE Standard for Standard SystemC Language Reference Manual, IEEE Std 1666–2011 (Revision of IEEE Std 1666–2005), pp. 1–638 (2012)
14. Arasteh, E.M., Dömer, R.: Improving parallelism in system level models by assessing PDES performance. In: *Forum on specification & Design Languages (FDL)*, vol. 2021, pp. 01–07. IEEE (2021)
15. Parmenter, S., Vukićević, V., Smith, A.: APNG Specification by Mozilla. [www.wiki.mozilla.org/APNG\\_Specification](https://www.wiki.mozilla.org/APNG_Specification). (Accessed 12 Aug 2021)
16. Deutsch, P., Gailly, J.-L.: Zlib compressed data format specification version 3.3, RFC, May, p. 1996. Tech, Rep (1990)
17. Sayood, K.: Chapter 9. Lossless compression handbook. Elsevier (2002)
18. Herdt, V., Große, D., Le, H.M., Drechsler, R.: Extensible and configurable RISC-V based virtual prototype. In: *Forum on Specification & Design Languages (FDL)*, vol. 2018, pp. 5–16. IEEE (2018)