

## Assignment 2: POSIX Programming & Benchmarking

# 1 Host Environment

For my host environment, I chose to run Ubuntu Server 24.04.2 LTS using a VirtualBox hypervisor. I chose this operating system as I have sufficient Linux experience to feel confident using an operating system with no graphical interface (as opposed to Ubuntu Desktop), and the absence of a GUI means a smaller ISO file, memory footprint, & CPU footprint. I chose Ubuntu specifically because it's a Linux system with which I have previous experience, and is well-documented with plenty of packages available to install if needed. Ubuntu also makes it easy to install the PREEMPT\_RT patches, which transform the standard Linux kernel into a fully preemptible, real-time kernel, which I felt was more suitable for this assignment, as the standard Linux kernel is not suitable for a hard real-time system due to its lack of preemption.

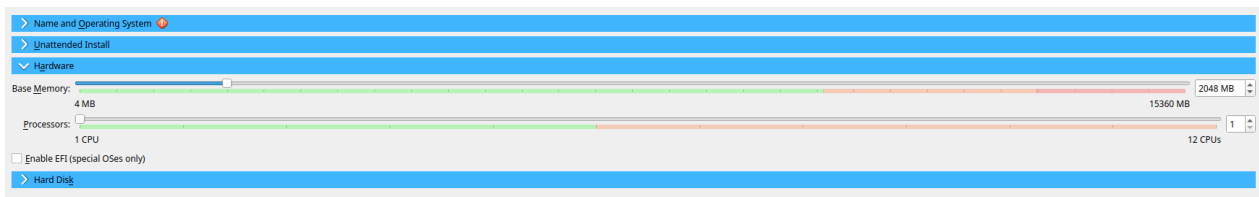


Figure 1: Virtual machine hardware configuration

I set the virtual machine to have a single CPU and set the amount of RAM to 2048MB which is the recommended minimum for Ubuntu Server<sup>1</sup>. I left the hard disk size at the default of 25GB as I saw no reason to change it. The real-time kernel with the PREEMPT\_RT patches installed is available with Ubuntu Pro, which is free for personal use. After setting up an Ubuntu Pro account, I enabled the real-time kernel using the pro command.

```
andrew-hayes@UbuntuRTS:~$ sudo pro enable realtime-kernel
One moment, checking your subscription first
Real-time kernel cannot be enabled with Livepatch.
Disable Livepatch and proceed to enable Real-time kernel? (y/N) y
No variant specified. To specify a variant, use the variant option.
Auto-selecting generic variant. Proceed? (y/N) y
Disabling incompatible service: Livepatch
Executing '/snap/bin/canonical-livepatch disable'
The Real-time kernel is an Ubuntu kernel with PREEMPT_RT patches integrated.

This will change your kernel. To revert to your original kernel, you will need
to make the change manually.

Do you want to continue? [ default = Yes ]: (Y/n) y
Configuring APT access to Real-time kernel
Updating Real-time kernel package lists
Updating standard Ubuntu package lists
Installing Real-time kernel packages
Real-time kernel enabled
A reboot is required to complete install.
andrew-hayes@UbuntuRTS:~$
```

Figure 2: Enabling the real-time kernel with the pro command

Finally, I transferred over the following C file (taken from the lecture slides) via scp<sup>3</sup> to the virtual machine to get the clock resolution, which is 1 nanosecond:

```
1 #include<unistd.h>
2 #include<time.h>
3 #include <stdio.h>
4
5 int main(){
6     struct timespec clock_res;
7     int stat;
8     stat=clock_getres(CLOCK_REALTIME, &clock_res);
9     printf("Clock resolution is %d seconds, %ld nanoseconds\n",clock_res.tv_sec,clock_res.tv_nsec);
10    return 0;
11 }
```

```

andrew-hayes@UbuntuRTS:~$ gcc -lrt res.c
res.c: In function 'main':
res.c:9:34: warning: format '%d' expects argument of type 'int', but argument 2 has type '__time_t' {aka 'long int'} [-Wformat=]
   9 |     printf("Clock resolution is %d seconds, %ld nanoseconds\n", clock_res.tv_sec, clock_res.tv_nsec);
     |                                ^~
     |                                |
     |                                int          __time_t {aka long int}
andrew-hayes@UbuntuRTS:~$ ./a.out
Clock resolution is 0 seconds, 1 nanoseconds
andrew-hayes@UbuntuRTS:~$

```

Figure 3: Getting the clock resolution of the virtual machine

## 2 Benchmarking Code

I combined the provided benchmarking programs `bm1.c` and `bm2.c` into a single file, and added logic to benchmark the `usleep()` function as well as outputting the relevant data to CSV files. Additionally, I updated the `#define ITERATIONS` constant to have value 10000 and I also tweaked the `while (!timer_expired)` loop to sleep for 100 nanoseconds in-between evaluations of the loop condition, as I found that the “busy waiting” was greatly slowing down the program when I ran it on my virtual machine. There is a potential drawback to this however: adding the `nanosleep()` to the while loop could artificially introduce a delay into obtaining the data, as there could be a maximum delay of 100 nanoseconds before the timer is registered as expired. However, since the busy wait was artificially increasing the runtime, and much more so than the version with the sleep, it too would introduce delay, and much more than the modified version, as the modified version ran around 10 times more quickly. Therefore, while this modification could potentially introduce noise to the data collected for the interval timer benchmark, it introduces less error than the busy wait, so I decided to include the modification.

```

1 // Compile code with gcc -o merged merged.c -lrt -Wall -O2
2 // Execute code with sudo ./merged
3
4 #include <stdio.h>           // Standard I/O functions
5 #include <stdlib.h>          // Standard library functions
6 #include <time.h>            // Time-related functions
7 #include <signal.h>          // Signal handling
8 #include <sys/mman.h>        // Memory locking
9 #include <unistd.h>          // POSIX standard functions
10 #include <sched.h>           // Scheduling policies
11 #include <errno.h>           // Error handling
12 #include <string.h>          // String manipulation
13 #include <limits.h>          // Limits of integral types
14
15 // Constants
16 #define ITERATIONS 10000    // Number of benchmark iterations
17 #define NS_PER_SEC 1000000000L // Nanoseconds per second
18
19 // Global Variables
20 timer_t timer_id; // Timer identifier
21 volatile sig_atomic_t timer_expired = 0; // Flag for timer expiration
22 volatile sig_atomic_t signal_received = 0; // Flag for signal reception
23 struct timespec start, end, sleep_time; // Time structures for benchmarking
24
25 // Function to save benchmark results to a CSV file
26 void save_results(const char *filename, long long *data) {
27     FILE *file = fopen(filename, "w");
28     if (!file) {
29         perror("fopen");
30         exit(EXIT_FAILURE);
31     }
32     fprintf(file, "Iteration,Latency/Jitter (ns)\n");
33     for (int i = 0; i < ITERATIONS; i++) {
34         fprintf(file, "%d,%lld\n", i, data[i]);
35     }
36     fclose(file);
37 }
38
39 // Signal handler for signal-based latency measurement

```

```

40 void signal_handler(int signum) {
41     signal_received = 1; // Mark signal as received
42     clock_gettime(CLOCK_MONOTONIC, &end); // Capture end time
43 }
44
45 // Timer signal handler
46 void timer_handler(int signum) {
47     timer_expired = 1; // Mark timer as expired
48     clock_gettime(CLOCK_MONOTONIC, &end); // Capture end time
49 }
50
51 // Configures real-time scheduling with FIFO priority
52 void configure_realtime_scheduling() {
53     struct sched_param param;
54     param.sched_priority = sched_get_priority_max(SCHED_FIFO);
55     if (sched_setscheduler(0, SCHED_FIFO, &param) == -1) {
56         perror("sched_setscheduler");
57         exit(EXIT_FAILURE);
58     }
59 }
60
61 // Locks memory to prevent paging for real-time performance
62 void lock_memory() {
63     if (mlockall(MCL_CURRENT | MCL_FUTURE) == -1) {
64         perror("mlockall");
65         exit(EXIT_FAILURE);
66     }
67 }
68
69 // Measures jitter of nanosleep function
70 void benchmark_nanosleep() {
71     long long jitter_data[ITERATIONS];
72     sleep_time.tv_sec = 0;
73     sleep_time.tv_nsec = 1000000; // 1 ms sleep
74
75     for (int i = 0; i < ITERATIONS; i++) {
76         clock_gettime(CLOCK_MONOTONIC, &start);
77         nanosleep(&sleep_time, NULL);
78         clock_gettime(CLOCK_MONOTONIC, &end);
79
80         jitter_data[i] = ((end.tv_sec - start.tv_sec) * NS_PER_SEC + (end.tv_nsec - start.tv_nsec)) -
            ↪ sleep_time.tv_nsec;
81     }
82     save_results("nanosleep.csv", jitter_data);
83 }
84
85 // Measures latency of sending and handling a signal
86 void benchmark_signal_latency() {
87     long long latency_data[ITERATIONS];
88     signal(SIGUSR1, signal_handler); // Register signal handler
89
90     for (int i = 0; i < ITERATIONS; i++) {
91         clock_gettime(CLOCK_MONOTONIC, &start);
92         kill(getpid(), SIGUSR1); // Send signal to itself
93         while (!signal_received); // Wait for signal to be handled
94
95         latency_data[i] = (end.tv_sec - start.tv_sec) * NS_PER_SEC + (end.tv_nsec - start.tv_nsec);
96         signal_received = 0;
97     }
98     save_results("signal_latency.csv", latency_data);
99 }

```

```

100
101 // Measures jitter of a real-time timer
102 void benchmark_timer() {
103     long long jitter_data[ITERATIONS];
104     struct sigevent sev;
105     sev.sigev_notify = SIGEV_SIGNAL;
106     sev.sigev_signo = SIGRTMIN;
107     sev.sigev_value.sival_ptr = &timer_id;
108
109     if (timer_create(CLOCK_MONOTONIC, &sev, &timer_id) == -1) {
110         perror("timer_create");
111         exit(EXIT_FAILURE);
112     }
113
114     struct itimerspec its;
115     its.it_value.tv_sec = 0;
116     its.it_value.tv_nsec = 1000000; // 1 ms
117     its.it_interval = its.it_value;
118     signal(SIGRTMIN, timer_handler);
119
120     if (timer_settime(timer_id, 0, &its, NULL) == -1) {
121         perror("timer_settime");
122         exit(EXIT_FAILURE);
123     }
124     clock_gettime(CLOCK_MONOTONIC, &start);
125     for (int i = 0; i < ITERATIONS; i++) {
126         while (!timer_expired) {
127             struct timespec ts = {0, 100};
128             nanosleep(&ts, NULL);
129         }
130
131         clock_gettime(CLOCK_MONOTONIC, &end);
132         jitter_data[i] = ((end.tv_sec - start.tv_sec) * NS_PER_SEC + (end.tv_nsec - start.tv_nsec)) -
            ↪ its.it_interval.tv_nsec;
133         timer_expired = 0;
134         start = end;
135     }
136     save_results("timer.csv", jitter_data);
137 }
138
139 // Measures jitter of usleep function
140 void benchmark_usleep() {
141     long long jitter_data[ITERATIONS];
142
143     for (int i = 0; i < ITERATIONS; i++) {
144         clock_gettime(CLOCK_MONOTONIC, &start);
145         usleep(1000); // Sleep for 1 ms
146         clock_gettime(CLOCK_MONOTONIC, &end);
147
148         jitter_data[i] = ((end.tv_sec - start.tv_sec) * NS_PER_SEC + (end.tv_nsec - start.tv_nsec)) -
            ↪ 1000000;
149     }
150     save_results("usleep.csv", jitter_data);
151 }
152
153 // Main function to execute all benchmarks
154 int main() {
155     configure_realtime_scheduling(); // Set high priority scheduling
156     lock_memory(); // Prevent memory paging
157
158     printf("Getting nanosleep benchmark\n");

```

```

159     benchmark_nanosleep();
160
161     printf("Getting signal benchmark\n");
162     benchmark_signal_latency();
163
164     printf("Getting timer benchmark\n");
165     benchmark_timer();
166
167     printf("Getting usleep benchmark\n");
168     benchmark_usleep();
169
170     return 0;
171 }

```

Listing 1: merged.c

### 3 CPU & Data-Intensive Applications

To develop my CPU & data-intensive programs, I chose to use Python for ease of development. I chose `htop`<sup>2</sup> as my resource-monitoring tool as I have often used it in the past, it has easy to read & understand output, and shows you exactly what proportion of the CPU & memory is in use at that time. It also allows you to list processes by CPU consumption or memory consumption which is a useful option to have for this assignment.

```

1  import multiprocessing
2  import time
3  import argparse
4  import os
5
6  def stress_cpu(workload: float):
7      """
8      Function to create CPU load. Uses a busy-wait method to simulate CPU usage.
9
10     :param workload: The fraction of time (0.0 to 1.0) the CPU should be busy.
11     """
12     cycle_time = 0.1 # Total cycle time (100ms per iteration)
13     busy_time = cycle_time * workload # Time to stay busy
14     idle_time = cycle_time - busy_time # Time to stay idle
15
16     while True:
17         start_time = time.time()
18         while (time.time() - start_time) < busy_time:
19             pass # Busy wait
20         time.sleep(idle_time) # Sleep to control CPU usage
21
22  def start_stress_test(load: str):
23      """
24      Starts CPU stress test based on load level.
25
26      :param load: 'medium' (~50% load) or 'high' (~100% load)
27      """
28     num_cores = os.cpu_count() or 4 # Use all available CPU cores
29     workload = 0.5 if load == "medium" else 1.0 # Set workload percentage
30
31     print(f"Starting {load.upper()} CPU stress test on {num_cores} cores...")
32
33     processes = []
34     for _ in range(num_cores):
35         p = multiprocessing.Process(target=stress_cpu, args=(workload,))
36         p.start()
37         processes.append(p)
38

```

```

39 try:
40     for p in processes:
41         p.join()
42 except KeyboardInterrupt:
43     print("Stopping stress test...")
44     for p in processes:
45         p.terminate()
46         p.join()
47
48 if __name__ == "__main__":
49     parser = argparse.ArgumentParser(description="CPU Stress Test Script")
50     parser.add_argument("--load", choices=["medium", "high"], required=True, help="Choose CPU load level
51     ↪ (medium or high)")
52     args = parser.parse_args()
53
54     start_stress_test(args.load)

```

Listing 2: stress\_cpu.py




Figure 4: htop output when running python3 stress\_cpu.py --load medium

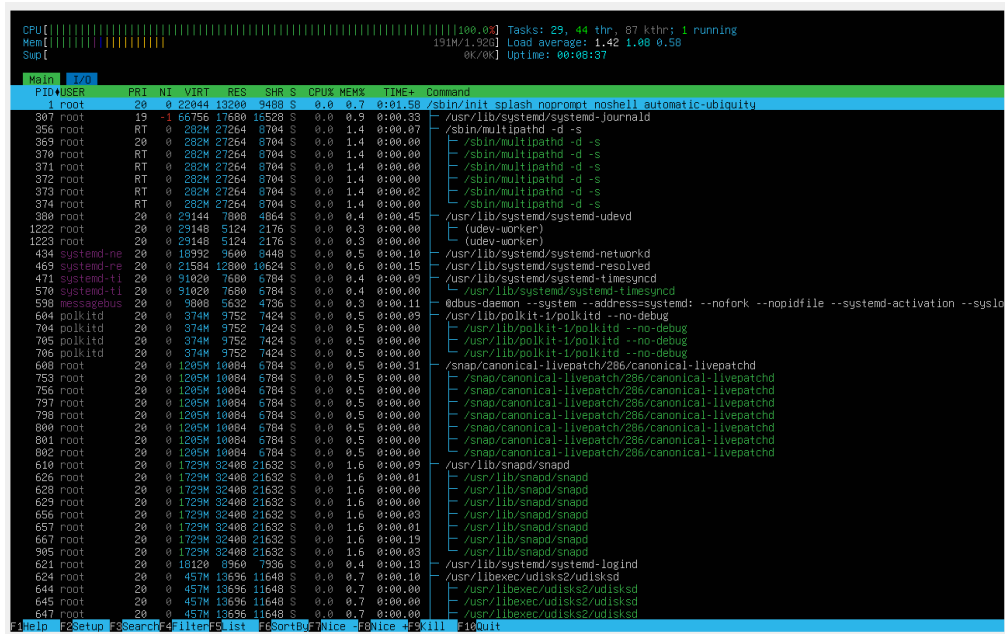


Figure 5: htop output when running python3 stress\_cpu.py --load high

```

1 import argparse
2 import time
3 import psutil
4
5 def stress_memory(target_usage: float):
6     """
7     Stress the system memory to a given percentage.
8
9     :param target_usage: Target memory usage (0.0 to 1.0, where 1.0 is 100%)
10    """
11    total_memory = psutil.virtual_memory().total # Get total RAM in bytes
12    target_memory = int(total_memory * target_usage) # Calculate target memory size
13
14    print(f"Total Memory: {total_memory / (1024**3):.2f} GB")
15    print(f"Target Memory Usage: {target_memory / (1024**3):.2f} GB ({target_usage * 100:.0f}%)")
16
17    try:
18        memory_hog = [] # List to store allocated memory chunks
19        chunk_size = 100 * 1024 * 1024 # Allocate in 100MB chunks
20
21        while sum(len(chunk) for chunk in memory_hog) < target_memory:
22            memory_hog.append(bytearray(chunk_size)) # Allocate memory
23            time.sleep(0.1) # Small delay to allow system response
24
25        print("Memory fully allocated. Holding...")
26        while True: # Keep the memory occupied
27            time.sleep(1)
28
29    except MemoryError:
30        print("Memory limit reached. Exiting...")
31    except KeyboardInterrupt:
32        print("Memory stress test stopped.")
33
34 if __name__ == "__main__":
35     parser = argparse.ArgumentParser(description="Memory Stress Test Script")
36     parser.add_argument("--usage", type=float, default=1.0, help="Target memory usage (default: 1.0 for 100%)")
37     args = parser.parse_args()

```

```
stress_memory(args.usage)
```

Listing 3: stress\_memory.py

I found that the maximum `--usage` value I could set without getting the process killed by the Linux kernel's Out-Of-Memory (OOM) killer was 0.85, so this is the value I used for my experiments.

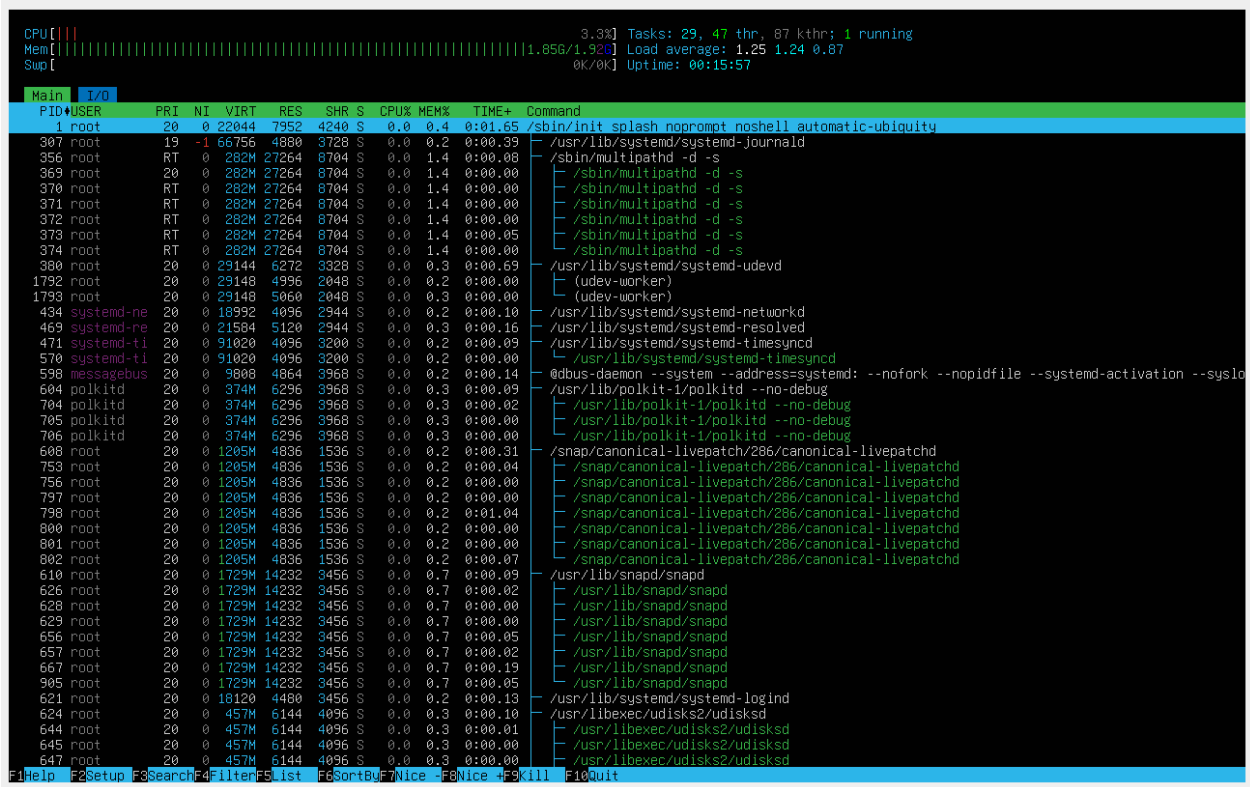


Figure 6: htop output when running `python3 stress_memory.py --usage 0.85`

## 4 Experiments

I ran the experiments in quick succession on the virtual machine by running the appropriate stresser script(s), forking it into the background using the `&` shell operator, and running the merged benchmark program. I then transferred the generated CSV files to my host machine using `scp`. To generate the plots, I wrote a Python script which will plot the mean, minimum, maximum, or standard deviation of the values collected in a bar chart for a number of given CSV files.

```
1 import pandas as pd
2 import matplotlib.pyplot as plt
3 import os
4 import argparse
5
6 parser = argparse.ArgumentParser(description="Plot specified metric from CSV files.")
7 parser.add_argument("metric", choices=["min", "max", "mean", "std"], help="Metric to plot (min, max, mean,
8 ↪ std)")
9 args = parser.parse_args()
10
11 metric_to_plot = args.metric.lower()
12 valid_metrics = {"min": "Min", "max": "Max", "mean": "Mean", "std": "Std"}
13
14 csv_files = [
15     ("../data/Locking Enabled/1. Low CPU Load, No Swap/usleep.csv", "Locking Enabled, Low CPU
16 ↪ Load, No Swap"),
17     ("../data/Locking Enabled/2. Medium CPU Load, No Swap/usleep.csv", "Locking Enabled, Medium CPU
18 ↪ Load, No Swap"),
```



```

16     ("../data/Locking Enabled/3. High CPU Load, No Swap/usleep.csv",      "Locking Enabled, High CPU
    ↪ Load, No Swap"),
17     ("../data/Locking Enabled/4. Medium CPU Load, Swap/usleep.csv",      "Locking Enabled, Medium CPU
    ↪ Load, Swap"),
18     ("../data/Locking Enabled/5. High CPU Load, Swap/usleep.csv",      "Locking Enabled, High CPU
    ↪ Load, Swap"),
19     ("../data/Locking Disabled/1. Low CPU Load, No Swap/usleep.csv",    "Locking Disabled, Low CPU
    ↪ Load, No Swap"),
20     ("../data/Locking Disabled/2. Medium CPU Load, No Swap/usleep.csv",  "Locking Disabled, Medium
    ↪ CPU Load, No Swap"),
21     ("../data/Locking Disabled/3. High CPU Load, No Swap/usleep.csv",    "Locking Disabled, High CPU
    ↪ Load, No Swap"),
22     ("../data/Locking Disabled/4. Medium CPU Load, Swap/usleep.csv",    "Locking Disabled, Medium
    ↪ CPU Load, Swap"),
23     ("../data/Locking Disabled/5. High CPU Load, Swap/usleep.csv",      "Locking Disabled, High CPU
    ↪ Load, Swap")
24 ]
25
26 column_name = "Latency/Jitter (ns)"
27
28 stats = {
29     "Metric": [],
30     "Label": [],
31     "Value": []
32 }
33
34 for file, label in csv_files:
35     if os.path.exists(file):
36         df = pd.read_csv(file)
37
38         if column_name not in df.columns:
39             print(f"Warning: Column '{column_name}' not found in {file}. Available columns:
    ↪ {list(df.columns)}")
40             continue
41
42         values = df[column_name].dropna()
43         if values.empty:
44             print(f"Warning: Column '{column_name}' in {file} is empty after removing NaN values.")
45             continue
46
47         stats["Metric"].append(valid_metrics[metric_to_plot])
48         stats["Label"].append(label)
49         if metric_to_plot == "min":
50             stats["Value"].append(values.min())
51         elif metric_to_plot == "max":
52             stats["Value"].append(values.max())
53         elif metric_to_plot == "mean":
54             stats["Value"].append(values.mean())
55         elif metric_to_plot == "std":
56             stats["Value"].append(values.std())
57     else:
58         print(f"Warning: File {file} not found.")
59
60 stats_df = pd.DataFrame(stats)
61
62 if stats_df.empty:
63     print("Error: No valid data found. Ensure the column name is correct and files are properly
    ↪ formatted.")
64 else:
65     fig, ax = plt.subplots(figsize=(16,4))
66     ax.bar(stats_df["Label"], stats_df["Value"], color="black")

```

```

67     ax.set_xticklabels(stats_df["Label"], rotation=45, ha="right")
68     ax.set_ylabel("Jitter (ns)")
69     ax.set_title(f"{valid_metrics[metric_to_plot]} usleep()")
70
71
72     plt.tight_layout()
73     plt.show()

```

Listing 4: barchart.py

It's important to note that the plots which display the mean value for each experiment could be misleading: if there was a high degree of variance in the collected results, with positive & negative values, they could cancel each other out and result in a deceptively small mean.

## 4.1 Signal Handling

The experimental data collected for the signal handling metric surprised me, as it did not match my expected results. Since this benchmark measures the latency between sending a signal to a process and the process executing it in its signal handler function, I would expect the mean latency to increase as CPU & memory load were increased. As the CPU load increases, processes can be delayed in their execution due to scheduling, and processes may be preempted, causing higher latency. I would expect high memory consumption to have similar effects, especially when memory locking is disabled, as the process data may then be swapped out, which is extremely slow & costly.

However, as can be seen in the figures below, this wasn't really the case for my collected data. The variance in my charted results seem to just be artefacts of noise in the system and fluctuations in the experimental conditions, as they don't seem to follow any discernible pattern. The main reason why I think this may have happened is because of the PREEMPT\_RT kernel patches that I installed, which turned the OS into a fully-preemptible RTS, resulting in more predictable response times, and better prioritisation of tasks; since the benchmark program runs with maximum priority, lower priority processes like my stresser scripts could get preempted in favour of the high priority benchmarking program, thus resulting in the benchmarking program not being majorly effected by the system load.

I found these results very surprising, but upon reflection, they make sense, and are indicative of the power of the Linux kernel for use in hard RTS applications when the PREEMPT\_RT patches are applied.




Figure 7: Mean latency for the signal handling benchmark




Figure 8: Minimum latency for the signal handling benchmark

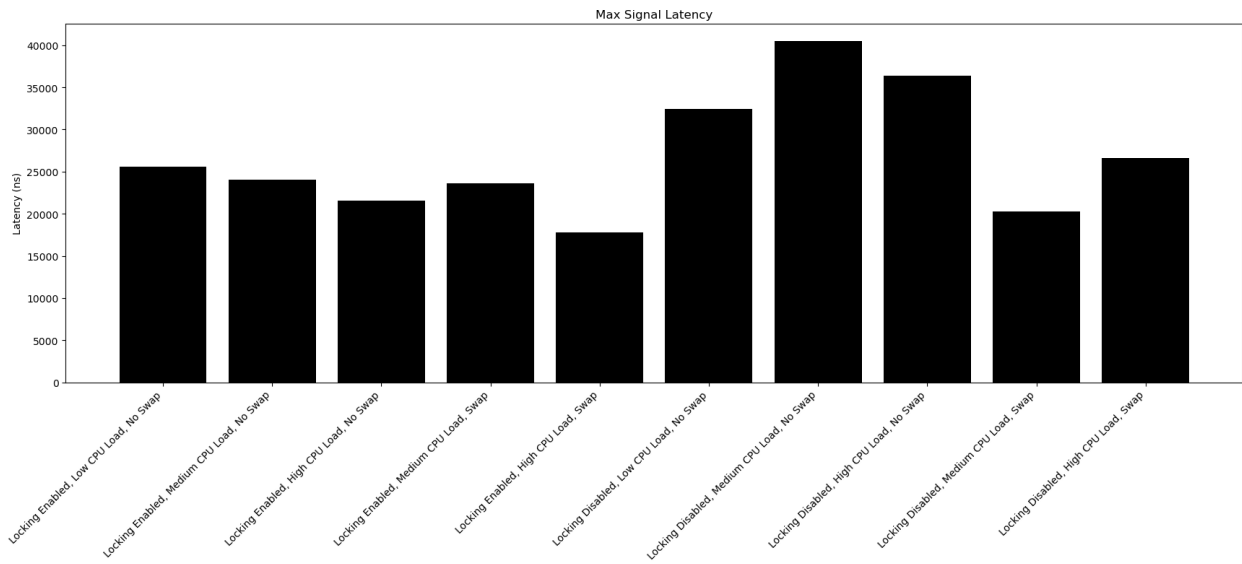


Figure 9: Maximum latency for the signal handling benchmark




Figure 10: Standard deviation of latency for the signal handling benchmark

## 4.2 Interval Timer

Since the interval timer benchmark uses a POSIX interval timer to trigger a signal at precise intervals, I would expect the time interrupts to be precisely scheduled under low CPU load, and greater delay to appear under higher CPU load due to the CPU being busy. I would also expect swapping to worsen the jitter, as accessing the memory will be in the order of milliseconds rather than microseconds. However, as previously discussed, the PREEMPT\_RT patches will help to mitigate these issues. We can see from the output data that, while not a clean trend upwards, there tends to be a higher jitter value for higher CPU loads. The most telling metric is the standard deviation; we can see from the standard deviation plot below that the variance in jitter trends upwards as CPU load & memory load increase, as one would expect.




Figure 11: Mean jitter for the interval timer benchmark




Figure 12: Minimum jitter for the interval timer benchmark




Figure 13: Maximum jitter for the interval timer benchmark




Figure 14: Standard deviation of jitter for the interval timer benchmark

### 4.3 nanosleep()

Since the `nanosleep()` benchmark measures the actual time elapsed versus the requested sleep duration, we would expect it to increase as the CPU load increases due to scheduling latency inducing jitter. Memory swapping adds large delays, and one would expect high CPU and high swap to cause erratic & unpredictable behaviour, making sleep times unreliable. The application of the `PREEMPT_RT` patches should increase the accuracy of sleep times, as the wake-ups will happen closer to the requested sleep duration and result in a lower maximum jitter value as the process can preempt other lower-priority tasks. The plotted charts don't bear a great deal of resemblance to the expected results, which is likely in large part due to the `PREEMPT_RT` patches, but also likely due to the large number of background tasks that are running at a given time on an Ubuntu system, which could be introducing noise into the data.




Figure 15: Mean jitter for the `nanosleep()` benchmark




Figure 16: Minimum jitter for the nanosleep() benchmark




Figure 17: Maximum jitter for the nanosleep() benchmark




Figure 18: Standard deviation of jitter for the nanosleep() benchmark

#### 4.4 usleep()

The `usleep()` function serves a similar role to `nanosleep()`, with the primary difference being that `usleep()` has precision in the microseconds (the `u` is an ASCII approximation of the  $\mu$  symbol typically used to symbolise the “micro” prefix) rather than in the nanoseconds, and is thus far less precise. For this reason, greater jitter is to be expected. At low CPU usage, we would expect slightly worse performance than `nanosleep()`, and for this performance to decrease as CPU usage increases; similar behaviour is to be expected as memory usage increases also. Since `usleep()` relies on signals internally, it could potentially suffer more greatly under high CPU strain. The `PREEMPT_RT` patches can help to improve response times due to the preemptible kernel, but swapping will still cause performance issues.

The most interesting plot for this benchmark is the standard deviation plot below, as it corresponds pretty much exactly to what we would expect; clearly, `usleep()` derives less performance benefit from `PATCHES_RT` than `nanosleep()`. The jitter is lowest when locking is enabled, there is low CPU load, and no swap, and increases as the CPU load & memory load are increased. When locking is disabled, there is greater performance degradations between the low CPU/memory experiment and the subsequent experiments, with the high CPU, high memory, no locking experiment yielding the greatest standard deviation, and thus the least predictability.




Figure 19: Mean jitter for the `usleep()` benchmark




Figure 20: Minimum jitter for the `usleep()` benchmark




Figure 21: Maximum jitter for the `usleep()` benchmark




Figure 22: Standard deviation of jitter for the `usleep()` benchmark

## 5 Conclusions

To conclude, as CPU load & memory load increase, performance in terms of jitter & latency are to be expected to degrade. Memory locking helps to mitigate the negative effects of high memory consumption, by preventing the memory from being swapped. Using a fully preemptible kernel like the Linux kernel with the PREEMPT\_RT patches applied can limit the negative effects of system strain, and help to ensure that deadlines are met, making such kernels a good choice for any kind of RTS, but particularly hard real-time systems.

## References

- [1] Canonical Group Ltd. *Basic Ubuntu Server Installation*. Accessed: 2025-03-18. 2025. URL: <https://documentation.ubuntu.com/server/tutorial/basic-installation/>.
- [2] Hisham Muhammad. *htop(1)*. Accessed: 2025-03-18. 2025. URL: <https://www.man7.org/linux/man-pages/man1/htop.1.html>.
- [3] Timo Rinne. *scp(1)*. Accessed: 2025-03-18. 2022. URL: <https://www.man7.org/linux/man-pages/man1/scp.1.html>.