The Offline Scheduler for Embedded Vehicular Systems

Raz Ben-Yehuda
BitBand LTD
Netanya, Israel
e-mail: razb@bitband.com

Yair Wiseman
Computer Science Department
Bar-Ilan University
Ramat-Gan, Israel
e-mail: wiseman@cs.biu.ac.il

Abstract—Nowadays, various transportation means use Linux as the operating system for their embedded control systems; however Linux uses all the processes in the hardware equally. This motivates some designer to develop an entire new Operating System for the vehicle; where a modification of Linux can produce the same and even a better product. This paper explains how this modification was designed and implemented in a commercial transportation company.

1. INTRODUCTION

Over the last years, Linux has been becoming the Operating system of choice for many transportation systems. The Linux adaption incorporates a wide range of transportation means from small cars [1] to long trains [2,3]. Also, new embedded systems of airplanes employ Linux [4] and some of them even a Linux distributed system [5].

The adaptation of Linux has many advantages; however, Linux like other current operating systems has no easy way to assign a processor to do an individual service using an undisturbed, accurate and fast way and still being a part of the operating system address space as long as the processor is an active part of an operating system. We suggest a technique of dedicating one processor of a running machine to a specific task.

Such a feature is very useful for systems that propose means of transportation that use one computer to do more than one task like the vehicle proposed in [6] that is able to detect the fire and to navigate its way to it using the very same computer. Such a feature will be also beneficial in [2] where the authors suggest using the very same computer to accomplish two tasks - monitoring a train and making sure the driver never opens the doors on the side without a passenger platform.

Linux is also common in use for navigation systems [7]. In view of the fact that navigation systems are becoming a standard in many means of transportation, an embedded computer knowing to perform the navigation task besides another tasks can be beneficial

Amdahl's Law asserts that a linear speed-up is not likely to be practicable [8,9,10], so this can motivate in some cases sparing a processor for certain tasks. That is to say we can use the Linux kernel ability to virtually hot plug an (SMT) processor [11]; instead of letting the processor mark time in endless "halts", assign the processor a service. The offline scheduler treats a processor as a device with computation ability and this is why the processor can be offloaded.

2. RELATED WORKS

A. CPU Sets

Current common technologies for service oriented systems are Linux CPU sets [12,13] and Solaris Resource Pool association [14,15]. Both of these technologies refer to assigning tasks to a set of processors, probably on the same memory node (in the NUMA case). CPU sets are actually a constraint on a (user space/kernel space) task to use resources available only in its set. CPU sets are used mainly in big systems, appliances, adapted to the client needs, where performance is an important issue. Assigning tasks to a CPU set is done using simple interfaces for maintaining memory policies and there is a very small overhead for the programmer. CPU sets was not designed to run Real-Time tasks and do not intent to reduce the interrupts over-head. This is unlike the offline scheduler who is a Real-Time scheduler. The offline scheduler may act as an alternate to softirqs and ISR. It is serialized, very accurate and provides a controlled environment to the service. The offline scheduler also protects the operating system in some cases. Therefore, Linux CPU sets and Solaris Resource Pool association serve different purposes than the offline scheduler.

B. INtime

INtime RTOS [16] for Microsoft windows is an extension for Microsoft Windows running on any platform that utilizes a standard Microsoft HAL: uniprocessor or multiprocessor. INtime RTOS has many components similar to the offline scheduler. INtime and the offline scheduler share this features:

- INtime is separated from Microsoft windows and offsched is separated from Linux; however both of the pairs share the same address space.
- Both INtime and offsched scheduler can allocate processor cores for exclusive use of INtime/offsched scheduler.
- Both INtime and offsched scheduler have high precision system timer of 50 μ s.

In spite of these, there are some other features that just one of the systems INtime or the offline scheduler maintains:

- INtime can identify user space tasks as real time processes whereas the offline scheduler cannot.
- INtime has a RealTime TCP-IP stack whereas the offline scheduler does not have.
- INTime can be used as a standalone RTOS whereas the offline scheduler cannot.
- INtime does not access Windows address space directly, whereas the offline scheduler accesses Linux address space directly.

The offline scheduler supports dynamic allocation of cores whereas INtime does not support.

C. IBM Logical partitions

IBM logical partitions [17,18] divide a server's resources into several logical units. The bus is shared; the memory and the disk storage may be partitioned and some other resources like communications can be also partitioned.

The partition model has a component similar to the offline scheduler; however as a concept the model is not relevant for the offline scheduler. In the offline scheduler system all the resources except the processor are shared. The offline scheduler can be depicted as a processor allocator; there is no administrator. Each processor has an access to some resources in the system, as long as it does not call "schedule()".

The component that is similar to the offline scheduler is a primary partition of IBM that may be depicted as CPU 0 and if an offloaded processor fails, the system will continue to run if the operating system memory is not corrupted.

D. RTOP

Any Linux machine may hang sometime. Hanging might happen due to a Kernel crash or a System overload. What can be done when a server is remote? To detect a kernel crash "netconsole" can be used; however what can be done when the system is overloaded? The server cannot be accessed because it is so

loaded that sshd (or any other daemon) does not get any cpu time. Utilities, vmstat, sar, top and so on, do provide details but under heavy load a user cannot access the machine remotely, or even locally. A user must be able to access the machine, and ask for the monitoring tool.

RTOP (Remote top) is monitoring software that provides top-like view of the system in any time. RTOP is based on the offline technology. This means that you must offline a processor (or enhance netconsole). RTOP is composed from RTOP utility ran in a remote Linux server and a driver that is invoked in the monitored server. RTOP driver pushes information out from the server. RTOP pulls information from the server.

3. THE OFFLINE SCHEDULER FOR VEHICULAR OPERATING SYSTEMS

Vehicular systems are more often than not Real-Time systems [19]; nevertheless, there are also sometimes Non-Real-Time tasks [20]. In [21] the author explains the relationship between the new trend of SMP machines and the boost of the use of Linux versions as a Real-Time operating system. The offline scheduler suggests a new version of Linux scheduler for Multiprocessors real time operating system. The new scheduler splits the processes into two groups - the processes of the traditional operating system and processes of the offline scheduler. Actually, the offline scheduler is a hybrid system, because the new system is both real time and still a standard Linux server. The real time property is mainly achieved by the nmi and the CPU isolation characteristics; however the offline scheduler still interacts with the operating system.

Caches typically maintain these properties: spatial locality, temporal locality and sequentiality. Keeping these properties in usual scheduling environments is not easy. In the offline scheduler platform, they can be easily maintained. The offline scheduler gives the programmer tools to control the cache miss rate. Low cache miss rate has a harmful effect on performance [22,23]; hence the offline scheduler adapts some rules to handle the caching:

- Spatial locality

Accesses to same or near memory locations in the lifetime of a program are very likely. The offline scheduler abides this principle in its architecture. A user may assign a task to a processor without setting a different program into that cache. This way, the cache lines are not evacuated and the program data set and code are highly available.

- Temporal locality

Given a sequence of references to n memory locations, it is likely that following reference will be made into that sequence. The offline scheduler abides this rule same as spatial locality.

- Sequentiality

Given a reference to location n , it is likely that a reference to location $n+1$ will be made in the near future. Processors utilize this rule by pre-fetching data. Prefetching can be done in software or by the hardware. There is no true benefit in the offline scheduler but the fact that running nmi make sure no undesirable prefetches happen.

Traditionally, programs' size has been increasing [24] Nowadays, programs' execution code size usually exceeds the L1i size. Most parts of a real time program are characterized as non real time whereas a relatively small portion of it is real time. When a real time code shares cache storage with a management code, or IO code, it is flushed from the instruction cache (L1i) as well as its data (L1d and L2). The offline scheduler's architecture enforces a programmer to make the distinction between the real time codes to the rest. This distinction maintains the real time portion small.

If the working set is large enough, it is very likely that several memory locations will refer to the same cache line and cache flushes and misses may happen. The offline scheduler has no need to maintain memory pools like in [25]. It has access to almost all the memory (exceptions are vmalloc'ed/ioremap memory). A program having its own memory pool might cause congestion in cache lines if memory is not shuffled enough, but this will be a fault of the Operating System; not of the offline scheduler.

Branch prediction is one of the processor's ways to speed up code execution. Problem with branch prediction is that a branch might get flushed too early. Branch flushing happens mainly due to interrupts,

exceptions, segments descriptor loads [26]. The offline scheduler is running nmi so less pipeline flushes happen.

Multiple processors systems are usually designed for two reasons, faults tolerance and program speedup. Current Linux versions have little respect to fault tolerance processors Programs' speedup is achieved by a program design and not enforced by the operating system. The offline scheduler is a system where all processors share the same memory. The offline scheduler architecture forces the programmer to write a program with a better speedup because concurrent instances of the sub units of the program run in each processor. Fault tolerance is partially achieved by the fact the even if the operating system stops (panics) or hangs the offline scheduler can still run and vise versa, if an offloaded processor failure does not corrupt the operating system's active memory the system will not stop. This is obviously cannot be compared to a failure in some user space task, so in this sense running a user space task is better. Multiprocessors systems are subjected to the same restrictions as a single processor system, such as cache size, branch prediction, BUS speed, etc; which were discussed previously.

Partitioning is the process of dividing a task into sub-tasks, each of which can be assigned to a single processor. Partitioning a process is done in compile time. Partitioning objective is to reveal maximum amount of parallelism possible without exceeding machine's hardware limitation.

A Partition analysis is preformed with some notion of a program overhead [24]. A program overhead (o) is the added time a task takes to be loaded into a processor prior to beginning execution. The larger the size of the minimum task, the smaller the partition overhead. The author of [24] describes program P work as follows: If a single processor P_1 program does operation O_1 , the parallel version of P_1 does operations O_p where $O_p \geq O_1$. For each task T_i , there is an associated number of overhead operations o_i , so that T_i takes O_i operations without overhead then:

$$O_p^T = \sum_i^p (O_i + o_i) \geq O_i$$

Where O_p is the total work that was done by P_p including overhead, and $o_i + 1 \geq o_i$. So, on one hand a better performance means maximizing parallelism and on the other hand finer code blocks increase the program overhead. Figure 1 depicts where the offline scheduler performs best.

Figure 1: effects of grain size

At the offline scheduler, $o_i \rightarrow 0$, so there is no program overhead. This is because there is no stack manipulation, and all tasks are in the processor's cache. Load may be imbalanced only if several processors handle different resources and tasks. E. g., if one processor handles network interface A and a second processor handles network interface B, and the load on the two interfaces is not the same [27]. The offline scheduler cares less for load imbalance as long as the service is being done. If not, there is something wrong with the design. When load is well balanced, we can say that:

$$O_i \approx O_1 \Rightarrow O_p^T = \sum_i^p (O_i + 0) = O_{1p}$$

Sometimes two tasks depends one on the other. In a single processor a typical scenario is:

1. T1 put message on T2 queue.
2. T1 schedule
3. Do something
4. T2 gets processor time
5. T2 replies on for T1
6. T2 schedules
7. Do something
8. T1 gets processor time
9. T1 gets T2's reply

The offline scheduler behaves very much the same. Yet, when using the offline scheduler, T1 may decide to schedule T2 directly, and only when T1 lets T2 will get a processor time; i. e. T1 chooses when T2 get processor time. Phases 7 and 3 do not exist in the offline scheduler. This small difference is extremely important when hard real time is required

Scheduling can be performed either dynamically or statically. Static scheduling is determined during compilation and dynamic is performed in real time. Ngai [28] observes that all major issues in runtime scheduling focus in insuring performance and reduce overhead losses. In multiple processors this is even a more important issue [29]. The following four run time scheduling major overheads include:

- Gathering information

Gather system and programs real time information. This includes, control structure, identification of critical paths and dependencies. Dynamic information includes work loads and resource availability. The offline scheduler has no need of gathering information. It is up to the designer to decide what information will be used. E. g. the offline scheduler timer uses "rdtsc" counters and Rdtsc has a very small overhead.

- Scheduling

In the usual course of events, scheduling is performed as shown in Table 1:

	Advantages	Disadvantages
Compile Time	Less overhead	Lack of fault tolerance Compiler lacks stall information
Run time	More efficient execution	Higher overhead

Table 1

When using the offline scheduler, scheduling is performed in run time and has a very little overhead. The offline scheduler has a single stack context [30]. The offline scheduler scheduling is mere procedure invocation. There are two scheduling procedure in the offline scheduler:

- `offsched_schedule`

A routine is a registered in the offline scheduler pending queue. In the offline scheduler there is both fast scheduling, simplicity and hardly any overhead.

- `The offsched reschedule`

A routine for assignment of a task into offsched run queue (calendar). It is lockless and very simple.

- Dynamic execution control

This includes clustering and process creation at run time. In the offline scheduler, A process creation is the actual assignment to a designated processor (offsched schedule); there is no creation of a context as in the familiar UNIX. There are no pids, gids or similar features. There is no notion of clustering of processes.

- Dynamic data management

This includes a minimization of memory overhead delay when accessing data. Delay may improve by assigning tasks by a policy of a minimum memory delay. In the offline scheduler, memory access is matter good scheduling design. Due to the offline scheduler serialized nature, memory low latency scheduling policy can be easily achieved.

Consider the following scenario: Process A assigns a message to process B. Process A schedules process B after itself and exit. Process B runs right after and the memory is in the processor cache. When process B is completed, it will schedule process A and will exit. During this process, there is no cache miss as long as the data working set is smaller than the processor's cache.

4. COMMUNICATION AND SYNCHRONIZATION

A typical programmer does not program when the processor cache or MESI (Illinois Protocol) in his mind. She does not know even on what hardware her program will be running on. So, her program is likely to have in-efficient memory accesses. If the program is intended to be running on multiprocessor machine, queues and synchronization primitives will be used. So cache misses, locks (atomic increment must seize the bus in a SMP system) and preemption are likely to load the system. The offline scheduler serialized nature reduces this contention.

One may ask, why not create a busy loop like:

```
While (can I do my thing)
    do my thing
```

This code will create a RCU starvation. RCU is a synchronization technique which enables multiple readers and multiple writers to access mostly accessed data structures from multiple processors. RCU starvation happens because each processor must walk through a quiescent state [31]. A quiescent state is when one of the bellow happens:

1. A processor performs a context switch.
2. A processor executes user mode.
3. A Processor executes the idle loop.

And neither of these will happen in a busy loop. The offline scheduler eliminates RCU starvation since the processor does not have to walk through a quiescent state as it is not part of the operating system.

Consider the following sequence of operations:

```
INC RAX
ADD RCX, RBX
```

These instructions are independent. In regular pipelined processors these two instructions are not likely to run concurrently. Some commands might be associated with interrupts and exceptions. If the interrupt service routine runs in the initiating processor context, RCX and RBX content will be lost unless some additional logic is added to the processor [32]. Either case, a processor has to do additional work that reduces performance. The NMI property of the offline scheduler actually makes programs running in the offline scheduler context faster than a in an active processor. When designing multi-threaded software, one has to understand the price of a context switch with respect to the system requirements.

Moore Law presents us a simple statement, if in 2008 we have 4 cores in a single die, in 2010 we will be having 8 cores, in 2012 16 cores, in 2014 32 cores etc. Does the current Linux operating system design fits the Multi core era? Should the operating system handle so many cores? As this paper argues, the answer is: "not always". Having an operating system balancing tasks between 16 processors is not necessarily good. The Linux kernel migrate pages and move tasks between processors repeatedly to achieve balance in the system. The offliine scheduler on other hand, does not care much for imbalance; it aims for responsiveness, accuracy and simplicity. The offline scheduler can produce linear speed-up as long as it is not bounded by BUS speed.

5. THE OFFLET

In the offline scheduler environment, we do not use the terms threads, processes or softirqs. The offline scheduler refers to a processor as a device running an offlet. Offlet is the context of an offlined processor. When an offlined processor is removed from the operating system, It is simply set to the "halt" machine instruction while interrupts are disabled. Right before the halt instruction is executed, the offline scheduler is invoked. From now on, we are in an offlet context. Offlet context has some restriction that a user must be aware of:

1. A user may not cause any page fault, meaning he cannot access vmalloc'ed memory (only by "walking on the pages").
2. A user may not "STI".
3. A user may not invoke any code path that ends up in "schedule".

This is why the author refers to the offline scheduler context as an offlet; it is because it has more constraints than any other kernel context. offlets may be scheduled in time T, when T can be any time starting from now. Offlets may be setup on the stack, schedule and released; very much like any other kernel entity but much faster.

The offline scheduler is based on the notion that processors are not an expensive resource. In general, when task A is in kernel context and wishes to seize resource X all it does:

- Step 0) ..
- Step 1) Lock
- Step 2) ..
- Step 3) Unlock
- Step 4) ..

If Step 4 is not serialized with respect to the previous steps then Offlet will suggest an asynchronous service, instead of spinning in step 1, the offline processor is asked to do steps 1 2 3 for the user. Offlet scheduling is very light so the scheduling cost is negligible in term of computation. Another added value of offlets is that a system may choose to serialize access to a device through an offlet. A good example is offline napi which is described below. This serialization actually relieves the kernel from contention on slow devices.

When any context enters the vmalloc area, it must generate a page fault in each processor this if the referenced address was not referenced before on the same processor. This is because vmalloc area is a dynamic mapped area and the processor's MMU has no reference to it. kmalloc address space is set on the kernel static page table that never change. For an offlet to access vmalloc area, or user pages, it must walk on the pages. Walking on the pages means dismantling the address to pages and calling kmap atomic on each page.

6. THE OFFLINE TIMER

The offline scheduler timer is a plain busy-pause loop running at the frequency of the CPU (in our tests it has been 2 GHZ). Each time unit is divided into N slots; E.g. 1ms is divided into slots of 1000us. There is still a need to adjust the timer, because it uses the rdtsc mechanism which might be inaccurate. If a 100us resolution meets the requirements, HPET overhead will be quite small, about just 3% cpu usage in our tests (Lenovo T61 laptop). So, in such a case the offline scheduler timer will be redundant.

In cases when just few timer handlings are needed i.e. few timer handlings every several hundreds milliseconds, one can use TICKLESS [33] timer implementation.

However, since an operating system does other tasks, like processing packets, an interrupt might be nested and thus be behind schedule. In addition, when the amount of timers is larger, the accuracy is less important.

7. OFFLINE NAPI

Napi [34] stands for New API. Napi is a technique for interrupts mitigation in networking in the Linux kernel. NAPI technique is a simple polling over incoming packet arrival queue; this eliminates the constant interrupt processing. By default, Linux uses receive interrupt for packet processing, and activate NAPI when a certain threshold of number of incoming packets is reached. This threshold is known as the network "weight". Only one processor may call poll, this is because only a single processor can get the initial interrupt. There are some flaws of NAPI.

1. Latency. In some cases, NAPI may introduce additional software IRQ latency.
2. IRQ Masking. On some devices, changing the IRQ mask may be a slow operation, or requires additional locking. This overhead may negate any performance benefits observed with NAPI.
3. Rotting Packet. In some cases a packet may rot in the RX ring of the device. This happens due to a race between the time we decide we have no packet in the RX ring and the time we are forced to enable interrupts.

The offline scheduler NAPI is aimed to solve the above problems and provide infrastructure to other services for packet processing.

- Latency

Very much like NAPI, the offline scheduler NAPI will process incoming packets from RX ring. Unlike NAPI, spinning processors will have very little latency (if any) so disabling RX interrupts entirely is possible, this way we eliminate both problems. In addition, since we have no initial interrupt (an interrupt that triggers NAPI) we can have several processors spinning over the same device.

- IRQ masking

The duration of the operation of IRQ masking is not a problem because it is done only once. Once IRQ masking is complete, the offline scheduler NAPI is not bothered with it again.

- Rotting Packet

Because interrupts are never enabled and the operating system continuously spins over the network devices, this problem does no longer exist.

- SMP IRQ Affinity

Current IRQ affinity is in a device granularity. In offline NAPI, granularity may be based on service affinity. Meaning, instead of having all packets of a device routed to a group of processors, packets of type A may be assigned to processor X, type B to processor Y and so on. The rational here is that there is no sense in putting together ARP queries on the same processor as the user's application's processor.

8. EVALUATION

We have generated traffic through eth1 to eth1. Each machine has been connected to two different network segments. Traffic has been generated on the 10.9 segment.

Figure 2: traffic configuration

Figure 2 depicts this test. The grey area is where the offline firewall process runs. It controls all interfaces. The test was conducted on a Supermicro PDMSI machine, with an Intel Pentium 3.4GHZ,

Hyper threading enabled. Receiving interface was Intel 1Gbps 82546EB. e1000 and Linux kernel version was 2.6.30. The Loader is generating traffic Linux packet generator driver, also is known as pktgen. Pktgen UDP port is 9, which is the discard port.

```
ifconfig eth1 10.9.0.62/24
ping 10.9.0.208 -I eth1 -c 3
insmod offsched.ko
mknod /dev/offschedctl c 252 0
offschedctl -a 1
offschedctl -p 1 --setservice napi
offschedctl -i
insmod offlet_napi.ko
insmod offletnet.ko    ransmission driver
insmod filter.ko       firewall driver
offschedctl --start napi
insmod offlet_srv.ko saddr=172.16.1.62
daddr=172.16.33.203 dport=7777 sport=7777
loads OFFSCHED server on source address eth0
port 7777. The destination address is not important; it is
use to get a routing interface pre-maturely.
```

Figure 3: Loading procedure

Figure 3 details the loading procedure. The test was first conducted with Linux Native NAPI. Then CPU1 was dropped and offsched was run.

Figure 4 shows the result of native NAPI vs. the offline scheduler. Using the offline scheduler, the hardware interrupts was 6% because receiving interrupts has not been disabled. When using Native NAPI, The consumption of CPU1 was 35%, whereas cpu0 is mostly idle. In the offline kernel CPU0 is 7% busy due to hardware interrupts.

```
Linux Native NAPI
Tasks: 66 total, 1 running, 65 sleeping, 0 stopped, 0 zombie
Cpu0 : 0.0%us, 0.0%sy, 0.0%ni, 100.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Cpu1 : 0.0%us, 0.0%sy, 0.0%ni, 65.3%id, 0.0%wa, 5.0%hi, 29.7%si, 0.0%st
Mem: 1025452k total, 153536k used, 871916k free, 5084k buffers
Swap: 0k total, 0k used, 0k free, 75568k cached

OFFLET SMT CPU1 is dropped and runs OFFSCHED.
Tasks: 58 total, 1 running, 57 sleeping, 0 stopped, 0 zombie
Cpu0 : 0.4%us, 0.4%sy, 0.0%ni, 92.9%id, 0.0%wa, 6.3%hi, 0.0%si, 0.0%st
Mem: 1025452k total, 194256k used, 831196k free, 22528k buffers
Swap: 0k total, 0k used, 0k free, 92808k cached
```

Figure 4: Native NAPI vs. offlsched

The Intel network Interface was loaded with 1Gbps of incoming traffic as observed by bmon. In terms of pure processor consumption, 107/200 is the total processing consumption of offline NAPI while in native NAPI it is 35/200. So, one may wonder, are not we wasting a processor? The answer is actually yes, we do. But what we do earn is an operating system running undisturbed. If this machine has been loaded with more traffic, then the operating system may be too loaded to even be accessed. In a machine

employing offline scheduler, this inaccessible state will not occur, because a denied packet processing is confined to the offloaded processor.

9. CONCLUSION

Linux is a good platform for embedded systems and particularly for vehicular autonomous systems [35,36]. Adding the offline scheduler will enhance the Linux abilities. The offline scheduler joins two different system types in a single machine, a Real Time system and a standard Linux machine. This is very helpful for systems like [2] where the authors use the embedded computer to monitor the train speed (a real time task) and to make sure the driver never opens the doors on the side without a passenger platform (a non-real-time task). Offloading processors to serve as dedicated engines might be very helpful in various real time scenarios. Using the offline scheduler, there is no need to buy expensive offloading network cards when the same feature can be achieved by commodity hardware. Who does need network processors for this feature, when every standard processor can do the same thing?

10. REFERENCES

- [1] D. Geer, "Survey: Embedded Linux Ahead of the Pack," IEEE Distributed Systems Online, vol. 5, no. 10, pp. 3, Oct. 2004
- [2] D. W. Carr, R. Ruelas, H. Salcedo-Becerra, and G. A. Ponce-Castaneda, "A Linux-based System to Monitor Train Speed and Doors for a Light-Rail System", Eight Real-Time Linux Workshop, Lanzhou, Gansu, China, October, 2006.
- [3] Z Jianhua, Design of Electrical Monitoring and Control Terminals for Trains Based on Linux, Industrial Control Computer, 2005.
- [4] V. Srovnal Jr, and J. Kotzian, Development of a Flight Control System for an Ultralight Airplane, International Multiconference on Computer Science and Information Technology. IMCSIT, 745-750, 2008.
- [5] Kepner J., Moore M., Travinin N., Kim H., Reuther A., Currie, T., McCabe A., Mathew B., Rabinkin D. and Rhoades A. Deployment of SAR and GMTI signal processing on a Boeing 707 aircraft using pMatlab and a bladed Linux cluster, Technical Report, Massachusetts Institute of Technology, Lexington Lincoln Lab, 2004.
- [6] Xu Y., Li D., Chen S. and Chang, H., Intelligent fire-fighting vehicle model based on embedded system, Journal of Computer Applications, Vol. 30(2), 2010.
- [7] Walchko K. J., Nechyba M. C., Schwartz E. and Arroyo A., Embedded low cost inertial navigation system, Florida Conference on Recent Advances in Robotics, FAU, Dania Beach, FL, 2003.
- [8] M. Hill and M. Marty, "Amdahl's Law in the Multicore Era" IEEE Computer, vol. 41(7), pp. 33-38, July 2008.
- [9] M. J. Flynn, "Computer Architecture Pipelined and Parallel Processor design", by Jones and Bartlett Publishers, 1995.
- [10] D. Eadline, "Multi-Core Melee", Linux Magazine, Issue #80, pp 40-41, July 2007.
- [11] Dave Boutcher and Dave Engebretsen, "Linux Virtualization on IBM POWER5", Systems. In Ottawa Linux Symposium, pp. 113-120, July 2004.
- [12] S. Derr, CPUSETS, <http://lxr.linux.no/linux+v2.6.26.3/Documentation/cpusets.txt>, BULL SA 2004.
- [13] P. Shinde, P. Sharma, S. Guntupalli, "Automated Process Classification Framework using SELinux Security Context," Third International Conference on Availability, Reliability and Security, 2008. ARES 08, pp.592-596, March 2008.
- [14] "Solaris dedicated CPU, "<http://docs.sun.com/app/docs/doc/8171592/gepsd?a=view>

- [15] W. Gentzsch, "Sun Grid Engine: Towards Creating a Compute Power Grid", pp.35, First IEEE International Symposium on Cluster Computing and the Grid (CCGrid'01), 2001.
- [16] INtime, tenasys, <http://www.tenasys.com/products/intime.php>.
- [17] "IBM Logical Partition Concept", <http://publib.boulder.ibm.com/infocenter/series/v5r3/index>.
- [18] J. Jann, L. M. Browning, R.S. Burugula, "Dynamic reconfiguration: Basic building blocks for autonomic computing on IBM pSeries servers", IBM Systems Journal, Vol. 42(1), pp. 29-37, 2003.
- [19] Y. Wiseman, Take a Picture of Your Tire!, Proc. IEEE Conference on Vehicular Electronics and Safety (IEEE ICVES-2010) Qingdao, ShanDong, China, pp. 151-156, 2010.
- [20] I. Grinberg and Y. Wiseman, Scalable Parallel Collision Detection Simulation, Proc. Signal and Image Processing (SIP-2007), Honolulu, Hawaii, pp. 380-385, 2007.
- [21] P. E. Mckenney, "SMP and embedded real-time", Linux Journal, issue 153, p. 1, Jan 2007.
- [22] H. Tomiyama, H. Yasuura, Code placement techniques for cache miss rate reduction, ACM Transactions on Design Automation of Electronic Systems (TODAES), Volume 2 (4), pp. 410 - 429, October 1997.
- [23] U. Drepper. "What every programmer should know about memory" Red hat INC Copyrights 2007, <http://people.redhat.com/drepper/cpumemory.pdf>
- [24] M. J Flynn. "Computer Architecture Pipelined and Parallel Processor Design", Jones and Bartlett Publishers Inc., 1995.
- [25] M. Reuven and Y. Wiseman, Medium-Term Scheduler as a Solution for the Thrashing Effect, The Computer Journal, Oxford University Press, Swindon, UK, Vol. 49(3), pp. 297-309, 2006.
- [26] "IA-32 Intel Architecture Software Developer Manual Volume 3 System programming guide", <http://pdos.csail.mit.edu/6.097/readings/intelv3.pdf>, pp. 7-34, 2002.
- [27] Y. Wiseman, K. Schwan and P. Widener, Efficient End to End Data Exchange Using Configurable Compression, Proc. The 24th IEEE Conference on Distributed Computing Systems (ICDCS 2004), Tokyo, Japan, pp. 228-235, 2004.
- [28] T. F. Ngai, "run time resource management in concurrent systems", PhD thesis, Department of Electric engineering, Stanford university, January 1992.
- [29] Y. Wiseman and D. G. Feitelson, Paired Gang Scheduling, IEEE Transactions on Parallel and Distributed Systems, Vol. 14(6), pp. 581-592, 2003.
- [30] Y. Wiseman, J. Isaacson and E. Lubovsky, "Eliminating the Threat of Kernel Stack Overflows," in Proceedings of IEEE International Conference on Information Reuse and Integration (IRI 2008), pp. 116-121, Las Vegas, USA, July 2008.
- [31] D. P. Bovet & M. Cesati. "Understanding the Linux Kernel". O'Reilly Media Inc. Copyrights 2006.
- [32] D. Harry, "Computer organization for multiple and out-of-order execution of condition code testing and setting instructions", US Patent 5630157, 1997.
- [33] "kerneltrap", <http://kerneltrap.org>
- [34] R. Bolla and R. Bruschi, "high-end Linux based Open Router for IP QoS networks: tuning and performance analysis with internal (profiling) and external measurement tools of the packet forwarding capabilities", Proc. of the 3rd International Workshop on Internet Performance, Simulation, Monitoring and Measurements (IPS MoMe 2005), pp. 203-214, 2005.
- [35] J. Williams and N. Bergmann, Embedded Linux as a platform for dynamically self-reconfiguring systems-on-chip, Proc. Int. Conf. on Engineering of Reconfigurable Systems and Algorithms, 2004.
- [36] K. Yaghmour, J. Masters, P. Gerum and G. Ben-Yossef, Building embedded linux systems, O'Reilly Media, Inc., 2008.