Optimization of Operating Systems towards Green Computing

Appasami.G
Assistant Professor, Department of CSE, Dr. Pauls Engineering College,
Affiliated to Anna University – Chennai, Villupuram, Tamilnadu, India
E-mail: appas_9g@yahoo.com

Suresh Joseph.K
Assistant Professor, Department of computer science,
Pondicherry University, Pondicherry, India
E-mail: sureshjosephk@yahoo.co.in

Abstract. Green Computing is one of the emerging computing technology in the field of computer science engineering and technology to provide Green Information Technology (Green IT / GC). It is mainly used to protect environment, optimize energy consumption and keeps green environment. Green computing also refers to environmentally sustainable computing. In recent years, companies in the computer industry have come to realize that going green is in their best interest, both in terms of public relations and reduced costs. Information and communication technology (ICT) has now become an important department for the success of any organization. Making IT "Green" can not only save money but help save our world by making it a better place through reducing and/or eliminating wasteful practices. In this paper we focus on green computing by optimizing operating systems and scheduling of hardware resources. The objectives of the green computing are human power, electrical energy, time and cost reduction with out polluting the environment while developing the software. Operating System (OS) Optimization is very important for Green computing, because it is bridge for both hardware components and Application Soft wares. The important Steps for green computing user and energy efficient usage are also discussed in this paper.

Keywords: Optimization of Operating Systems, Green Computing, Green Information Technology, Energy Efficiency.

1 Introduction

1.1. Green computing

Green computing is the study and practice of efficient and eco-friendly computing resources with conservation of power energy. In recent years, companies in the computer industry show much more interest in green computing because it saves energy and expenditure cost. In 1992, the U.S. Environmental Protection Agency (EPA) launched Energy Star, a voluntary labeling program which is designed to promote and recognize energy-efficiency in monitors, climate control equipment, and other hardware and Software technologies. This resulted in the widespread adoption of sleep mode among consumer electronics. The term "green computing" was probably coined shortly after the Energy Star program began; there are several USENET posts dating back to 1992 which use the term in this manner. Concurrently, the Swedish organization TCO Development launched the TCO Certification program to promote low magnetic and electrical emissions from CRT-based computer displays; this program was later expanded to include criteria on energy consumption, ergonomics, and the use of hazardous materials in construction. In the article Harnessing Green IT: Principles and Practices, San Murugesan defines the field of green computing as "the study and practice of designing, manufacturing, using, and disposing of computers, servers, and associated subsystems such as monitors, printers, storage devices, and networking and communications systems efficiently and effectively with minimal or no impact on the environment [1] [10] [9].

1.2. Energy Efficiency

We will very briefly outline some of the important aspects through which energy efficiency can be pursued in computers. More focus has been put on algorithmic efficiency and optimization. A fast algorithm doesn't only save time; it also saves loads of computer resources and in turn, energy. Computer power management issues are directly related to the operating system which is used. Most modern operating systems and hardware support the Advanced Configuration and Power Interface (ACPI) standard, which allows power consumption to be enormously reduced when the computer is idle or under low load. Voltages and operating frequencies of most computer parts are scalable in real time, meaning that performance (and power consumption) can be easily adjusted as needed.

Big datacenters resort to cooling solutions which involve less air conditioning, like underground installations. The "80 PLUS" power supply industry initiative certifies power supplies that are at least 80% efficient, rather than the typical 70-75%. Power consumption of storage devices can be reduced either by using 2.5 inch laptop hard disks or, even better, by using the ever-cheaper solid state drives which except for more energy efficient, are also significantly faster. Computer graphics cards consume growing amounts of energy. Many datacenters use computers with absolutely no graphics card as the administration is usually done remotely and no physical display is needed. The nVidia Optimus technology introduced recently also offers power savings by automatically switching from your energy-hungry beast to a more friendly onboard graphics card. Computer monitors are also becoming more energy efficient with the recent introduction of LED backlighting and the less recent introduction of LCD monitors which rendered CRT monitors obsolete. Recycling computer parts, batteries and printer cartridges is also of vital importance [1] [2] [4].

1.3. General rules to be followed for green computing:

- Switching off your computer and turning it on again won't damage it. In fact, it's even good for your PC, as computers slow down if they are not rebooted every now and again.
- When you are buying a printer, a scanner, or any other device, make sure that it has the Energy Star logo on it. Only the most energy-efficient devices are approved by Energy Star. Using these devices will help you significantly reduce power consumption.
- If your computer RAM is struggling, your PC requires more power. The best way to fight this problem is to add more RAM if your computer has the capability. If you are using a 32-bit operating system, then your computer could support up to 3GB of RAM. 64-bit operating systems can support a lot more memory.
- In addition to adding RAM, disable unnecessary programs launching on startup, as programs running in the background can consume a lot of memory.
- Another way to free up RAM is to optimize Windows services that run in the background. A lot of the default Windows services are never required for home use. Disabling them will make your computer faster and more energy efficient.
- Energy efficiency depends on the performance of your computer. If your PC is slow, has lots of errors, and is generally resource-hungry, it consumes a lot of extra energy that could be saved. Keeping your computer well-maintained and optimized will make it more energy efficient and ultimately help the planet.

Following these simple tips will help you reduce your monthly electricity bill and do something useful for the planet.
The key to green computing can be summarized in one word: efficiency. This word dominates the design, manufacturing, use and disposal phases of all green computing products. Particular emphasis is given to the longetivity of products, since their construction accounts for more than 70% of the ecological footprint a product leaves behind in its lifetime. By increasing its lifetime, this footprint is spread across many years [4] [5] [7].

1.4. Tips for Green Computing User

1. Turn off your computer at night so it runs only eight hours a day— will reduce energy use by 810 kWh per year and net a Sixty seven percent (67%) annual savings.
2. Plug your computer into a surge protector with a master control outlet, which automatically senses when the computer is not in use and cuts power to it and all your peripherals.
3. Purchase flat-screen monitors—they use significantly less energy and are not as hard on your eyes as CRTs.

4. Purchase an Energy Star–compliant computer. Note that laptop models use much less energy than desktop units.
5. Plan your computer-related activities so you can do them all at once, keeping the computer off at other times.
6. Consider a smaller monitor—a 14-inch display uses 40 percent less energy than a 17-inch one.
7. Enable the standby/sleep mode and power management settings on your computer.
8. Forgo the screen saver—it doesn't save energy or your screen unless you're using an old monochrome monitor.
9. Review document drafts and e-mails onscreen instead of printing them out.
10. Power off your monitor when you are not using it, instead of using screen savers.
11. Consider using an ink-jet printer—although a bit slower than laser printers, inkjets use 80 to 90 percent less energy.
12. Buy vegetable or non-petroleum-based inks— they are made from renewable resources, require fewer hazardous solvents, and often produce brighter, cleaner colors.
13. Turn off all printers and peripherals unless you are using them.
14. Do not leave the computer running overnight or on weekends.
15. Choose dark backgrounds for your screen display—bright-colored displays consumer more power.
16. Reduce the light level in your room when you are working on your computer.
17. Network and share printers where possible.
18. Print on recycled-content paper. Look for non-chlorine bleached papers with 50 to 100 percent post-consumer waste.
19. Use double-sided printing functions.
20. E-mail communications as an alternative to paper memos and fax documents.

These are all important tips for every green computing user for energy efficient green environment [4] [5] [6].

2. EXISTING GREEN APPROACHES

Green IT is the Principles and Practices, San Murugesan defines the field of green computing as "the study and practice of designing, manufacturing, using, and disposing of computers, servers, and associated subsystems such as monitors, printers, storage devices, and networking and communications systems, efficiently and effectively with minimal or no impact on the environment." the environmental affects of green computing should be addressed: Green use, green disposal, green design, and green manufacturing [6].

2.1. Network communication

Huge amount of data are transferred from one location to another location, so there is always heavy traffic which leads to delay in sending and receiving data. Before sending data it should be compressed and encrypted, so that the size of data will be reduced and network traffic also reduced. Modern IT systems rely upon a complicated mix of people, networks and hardware; as such, a green computing initiative must cover all of these areas as well. A solution may also need to address end user satisfaction, management restructuring, regulatory compliance, and return on investment (ROI). There are also considerable fiscal motivations for companies to take control of their own power consumption; "of the power management tools available, one of the most powerful may still be simple, plain, common sense [7] [8].

2.2. Heat Reduction

Computer cooling is required to remove the waste heat produced by computer components, to keep components within their safe operating temperature limits. Various cooling methods help to improve processor performance or reduce the noise of cooling fans [9].

Components which produce heat and are susceptible to performance loss and damage include integrated circuits such as CPUs, chipset and graphics cards, along with hard drives (though excessive cooling of hard drives has been found to have negative effects). Overheated parts fail early and may give sporadic problems resulting in system freezes or crashes. Both integral and peripheral means are used to keep the temperature of each component at a safe level. With regard to integral means, CPU and CPUs are designed with energy efficiency, including heat dissipation, in mind; though improved efficiency may only allow increased performance instead of reduced heat. Peripheral means include heat sinks to increase the surface area which dissipates heat, fans to speed up the exchange of air heated by the computer parts for cooler ambient air, and in some cases soft cooling, the throttling of computer parts in order to decrease heat generation [4] [5].

As a safety measure, many computers are designed to turn themselves off if the internal temperature exceeds a certain point. Alternatively, some have an option in their BIOS that allows the user to determine if the system emits an alarm beep or shuts itself down when the core temperature reaches the level set by the user. However, setting this incorrectly can result in hardware damage or erratic system behavior [4] [5].

2.3. Hardware Product longevity

Gartner maintains that the PC manufacturing process accounts for 70 % of the natural resources used in the life cycle of a PC. Info World July 06, 2009; Therefore, the biggest contribution to green computing usually is to prolong the equipment's lifetime. Another report from Gartner recommends to "Look for product longevity, including upgradeability and modularity.", For instance, manufacturing a new PC makes a far bigger ecological footprint than manufacturing a new RAM module to upgrade an existing one, a common upgrade that saves the user having to purchase a new computer. Power and cost are saved very much by making use of reusable components and long life hardware [4] [5].

2.4 Power supply

Desktop computer power supplies (PSUs) are generally 70–75% efficient, dissipating the remaining energy as heat. An industry initiative called 80 PLUS certifies PSUs that are at least 80% efficient; typically these models are drop-in replacements for older, less efficient PSUs of the same form factor.80 PLUS As of July 20, 2007, all new Energy Star 4.0-certified desktop PSUs must be at least 80% efficient [4] [9].

3. SOFTWARE DEPLOYMENT OPTIMIZATION

3.1. Algorithmic efficiency

The efficiency of algorithms has an impact on the amount of computer resources required for any given computing function and there are many efficiency trade-offs in writing programs. As computers have become more numerous and the cost of hardware has declined relative to the cost of energy, the energy efficiency and environmental impact of computing systems and programs has received increased attention. A study by Alex Wissner-Gross, a physicist at Harvard, estimated that the average Google search released 7 grams of carbon dioxide (CO_2). However, Google disputes this figure, arguing instead that a typical search produces only 0.2 grams of CO_2 [10].

3.2. Resource allocation

Algorithms can also be used to route data to data centers where electricity is less expensive. Researchers from MIT, Carnegie Mellon University, and Akamai have tested an energy allocation algorithm that successfully routes traffic to the location with the cheapest energy costs. The researchers project up to a 40 percent savings on energy costs if their proposed algorithm were to be deployed. Strictly speaking, this approach does not actually reduce the amount of energy being used; it only reduces the cost to the company using it. However, a similar strategy could be used to direct traffic to rely on energy that is produced in a more environmentally friendly or efficient way. A similar approach has also been used to cut energy usage by routing traffic away from data centers experiencing warm weather; this allows computers to be shut down to avoid using air conditioning [4]. Larger server centers are sometimes located where energy and land are inexpensive and readily available. Local availability of renewable energy, climate that allows outside air to be used for cooling, or locating them where the heat they produce may be used for other purposes could be factors in green sitting decisions [10] [11].

4. OPERATING SYSTEMS TOWARDS GREEN COMPUTING

The dominant desktop operating system, Microsoft Windows, has included limited PC power management features since Windows 95. These initially provided for stand-by (suspend-to-RAM) and a monitor low power state. Further iterations of Windows added hibernate (suspend-to-disk) and support for the ACPI standard. Windows 2000 was the first NT based operation system to include power management. This required major changes to the underlying operating system architecture and a new hardware driver model. Windows 2000 also introduced Group Policy, a technology which allowed administrators to centrally configure most Windows features. However, power management was not one of those features. This is probably

because the power management settings design relied upon a connected set of per-user and per-machine binary registry values, effectively leaving it up to each user to configure their own power management settings [12].

This approach, which is not compatible with Windows Group Policy, was repeated in Windows XP. The reasons for this design decision by Microsoft are not known, and it has resulted in heavy criticism Microsoft significantly improved this in Windows Vista by redesigning the power management system to allow basic configuration by Group Policy. The support offered is limited to a single per-computer policy. The most recent release, Windows 7 retains these limitations but does include refinements for more efficient user of operating system timers, processor power management, and display panel brightness. The most significant change in Windows 7 is in the user experience. The prominence of the default High Performance power plan has been reduced with the aim of encouraging users to save power [10].

There is significant market in third-party PC power management software offering features beyond those present in the Windows operating system. Most products offer Active Directory integration and per-user/per-machine settings with the more advanced offering multiple power plans, scheduled power plans, anti - insomnia features and enterprise power usage reporting [4] [6].

4.1. Virtualization

We would like to look at processing virtualization and the role it plays in green computing. Processing virtualization, by the way, is using hardware and software to break the link between an application, application component, system service or whole stack of software and the underlying machine(s). There are a range of different types of virtualization technology in the virtual processing software category including the following [10] [13] [14].

- High performance computing software that allows applications or data to be segmented (sometimes the word "decomposed" is used) into units that can be run simultaneously on many machines to reduce the amount of time necessary to execute applications. There are several different communities using this type of software. Some call it "grid computing," some call it "parallel processing, and others make references to the marching brooms in the Disney film "Fantasia." This one application to many machines approach is typically found supporting scientific, research, geophysical, weather and content creation applications.
- High availability/clustering software that is used to monitor an entire application stack or service. If the application or service slows down or stops completely, a replica of the application or service is started on another machine. This approach is commonly used to support a monolithic function such as a collaborative application or a database management engine.
- Virtual machine software that gets beneath an then encapsulates an entire stack of software from the operating system, data management software, middleware and applications themselves. This software then allows multiple "capsules" or virtual machines, each running its own operating system, to run on the same piece of physical hardware. This approach is commonly deployed as part of a server consolidation strategy. Interestingly enough, it might also be used to support client consolidation.
- Partitioned operating systems this software segments a single operating system so multiple stacks of software can share the same operating system without interfering with one another or even knowing that other stacks or Zones exist. As with virtual machine software, this type of software is often part of a server consolidation strategy.
- Each of these approaches abstract processing away from the underlying hardware but, the goals for the use of this software varies depending upon the type of software in question. High performance computing software is focused on getting things done faster with the help of more brooms, err, I mean more systems. High availability/clustering software is focused on agility, scalability and reliability. Virtual machine software and partitioned operating systems are two different ways to optimize the use of systems.
- Consolidation in the past, the scalability and reliability of some operating systems made it seem necessary to assign each application or service to its own system. Other, more mature, operating systems were able to support many different applications or services without having any scalability or reliability problems. As machines become more and more powerful, this approach wasted much of the computing power available. When utilization could be measured as only hitting the low single digits or in the low double digits, organizations started to realize they could do better. Virtual machine software or a partitioned operating system made it possible for multiple applications or services to reside on a single system. This approach clearly could reduce the number of systems needed, the amount of power required, the heat produced and, as a wonderful side effect, would reduce some of the operational and administrative costs of physical systems.

- Workload management high performance computing software and high availability/clustering software often making it possible for work to be moved from machine to machine based upon service level agreements or upon the workload being imposed by users on the network. This is also seen when workload management software is combined with either virtual machine software or partitioned operating system software. Since the link between the application and the actual physical system has been broken software, systems can be selected based upon what's most appropriate for the task at hand, the time of day, the phase of the moon, or some other important criteria. Fewer high performance, costly, power hungry and heat producing systems might be required. Less costly systems that consume less power and produce much less heat could be used when and wherever possible. This approach could help organizations reduce the hardware costs and consume less power.

As with other forms of virtualization, this is not new. It's been an important part of most datacenters for the past 30 years or so. Suppliers such as HP, IBM, Sun and others have made this software available for years. Suppliers such as Cassatt, DataSynapse, Scalent, SteelEye, Surgient, VMware, XenSource and quite a few others offer this technology on industry standard, high volume clients, servers and blades [10].

Computer virtualization refers to the abstraction of computer resources, such as the process of running two or more logical computer systems on one set of physical hardware. The concept originated with the IBM mainframe operating systems of the 1960s, but was commercialized for x86-compatible computers only in the 1990s. With virtualization, a system administrator could combine several physical systems into virtual machines on one single, powerful system, thereby unplugging the original hardware and reducing power and cooling consumption. Virtualization can assist in distributing work so that servers are either busy, or put in a low power sleep state. Several commercial companies and open-source projects now offer software packages to enable a transition to virtual computing. Intel Corporation and AMD have also built proprietary virtualization enhancements to the x86 instruction set into each of their CPU product lines, in order to facilitate virtualized computing [10] [14].

4.2. Terminal servers

Terminal servers have also been used in green computing. When using the system, users at a terminal connect to a central server; all of the actual computing is done on the server, but the end user experiences the operating system on the terminal. These can be combined with thin clients, who use up to 1/8 the amount of energy of a normal workstation, resulting in a decrease of energy costs and consumption. There has been an increase in using terminal services with thin clients to create virtual labs. Examples of terminal server software include Terminal Services for Windows and the Linux Terminal Server Project (LTSP) for the Linux operating system.

4.3. Shared memory

In computing, shared memory is memory that may be simultaneously accessed by multiple programs with an intent to provide communication among them or avoid redundant copies.

In computer hardware, shared memory refers to a (typically) large block of random access memory that can be accessed by several different central processing units (CPUs) in a multiple-processor computer system.

A shared memory system is relatively easy to program since all processors share a single view of data and the communication between processors can be as fast as memory accesses to a same location.

The issue with shared memory systems is that many CPUs need fast access to memory. Resource usage is minizized.

4.4. Power management

The Advanced Configuration and Power Interface (ACPI), an open industry standard, allows an operating system to directly control the power-saving aspects of its underlying hardware. This allows a system to automatically turn off components such as monitors and hard drives after set periods of inactivity. In addition, a system may hibernate, where most components (including the CPU and the system RAM) are turned off. ACPI is a successor to an earlier Intel-Microsoft standard called Advanced Power Management, which allows a computer's BIOS to control power management functions.

Some programs allow the user to manually adjust the voltages supplied to the CPU, which reduces both the amount of heat produced and electricity consumed. This process is called undervolting. Some CPUs can automatically under volt the processor

depending on the workload; this technology is called "SpeedStep" on Intel processors, "PowerNow!"/"Cool'n'Quiet" on AMD chips, LongHaul on VIA CPUs, and LongRun with Transmeta processors.

4.5. Storage Management

Smaller form factor (e.g. 2.5 inch) hard disk drives often consume less power per gigabyte than physically larger drives. Unlike hard disk drives, solid-state drives store data in flash memory or DRAM. With no moving parts, power consumption may be reduced somewhat for low capacity flash based devices. Super Talent's 2.5" IDE Flash hard drive - The Tech Report - Page 13Power Consumption - Tom's Hardware: Conventional Hard Drive Obsoletism. Samsung's 32 GB Flash Drive Previewed
In a recent case study, Fusion-io, manufacturers of the world's fastest Solid State Storage devices, managed to reduce the carbon footprint and operating costs of MySpace data centers by 80% while increasing performance speeds beyond that which had been attainable via multiple hard disk drives in Raid 0, MySpace was able to permanently retire several of their servers, including all their heavy-load servers, further reducing their carbon footprint [6].

As hard drive prices have fallen, storage farms have tended to increase in capacity to make more data available online. This includes archival and backup data that would formerly have been saved on tape or other offline storage. The increase in online storage has increased power consumption. Reducing the power consumed by large storage arrays, while still providing the benefits of online storage, is a subject of ongoing research.IBM chief engineer talks green storage, SearchStorage – TechTarget [3] [10].

4.6. Video card

A fast GPU may be the largest power consumer in a computer [17]. X-bit labs: Faster, Quieter, Lower: Power Consumption and Noise Level of Contemporary Graphics Cards
Energy efficient display options include:

- No video card - use a shared terminal, shared thin client, or desktop sharing software if display required.
- Use motherboard video output - typically low 3D performance and low power.
- Select a GPU based on low idle power, average wattage or performance per watt.

4.7. Display

CRT monitors typically use more power than LCD monitors. They also contain significant amounts of lead. LCD monitors typically use a cold-cathode fluorescent bulb to provide light for the display. Some newer displays use an array of light-emitting diodes (LEDs) in place of the fluorescent bulb, which reduces the amount of electricity used by the display. Fluorescent back-lights also contain mercury, whereas LED back-lights do not [10].

4.8. Computer multitasking

In computing, multitasking is a method where multiple tasks, also known as processes, share common processing resources such as a CPU. In the case of a computer with a single CPU, only one task is said to be running at any point in time, meaning that the CPU is actively executing instructions for that task. Multitasking solves the problem by scheduling which task may be the one running at any given time, and when another waiting task gets a turn. The act of reassigning a CPU from one task to another one is called a context switch. When context switches occur frequently enough the illusion of parallelism is achieved. Even on computers with more than one CPU (called multiprocessor machines), multitasking allows many more tasks to be run than there are CPUs.

Operating systems may adopt one of many different scheduling strategies, which generally fall into the following categories:

- In multiprogramming systems, the running task keeps running until it performs an operation that requires waiting for an external event (e.g. reading from a tape) or until the computer's scheduler forcibly swaps the running task out of the CPU. Multiprogramming systems are designed to maximize CPU usage.
- In time-sharing systems, the running task is required to relinquish the CPU, either voluntarily or by an external event such as a hardware interrupt. Time sharing systems are designed to allow several programs to execute apparently simultaneously. The expression 'time sharing' was usually used to designate computers shared by interactive users at terminals, such as IBM's TSO, and VM/CMS

- In real-time systems, some waiting tasks are guaranteed to be given the CPU when an external event occurs. Real time systems are designed to control mechanical devices such as industrial robots, which require timely processing.

The term time-sharing is no longer commonly used, having been replaced by simply multitasking, and by the advent of personal computers and workstations rather than shared interactive systems [4] [5] [6].

4.9. Parallel Processing in Computers

The simultaneous use of more than one CPU or processor core to execute a program or multiple computational threads are considered as parallel processing . Ideally, parallel processing makes programs run faster because there are more engines (CPUs or cores) running it. In practice, it is often difficult to divide a program in such a way that separate CPUs or cores can execute different portions without interfering with each other. Most computers have just one CPU, but some models have several, and multi-core processor chips are becoming the norm. There are even computers with thousands of CPUs.

With single-CPU, single-core computers, it is possible to perform parallel processing by connecting the computers in a network. However, this type of parallel processing requires very sophisticated software called distributed processing software. Note that parallelism differs from concurrency. Concurrency is a term used in the operating systems and databases communities which refers to the property of a system in which multiple tasks remain logically active and make progress at the same time by interleaving the execution order of the tasks and thereby creating an illusion of simultaneously executing instructions. Page text Parallelism, on the other hand, is a term typically used by the supercomputing community to describe executions that physically execute simultaneously with the goal of solving a problem in less time or solving a larger problem in the same time. Parallelism exploits concurrency.

Parallel processing is also called parallel computing. In the quest of cheaper computing alternatives parallel processing provides a viable option. The idle time of processor cycles across network can be used effectively by sophisticated distributed computing software [4] [5].

4.10. Parallel computing

Parallel computing is a form of computation in which many calculations are carried out simultaneously, operating on the principle that large problems can often be divided into smaller ones, which are then solved concurrently ("in parallel"). There are several different forms of parallel computing: bit-level, instruction level, data, and task parallelism. Parallelism has been employed for many years, mainly in high-performance computing, but interest in it has grown lately due to the physical constraints preventing frequency scaling. As power consumption (and consequently heat generation) by computers has become a concern in recent years, parallel computing has become the dominant paradigm in computer architecture, mainly in the form of multicore processors.

Parallel computers can be roughly classified according to the level at which the hardware supports parallelism—with multi-core and multi-processor computers having multiple processing elements within a single machine, while clusters, MPPs, and grids use multiple computers to work on the same task. Specialized parallel computer architectures are sometimes used alongside traditional processors, for accelerating specific tasks. Parallel computer programs are more difficult to write than sequential ones, because concurrency introduces several new classes of potential software bugs, of which race conditions are the most common. Communication and synchronization between the different subtasks are typically one of the greatest obstacles to getting good parallel program performance. The maximum possible speed-up of a program as a result of parallelization is observed as Amdahl's law. Computer software has been written for serial computation. To solve a problem, an algorithm is constructed and implemented as a serial stream of instructions. These instructions are executed on a central processing unit on one computer. Only one instruction may execute at a time, after that instruction is finished, the next is executed.

Parallel computing, on the other hand, uses multiple processing elements simultaneously to solve a problem. This is accomplished by breaking the problem into independent parts so that each processing element can execute its part of the algorithm simultaneously with the others. The processing elements can be diverse and include resources such as a single computer with multiple processors, several networked computers, specialized hardware, or any combination of the above. Frequency scaling was the dominant reason for improvements in computer performance from the mid-1980s until 2004. The runtime of a program is equal to the number of instructions multiplied by the average time per instruction. Maintaining everything else constant, increasing the clock frequency decreases the average time it takes to execute an instruction. An increase in frequency thus decreases runtime for all computation-bounded programs.

However, power consumption by a chip is given by the equation P = C × V2 × F, where P is power, C is the capacitance being switched per clock cycle (proportional to the number of transistors whose inputs change), V is voltage, and F is the processor frequency (cycles per second). Increases in frequency increase the amount of power used in a processor. Increasing processor power consumption led ultimately to Intel's May 2004 cancellation of its Tejas and Jayhawk processors, which is generally cited as the end of frequency scaling as the dominant computer architecture paradigm.

Moore's Law is the empirical observation that transistor density in a microprocessor doubles every 18 to 24 months. Despite power consumption issues, and repeated predictions of its end, Moore's law is still in effect. With the end of frequency scaling, these additional transistors (which are no longer used for frequency scaling) can be used to add extra hardware for parallel computing [4] [5].

4.11. Amdahl's law and Gustafson's law.

Optimally, the speed-up from parallelization would be linear—doubling the number of processing elements should halve the runtime, and doubling it a second time should again halve the runtime. However, very few parallel algorithms achieve optimal speed-up. Most of them have a near-linear speed-up for small numbers of processing elements, which flattens out into a constant value for large numbers of processing elements.

The potential speed-up of an algorithm on a parallel computing platform is given by Amdahl's law as shown in figure 1, originally formulated by Gene Amdahl in the 1960s. It states that a small portion of the program which cannot be parallelized will limit the overall speed-up available from parallelization. A program solving a large mathematical or engineering problem will typically consist of several parallelizable parts and several non-parallelizable (sequential) parts. If   is the fraction of running time a sequential program spends on non-parallelizable parts, then:

$$S = \frac{1}{\alpha}$$

(1)

S is the maximum speed-up with parallelization of the program. If the sequential portion of a program accounts for 10% of the runtime, we can get no more than a 10× speed-up, regardless of how many processors are added. This puts an upper limit on the usefulness of adding more parallel execution units. "When a task cannot be partitioned because of sequential constraints, the application of more effort has no effect on the schedule. The bearing of a child takes nine months, no matter how many women are assigned.
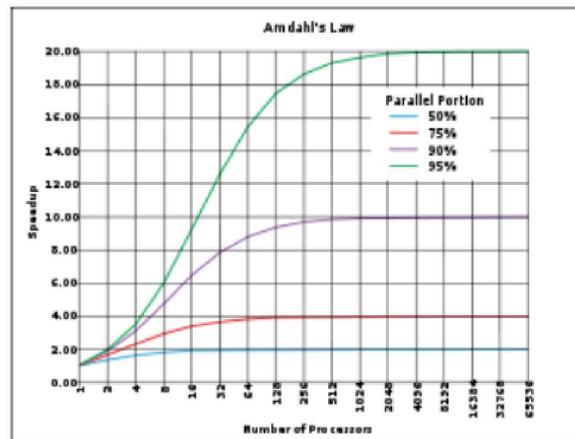

Figure 1. A graphical representation of Amdahl's law.

The speed-up of a program from parallelization is limited by how much of the program can be parallelized. For example, if 90% of the program can be parallelized, the theoretical maximum speed-up using parallel computing would be 10x no matter how many processors are used.

Gustafson's law is another law in computing, closely related to Amdahl's law. It states that the speedup with P processors is

$$S(P) = P - \alpha(1 - P)$$

(2)

Amdahl's law assumes a fixed problem size and that the running time of the sequential section of the program is independent of the number of processors, whereas Gustafson's law does not make these assumptions [10].

4.12. Software Pipeline

In software engineering, a pipeline consists of a chain of processing elements (processes, threads, co-routines, etc..), arranged so that the output of each element is the input of the next. Usually some amount of buffering is provided between consecutive elements. The information that flows in these pipelines is often a stream of records, bytes or bits.

The concept is also called the pipes and filters design pattern. It was named by analogy to a physical pipeline.

Pipelines are often implemented in a multitasking OS, by launching all elements at the same time as processes, and automatically servicing the data read requests by each process with the data written by the upstream process. In this way, the CPU will be naturally switched among the processes by the scheduler so as to minimize its idle time. In other common models, elements are implemented as lightweight threads or as co-routines to reduce the OS overhead often involved with processes. Depending upon the OS, threads may be scheduled directly by the OS or by a thread manager. Co-routines are always scheduled by a co-routine manager of some form.

Usually, read and write requests are blocking operations, which means that the execution of the source process, upon writing, is suspended until all data could be written to the destination process, and, likewise, the execution of the destination process, upon reading, is suspended until at least some of the requested data could be obtained from the source process. Obviously, this cannot lead to a deadlock, where both processes would wait indefinitely for each other to respond, since at least one of the two processes will soon thereafter have its request serviced by the operating system, and continue to run [4] [5].

5. ARCHITECTURE OF OPTIMIZED OPERATING SYSTEM TOWARDS GREEN COMPUTING

There are many ways to optimize a client environment to save energy and reduce environmental impact barriers. The Architecture of optimized operating system towards green computing is shown as in Figure 2. It has three layers namely Physical hardware services layer, Operating system services layer and Application services layer.
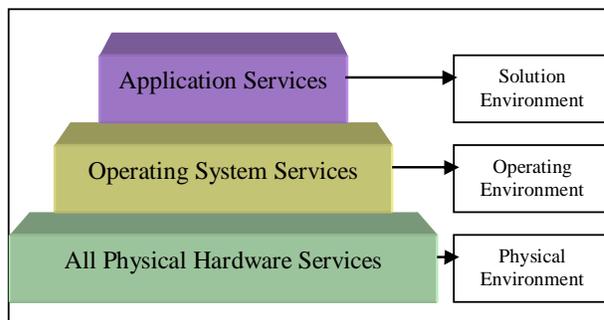
Figure 2. Three-Tier Architecture for green computing

Each layer has its own services and provides services to the above layer. Each layer has to work coordinately with other layers. In hardware layer we can not do anything, because it was made by hardware manufacturers. But we can apply green computing in both Operating System Services and application services.

5.1. Physical Services environment

Acquiring an Energy Star 4.0 system which recognizes ACPI 3.0 power management capabilities from Windows Vista allows the operating system to manage power for processors (and multiple cores), attached devices, and allows advanced capabilities for hibernation and sleep. Also, administrators can use group policies to throttle back the maximum CPU load to reduce energy consumption when needed.

5.2. Operating System services environment

To leverage advanced energy efficient computer hardware, the operating environment must be capable of using the new ACPI 3.0 hardware functionality, and it must deliver advanced performance and power management capabilities for the user and administrator. The operating execution environment is the configuration and standardization of the operating system and the supporting utilities. It is crucial to leverage the most aggressive power savings capabilities possible while accomplishing computing goals of the organization. When setting up a standardized configuration, minimize the number of running system services to reduce energy consumption. Operating system also provides services for both the upper and lower layers. Resources allocation and efficient energy saving are the key services.

5.3. Application services environment

To reduce the amount of resources a client must use to run a fully installed application, architect teams can leverage client application virtualization from solutions such as Microsoft's Application Virtualization solution for client systems and remote client interface solutions from new Windows Server 2008 Terminal services. However, this takes careful planning and works in a focused set of scenarios. Power-aware Windows Presentation Foundation (WPF) applications can use less power intensive presentation experience based on the power state of the client. Also, some are aggregating application development best practices to minimize energy resource consumption on the client environment [1] [2] [7].

6. GUIDELINES TO ENABLE COMPUTER'S POWER MANAGEMENT FEATURES IN OPERATING SYSTEMS

6.1. Windows Vista

1. Select Start > Settings > Control Panel from the Start Menu.
2. The Control Panel has two views. If in the Theme View, select System Maintenance and then Power Options.
3. If you are in Classic View, simply double click on Power Options.
4. Once you are in Power Options, click Create a Power Plan link.
5. Enter a name for your new plan in the Plan Name field and click next.
6. Set Turn Off the Display for 15 minutes of less.
7. Set Put Computer to Sleep for 30 minutes or less.
8. Click the Create button to create and activate your new Power Plan.
9. To make additional changes, click on Change Plan Settings underneath the name of the plan just created.
10. Click on Change Advanced Power Settings. Here you will be able to enable or set values for a variety of Power Management options.
11. Set Turn Off Hard Disks After for 15 minutes or less and Click OK.
12. Click Save Changes.

6.2. Windows XP/2000

1. Select Start > Settings > Control Panel from the Start menu.
2. In the Display Properties window, click the Screen Saver tab, then click the Power button in the Energy Saving Features of Monitor box.
3. Under Power Schemes, select Home/Office Desk.
4. Set Turn Off Monitor for 15 minutes or less.
5. Set Turn Off Hard Disks for 15 minutes or less.
6. Set System Standby for 30 minutes or less and Click OK.

6.3. Windows 98.

1. Select Start > Settings > Control Panel from the Start Menu.
2. Double-click the Display icon in the Control Panel.
3. In the Display Properties window, click the Screen Saver tab, then click the Settings button in the Energy Saving Features of Monitor box.
4. Under Power Schemes, select Home/Office Desk.

5. Set Turn Off Monitors for 15 minutes or less.
6. Set Turn Off Hard Disks for 15 minutes or less.
7. Set System Standby for 30 minutes or less.
8. Click OK.

6.4. Mac OS X.

1. Click the apple symbol in the upper left corner of screen.
2. Go to System Preferences, click Show All, and select Energy Saver from the Hardware row.
3. Click on the Show Details button. • Select Power Adapter under Settings For.
4. Using the slider, set Put the Computer to Sleep When It Is Inactive for 30 minutes or less.
5. Check the box for Use Separate Time to Put the Display to Sleep and set the time for 15 minutes or less, using the slider.
6. Check Put the Hard Disk to Sleep When Possible.

6.5. Mac OS 9

1. Click on the apple symbol (Apple Menu) in the upper left corner of your screen.
2. Go to Control Panels and select Energy Saver.
3. Click on the Sleep Setup tab.
4. Using the slider, set Put the System to Sleep Whenever It's Inactive for 30 minutes or less.
5. Check the box for Separate Timing for Display Sleep and set the time for 15 minutes or less, using the slider.
6. Check the box for Separate Timing for Hard Disk Sleep and set the time for 15 minutes or less, using the slider.

Each OS has different method of power management scheme. We have to use it in proper way to save energy. Power reduction and less consumption are the key points for green computing. Green user has to follow the above steps to keep the environment as green environment [4] [5] [6].

7. CONCLUSION

Green computing not only reducing cost, but also saves energy and optimize the resource utilization. Reducing Energy Consumption is the key goal of Green Computing. If public and private sectors takes more interest in green computing, definitely we can save our environment and maintain green environment. Every body in this world should be either a green computing person or a green user in order to keep our environment as green environment. Green computing can be easily adapted to any computing to make green environment. Now a days advanced operating system developers are very much interested to wards green computing to attract customer as well as to protect environment by saving power energy.

References

1. Sheikh, R.A. and Lanjewar, U.A. Green Computing- Embrace a Secure Future", International Journal of Computer Applications Vol. 10, No. 4 (2010) 22-27.
2. Appasami. G and Suresh Joseph.K. Design and Architecture of Secured Flash Solid State Device File System, IETECH Journal of Communication Techniques, Vol.4, No. 2 (2009) 48-54.
3. Chakraborty, P., Bhattacharyya, D., Nargiza, S. and Bedajna, S. Green computing: Practice of Efficient and Eco-Friendly Computing Resources. International Journal of Grid and Distributed Computing, Vol. 2, No. 3 (2009).
4. Harris, J. Green Computing and Green IT Best Practices on Regulations and Industry Initiatives, Virtualization, Power Management, Materials Recycling and Telecommuting. Emereo Publishing, 2008.
5. Lamb, J. The Greening of IT: How Companies Can Make a Difference for the Environment. IBM Press, 2009.
6. Velte, T., Velte, A. and Elsenpeter, R. Green IT: Reduce Your Information System's Environmental Impact While Adding to the Bottom Line. McGraw-Hill Osborne Media, 2008.
7. Leonhard, W., Murray, K. Green Home Computing For Dummies. For Dummies, 2009.
8. Di Kacangin, http://dikacang.in/download/green-computing (July 1, 2011).
9. Google Sites, http://sites.google.com/site/greencomputingproject/book-journal (July 1, 2011).
10. Wikipedia. http://en.wikipedia.org/wiki/Green_computing (July 1, 2011).

Appasami.G and Suresh Joseph.K / Optimization of Operating Systems towards Green Computing, IJCOPI Vol. 2, No. 3, Sep-Dec 2011, pp. 39-51.

11. MyTechSupportStore   http://www.mytechsupportstore.com/articles/green-computing-a-reformed-tech-support.html (July 1, 2011).
12. Article Directory http://www.internetionalmedia.com/Article/The-Advantages-of-Green-Computing-For-Businesses-and-Individuals/6007 (July 1, 2011).
13. Aspen Systems, Inc. http://www.aspsys.com/content/green_computing (July 1, 2011).
14. ZDNET. http://www.zdnet.com/blog/virtualization/processing-virtualization-and-green-computing/143 (July 1, 2011).
15. http://green-computing.co.tv/ (July 1, 2011).
16. School of Computer Science and Engineering, Bharathidasan University. http://www.csbdu.in/MTH10101.php (July 1, 2011).
17. X-bit labs. http://www.xbitlabs.com/articles/video/display/power-noise.html (July 1, 2011).