Scientific Journal of Impact Factor (SJIF): 5.71
e-ISSN (O): 2348-4470
p-ISSN (P): 2348-6406

International Journal of Advance Engineering and Research Development

Volume 5, Issue 04, April -2018

GREEN COMPUTING - AN INSIDE ANALYSIS

Vandna Dahiya

Maharshi Dayanand University, Haryana,

Abstract:- *Green computing is an all-encompassing term that covers computing practices from the birth to death of computing devices. The emphasis is two-fold - the design, manufacturing, use and disposal of computers and related equipment as well as networking and communication equipment in a carbon-neutral and environmentally safe way. Of the first of these objectives, most corporations today have to adhere to environmental regulations that limit their carbon footprint. Because more energy consumption by data centers, which are notorious as the SUVs of the tech world, directly translates into a larger carbon footprint for the corporation, there is tremendous incentive for best practices to limit power consumption of data centers. The second is limiting and reversing inefficiencies in design, design integration and energy consumption. In gigantic power systems as well as for fixed computing devices in homes, this would translate into hundreds of dollars saved on decreased power consumption. For mobile devices, this would mean increased battery lifespan and smaller battery sizes, which would leave space on these devices for more functions. This research, while extremely abstract, has given computer and IT scientists a definite head start in the broader game to reduce carbon footprints for the IT industry.*

Keywords: *Green computing; Cloud computing; Green products; Computer architecture.*

I. Introduction

Talking about data centers is focusing on the big units in the strategic game of green computing. A lot of green computing analysis begins with power management at the level of components, and that involves a lot of basic computer science. The US science funding National Science Foundation has funded two workshops in power management as part of a larger effort to solve basic questions like exactly how much power one really need to get some basic computation unit task done.

There are also other fundamental factors. Other industries are not as influenced by their products' performances on the green parameters as the IT industry is. For example, mobile devices such as cellphones and laptops whose battery lives are greatly improved if more attention is paid to their power management style will find a greater market. Thermal constraints on microprocessors, which have grown tremendously in computing power with every passing generation, are also attributed as a big factor. This impetus for research has given experts working with power management a number of tools to work with on the central processing unit (CPU). Although much of this effort is designed for cross-platforms that work on servers, laptops and computers, the default power management protocols in most data centers remain the same. There are some independent efforts - software called Granola designed by Kirk Cameron, who is associate professor at Virginia Polytechnic Institute, uses predictive algorithms to scale frequency and voltage in real time. According to some estimates, back power usage can be cut in systems by 2-18%, working best by cutting back power when the system is not in use. The dynamic voltage and frequency scaling (DVFS) technology is apparently very stable, and, according to Cameron, is "low hanging fruit" which is used in the Green IT business for best practices that are easy to enforce but produce quick and tangible results.

II. Multicore architecture

While DVFS is quite popular with the green computing community, most academics including the creator of DVFS agree that with the advent of multicore computing, the basic way of looking at power management and optimization should change. What is now understood is that the correct way to look at speed and frequency scaling in single processor systems will not hold for parallel computing. This includes the famous Amdahl's law, which states that any parallelized algorithm can be sped up at the percentage of a given task within that program not run serially. Multicore architecture will present new challenges for power management including brand new algorithmic issues, which will require fresh ideas in core computer science.

III. New Ideas

What green power management can really learn from are sensor networks and embedded systems that have faced power constraints for a longer time. What has been pointed out time and again is the fact that most processors are idle for most of the time, and are power inefficient while being idle. Many microprocessors are gauged for performance by how well they do when they are active, but power management efficiency is better measured by how well they do when they are waiting for instructions. Shutting off processors while they await commands is a valid theory, and there are two schools of thought as far as that goes - Token Ring or a similar time division technique, or the Carrier Sense Multiple Access (CSMA) approach that will look similar to the Ethernet approach. The CSMA approach is intuitive in the sense that it allows processor nodes to sense if the network they are transmitting to be idle before transmitting. In all these approaches, what is fundamentally in question is how well the system sits idle, especially in the modern computer architecture that will not adhere to the usual laws of computer science.

However, one drawback for research in dynamically provisioned microprocessors is their remote benefits that are apparent to the end-consumers. A typical home in Western Europe or US will care very deeply about bulk power management that saves these homes thousands of dollars every year via more efficient distribution systems and cooling architectures. However, much of chip-level energy efficiency is now being linked with bulk efficiency, especially with discovery and predictive algorithms that combine and integrate functions in a single chip to achieve both objectives with the same resources.

IV. **Where is this energy wasted?**

When talking power management, whether that is at the level of supercomputers, data centers, PCs, laptops or mobile devices, the first question to ask you is where the sources of energy waste are. Planners, engineers and architects should be well versed with this, but as it turns out, the rise of green computer science has exposed gaping holes in our understanding of these seemingly basic ideas. RN Mayo was the first person to pose and measure the question of how much energy is actually consumed in performing the seemingly simple tasks of, say, listening to music or sending a text message or an email or if you are using a cell phone, the energy spent in making a phone call. What he found was that the answer was not unambiguous but rather very strongly dependent on the platform used. What was striking in his results was that the differences in energy consumed was often different by several orders, even when one tried to duplicate the experiment on the best and worst performing devices. The green computing theories of today attribute the inefficiencies in some systems not to just the poor power management architecture but also to the several complex tradeoffs that are designed into systems so that their idle state is also inefficient. None of the computing circuits in question exist in isolation. Therefore, it is impossible to talk about the power management of one circuit without taking into account the influence that is exerted by all the other circuits. Most systems are designed, for instance, to handle the most aggressive workload or recover from the most fatal errors, which of course lead to redundancies that of course are designed to handle failure but lead to over allotment of precious resources when the demands are mild and non-fatal. As another example, one can talk about the prevalence of modularity in software engineering and design. While it may help to solve real life problems faster by breaking it up into smaller problems, what this means for power management is that the various modules, even if engineered to be efficient on their own, end up creating competing demands that lead to gross inefficiencies overall. The way out of this kind of mess is general-purpose systems that seek to integrate all common tasks. These systems have a great market appeal - mobile phones that integrate cameras, iPods, email and all kinds of PC functions have been a great hit in recent times.

V. **Smart Planning**

One stunning piece of data that has driven much action in green computing is that the average utilization levels in data centers across the US and Europe can be as measly as 10-30%. While data centers eat up billions of dollars worth of power, the same problem afflicts mobile phones whose circuits are designed to handle the worst case load in terms of functions but they sit idle for most of their time. What researchers are proposing now are designs that are smart in terms of how they scale their energy and power consumption with the amount of resources that are being brought into play on the microprocessor, so idle circuits do not end up consuming as much power as they do now. However, this requires a greater degree of predictive and smart algorithms from engineers.

Overprovisioning is also a feature when you examine plans for future growth of typical installations like data centers. Most facilities are built for three-to-five year depreciation cycles, and planners have to ensure there is sufficient capacity to encompass future growth of the facility. This often leads to significant inefficiencies - if a typical system that is built to cool for a total power of 1 megawatt is operated at a tenth of that power, it will experience significant inefficiencies than if the same system was provisioned to cool for say 100 kilowatts of power.

The first step in a good power management scheme that belongs to the next generation is consideration of interactions across structures in a given hardware setup. For example, a CPU has functionality that is divided between different stacks such as CPU, memory, chipset, hard disk and so on, and power management is considered in each of these units. The new school of thought aims for a holistic approach where architects and engineers look at the way the different stacks interact with each other in their power management schemes. One of the fundamental problems in modulation of system functionality is that each stack will make worst-case assumption about the functionality of the other stacks. This leads to obvious considerable waste. To solve this and the other problems, planners have to aim for a system that allows information to be exchanged within the various stacks so overprovisioning is avoided and inefficiencies are prevented.

VI. Tethered Systems

While most of the discussion here pertained to mobile as well as tethered systems, there are some problems that pertain specifically to fixed devices that produce unique inefficiencies. Much of tethered systems are designed such that the aim is to achieve the highest performance that they can, without regard to the power consumed or the battery life depleted. Processor-architecture designs have historically implemented design solutions that achieve design incremental improvements that do not match up with the additional power that goes into achieving said increments in performance. To illustrate this further, designs with user interfaces that focus squarely on energy delay are cited to be an example of one strategy that is much more energy efficient than designs that focus squarely on performance, while at the same time offering only marginally worse performance than those brute force designs.

VII. Reducing Waste: Some Suggestions

In this last section, we want to round up with some very broad suggestions as to where waste can be cut back in power management in computing systems. Simply using more power efficient systems is an obvious solution. Replacing conventional hard disks with the more energy efficient non-volatile memory drives will have a direct impact on power consumption, especially in large facilities such as data centers. Predictive algorithms for energy and frequency scaling in circuits that go unused for long periods but are built for high demand performances will also go a long way in ensuring they continue to perform near their peak efficiency in terms of power management. One innovative way is to club multiple tasks into a single energy task. The best illustration of this is in the reading of a disk drive - if many disk reads are coalesced onto a single spin, it will reduce the energy consumption by many orders.

Power management is still in its infancy and much of research being carried out is through academic projects funded by the EPA and EU agencies. While there is tremendous future incentive for corporations to invest in effective power management, it is to be understood that predictive algorithms for multicore processors will require a change in the way concepts are conceived and adopted in green computing.

References

1. Goth G., Chipping Away at Greenhouse Gases, doi: 10.1145/1897816.1897823.
2. Amdahl, Gene M. (1967). "Validity of the Single Processor Approach to Achieving Large-Scale Computing Capabilities", AFIPS Conference Proceedings (30): 483–485. doi: 10.1145/1465482.1465560.
3. Mayo, R.N. and Ranganathan, P. "Energy Consumption In Mobile Devices: Why Future Systems Need Requirements-Aware Energy Scale-Down," In Proceedings of the Workshop on Power-Aware Computing Systems, pp: 26–40, 2003
4. Parthasarthy Ranganathan, "Recipe for Efficiency: Principles of Power-Aware Computing," doi: 10.1145/1721654.1721673.
5. Patrick Kurp, "Green Computing," Communications of the ACM, Volume: 51, Issue: 10, 2008, pp: 1-19 doi:10.1145/1400181.1400186.
6. R.Varatharajan, M. Sundharajan, Xiao-ZhiGao, Jimming Wen, "Introduction to special issue on Green Computing Techniques for Wireless Sensor Networks," Sustainable Computing: Informatics and Systems, ScienceDirect, 2018, doi: 10.1016/j.suscom.2018.04.001
7. Ghossoon M.W. AI Sadoon, Hussan A. Makki, Ali Radih Saleh, "Green computing system, health and secure environment management system," 4th IEEE International Conference on Engineering Technologoes and Applied S iences, Bahrain, 2017, doi: 10.1109/ICETAS.2017.8277847
8. Seema Rawat, Praveen Kumar, Seemant Sagar, Iqbal Singh, Kartikay Garg, "An analytical evaluation of challenges in Green cloud computing", International conference on infocom technologies and unmanned systems, UAE, 2017, doi:10.1109/ICTUS.2017.8286031
9. Aanchal, Sushil Kumar, Omparkash Kaiwartya, Abdul Hanan Abdullah, "Green computing for wireless sensor networks: Optimization and Huffman coding approach," Springer, Volume: 10, Issue: 3, pp: 592-609, 2017
10. Laurent Lefevre, Jean-Marc Pierson, "Introduction to special issue on sustainable computing for ultrascale computing," Sustainable Computing: Informatics and Systems, ScienceDirect, Volume: 17, pp: 25-26, 2018

