A Mechanism for Load Proportional Energy Use in Wireless Local Area Networks

Jie Chen and Biplab Sikdar
Department of Electrical, Computer and Systems Engineering
Rensselaer Polytechnic Institute, Troy, NY, 12180
Email: {chenj24,sikdab}@rpi.edu

***Abstract*—From the perspective of energy efficiency, a major drawback of existing wireless networks is that the power consumption rate of the networking hardware stays at levels close to the maximum, even when the offered traffic load is low. To address this issue, this paper presents a mechanism for load proportional energy usage in wireless networks. The proposed mechanism is based on changing the operating clock frequency of the network interface cards as a function of the offered load, in order to reduce to energy consumption. To this end, we propose a frequency selection mechanism that tries to ensure the stability of the queues while minimizing the energy consumption. The proposed mechanism has been evaluated through simulations.**

I. Introduction

The energy efficiency of various aspects of information and communications technology (ICT) has become a topic of increasing concern in recent years. Traditionally, networking hardware was not designed with power management in mind and as a result, their power consumption rate stays very close to the maximum rating, irrespective of the network load (for example, WiFi access points [1]). Load proportional energy usage, where the energy consumed by the networking devices is proportional to the traffic or load they carry, is thus a promising approach to improve the power efficiency of network equipment. For wireless networks, the issue of power management is particularly important, given that the terminals may be portable and the batteries have limited storage. This paper presents a methodology for achieving load proportional energy usage in wireless local area networks (WLANs), while maintaining system stability, by developing a strategy for dynamic frequency scaling in the wireless devices.

The design of energy efficient or green wireless networks has received increasing attention in the recent past [2], [3], [4]. Existing approaches to improve the power efficiency of wireless networks broadly belong to the following categories: hardware improvements such as improved amplifier design, sleep scheduling, power control, adjustment of cell sizes, energy efficient routing etc.. However, the rate of energy consumption after the application of these approaches is independent of the load (i.e. traffic) carried by the devices in most cases. For example, if a base station is on, the energy it consumes has a very minor variation with the load. In most cases, for any given network configuration, the rate of power consumption of the devices shows less than 15% variation with the load [1], [5].

This paper addresses the issue of developing a mechanism for facilitating load proportional energy consumption in WLANs. The approach taken by this paper is to vary the hardware's operating speed as a function of the traffic load. The use of variable voltage and frequency settings (also known as voltage and frequency scaling (VFS)) to control the operating characteristics and thus the energy consumption of devices as a function of the workload has been well studied in the area of microprocessors [6], [7]. The fundamental idea of VFS is that when the load of tasks is not heavy, the microprocessor can slow down the rate at which it does work, so as to save energy.

In this paper, we apply the same principles to control the operational characteristics of the functional blocks inside a wireless network interface card in order to conserve energy. However, a key challenge here is to maintain stability of the queues, while ensuring that the energy savings are maximized. When the operating frequency is reduced to save energy, it is accompanied by an increase in the queue length since the device now takes longer to process and transmit a packet. Consequently, the likelihood of the device becoming unstable (i.e. the queues becoming full) increases. In this paper, we present a mechanism that tries to ensure the stability of the queues while mamimizing the energy savings that are possible though dynamic frequency scaling (DFS). The proposed approach is based on observing the queue length at the wireless device and choosing the operating frequency based on the queue length. Simulation results are presented to evaluate the effectiveness of the proposed mechanism. Our results show that the system with the proposed DFS technique outperforms conventional systems in terms of the energy consumed as well as the packet loss rate.

The rest of the paper is organized as follows. Section II presents the related work and the system model assumed in this paper. Section III presents the proposed mechanism for load proportional energy usage. Section IV presents the simulation results to evaluate the proposed mechanism. Finally, Section V concludes the paper.

II. Background and Related Work

In this section we present the related work. We also present the system model assumed in this paper and an overview of the background material.

A. Related Work

The use of VFS to reduce the energy consumption of electronic devices under different scenarios has been widely reported in existing literature. The main focus of the work in this area is to select the operating frequency and/or voltage in order to minimize the energy consumption. In [9], the selection is based on decomposing the workload into off central processing unit (CPU) and on CPU periods. In [6], the selection is done with further consideration of the system performance and the due time of the task. The VFS problem is investigated from the perspective of software level program codes in [15]. The authors decompose the tasks into different program regions and collect the experimental data on energy consumed under different VFS states and a branch and bound method is then used to solve the selection problem. The focus of these papers is specific to the scenario they consider and they are not directly applicable to the problem considered in this paper.

In addition to microprocessors, the technique of dynamic voltage and frequency scaling (DVFS) has also been applied to network devices. In [8], a predictive task queue based energy conservation scheme has been proposed for a specific type of device, the Gigabit Ethernet controller. In [11] a methodology for selecting the operating frequency and voltage based on a prediction of the future traffic load from the previous statistics is presented. The feasibility of implementing DFVS schemes in local area network (LAN) switches is presented in [12] using statistical data collected from real time experiment set-ups. The authors of [12] also propose an algorithm to optimally transition between idle and busy modes. In [13], three schemes have been proposed to select the suitable router frequency so as to minimize the power consumption. The authors of [14] discuss the implementation of DFVS schemes in parallel network processors by modeling the system as a $M^x/D/1/SET$ queuing system, and obtain the best frequency and voltage setting by solving an optimization problem with constraints on packet waiting time and energy consumed. While these papers all have the energy efficiency as their objective, the constraints on the design are different from those assumed in this paper. In contrast to the existing papers, the objective of this paper is to ensure the stability of the queue while maximizing the energy savings that are possible. Also, many of the schemes proposed in literature (e.g. prediction based schemes) are based on the knowledge of the traffic arrival process. Our mechanism does not require prior knowledge of characteristics of the arrival process.

B. System Model

In this paper we consider a general wireless local area network without assuming any particular medium access control (MAC) protocol. We assume that each node in the network has a wireless network interface card. As shown in Figure 1, a generic wireless network interface card typically consists of: functional blocks for radio frequency (RF) transmit and receive operations, MAC and physical (PHY) layers for the wireless network, transmit (TX) and receive (RX) buffers that are controlled by a buffer manager, units for management and Internet Protocol Security (IPSec), and a control unit. In additional, the network card either has functional blocks for communicating with the host computer over a bus (e.g. Peripheral Component Interconnect (PCI) bus) or, in the case of an access point, PHY and MAC layers for the wired network interface (e.g. Ethernet).

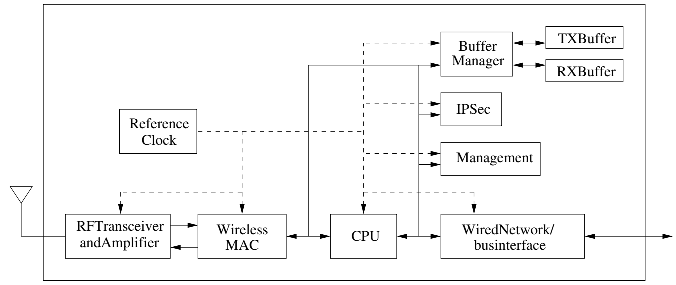

Fig. 1. Simplified block diagram of wireless network interface.

Traditionally, the clock rate and the voltage levels supplied to the various functional blocks are kept constant, resulting in a fairly constant level of power consumption across various traffic loads. With DVFS, the clock frequency and/or the supply voltages to some (or all if possible) of the blocks is varied in order to reduce the energy consumption. This paper assumes that the energy consumption of the blocks is proportional to the square of the operating voltage and directly proportional to the operating frequency [9].

While both the frequency and the voltage may be varied to achieve energy savings, practical implementation of voltage scaling is challenging and associated with considerable overhead [8]. In particular, voltage scaling is comparatively slower due to the response time of the voltage control loop and also suffers from the limitations of current DC-DC converters (e.g. load current change after mode transitions that lead to output voltage drops). Consequently, in this paper we only consider dynamic frequency scaling.

III. Framework for Load Proportional Energy Use

This section presents the proposed framework for load proportional energy use in wireless networks. The proposed approach is based on a dynamic frequency scaling (DFS) mechanism that uses the current queue length at the device to choose the frequency setting.

An overview of the proposed mechanism for load proportional energy use is shown in Figure 2. The proposed system is based on changing the operating conditions (and thus the energy consumption) of the wireless network interface according to the current network conditions. The proposed system has three major components: (i) a mechanism for selecting the appropriate operating frequency, (ii) hardware support for DFS, and (iii) the functional blocks that make up a wireless interface, such as those shown in Figure 1.

The operation of the proposed system is as follows. The network interface card stores the packets that it wishes to

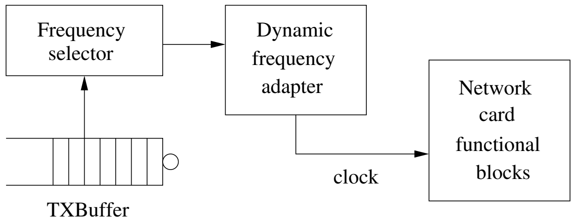

Fig. 2. Overview of the proposed system.

transmit in the TX buffer. Whenever the network interface initiates a new transmission (e.g. when the interface card starts channel access and carrier sensing in IEEE 802.11), it first checks the current length of the TX buffer. Using the queue length, the interface card then uses the mechanism in Section III-A to select the operating frequency that leads to the highest energy savings while ensuring stability of the queue. The selected operating frequency is then provided and used to control the functional block that generates the clock signals. The new clock signals are then used to operate the remaining functional blocks in the network interface card. Note that hardware implementations of dynamic frequency adapters are already in existence. Thus this paper focuses on developing the methodology for selecting the operating frequency and the performance evaluation of the proposed scheme.

A. Frequency Selection

Consider a system where the frequency adapter can provide a set of operating frequencies $\mathcal{F}$, with $|\mathcal{F}| = N$, i.e., we have a set of N frequencies to choose from. Let the frequencies be arranged in decreasing order, $f_1, f_2, \cdots, f_N$, with $f_1 > f_2 > \cdots > f_N$. Let the rate of energy consumption of the network interface card when frequency f_i, $1 \leq i \leq N$, is selected by denoted by P_i. Since the power consumption is directly related to the operating frequency (see Section III-C for the exact relation), we have $P_1 > P_2 > \cdots > P_N$. Let the *reward* r_i associated with the use of frequency f_i be defined as the energy saved when frequency f_i is used, as compared to frequency f_1, i.e.,

$$r_i = P_1 - P_i, \tag{1}$$

for $1 \leq i \leq N$. Note that $r_i < r_2 < \cdots < r_N$.

The frequency selection policy chooses the operating frequency in the following way. If the current length of the TX buffer at the sender is Q, frequency f_i is chosen if $(N-i)\xi < Q \leq (N-i+1)\xi$, where ξ is a fixed positive integer. For $Q \geq N\xi$, f_1 is selected as the operating frequency. This strategy maps queue sizes to the frequency in blocks of ξ and the simplest choice for ξ is $\lfloor Q_{max}/N \rfloor$ where Q_{max} is the maximum queue size. The mapping above ensures that as the queue length becomes large, the network interface card operates at the highest possible frequency in order to clear the backlog and move the system towards stability. Also, when the queue length is small, lower operating frequencies are selected in order to gain energy savings.

B. Stability and Energy Optimality

In this section we discuss the stability and performance (in terms of the energy consumption) of the proposed frequency selection scheme. We start with the definitions and then present an intuitive explanation of the issues related to the stability of the proposed scheme.

A system is called *stable* if the expected queue length at the node is bounded and any frequency selection policy that leads to a stable system is termed a *stable selector*. The packet departure rate equals the packet arrival rate for any stable scheduler. A stable selector Π is called *ϵ-optimal* if no other stable selector can achieve an energy consumption that is less than ϵ from to the energy consumption under Π (i.e. if the energy consumption under Π is P^S, no other selector can achieve an energy consumption of less than $P^S - \epsilon$). Finally, we note that a wireless interface card that does not use use DFS can be considered to use a selector that always selects a single frequency.

Intuitively, the frequency selection policy described in Section III-A maps the queue length to the set of available frequencies (or equivalently, power consumption modes), with smaller queue lengths mapped to lower values of the operating frequency. When the queue length is small, the policy tries to maximize the energy savings by setting the selected frequency to low values. In this case, even if the node has to spend a longer time on the pprocessing and transmission of each packet, the likelihood of the queue becoming unstable is small. On the other hand, the queue is more likely to overflow if the current queue length is large. In such cases, the policy selects the operating frequency to larger values so that the packets are processed and transmitted quickly. Consequently, the queue length reduces and the node moves back towards a stable region.

We conjecture that the proposed frequency selection policy is ϵ-optimal. Also, we note that the proposed policy achieves the stability region since each node can operate at the highest frequency whenever the queue length increases.

C. Energy Model

This section presents a methodology to evaluate the energy consumption of the wireless network interface card.

The use of DVFS for energy savings exploits the fact that the energy consumption of electronic devices is directly related to the operating voltage and frequency, and the effective capacitance of the device. The energy consumption of a device can be written as [10]:

$$P = \alpha C_{eff} V^2 f \tag{2}$$

where α is the switching factor, C_{eff} is the effective capacitance, V denotes the operating voltage and f is the operating frequency. In this paper, as in [10], we further assume that

$$V^2 \propto f. \tag{3}$$

Hence,

$$P \propto f^2. \tag{4}$$

TABLE I
PACKET DEQUEUEING TIME

Index	Interval (s)
1	0.001
2	0.0015
3	0.002
4	0.0025
5	0.003

Then, for operating frequency f_i, (2) can be written as

$$P_i = k{f_i}^2, \quad (5)$$

where k is a constant.

IV. SIMULATION RESULTS

In this section we present simulation results to evaluate the proposed DFS scheme for achieving load proportional energy use. The simulations were conducted in the NS2 simulation platform. The proposed system was implemented in the NS2 platform and support was provided to choose from a set of 5 operating frequencies. An energy model based on the expressions of Section III-C was also implemented. For the implementation, we consider the case where the nodes each have a physical queue where the packets are stored and a network processor that controls the operation of all the functional blocks of the wireless interface.

In our simulations, the nodes used the user datagram protocol (UDP) as the transport layer. Simulations were carried for two types of traffic: constant bit rate (CBR) and variable bit rate (VBR). Each simulation run was for 100 seconds and each reported value is the average of 10 runs. The wireless link bandwidth was set at 10Mbps and the size of each packet was 1000 bytes. Slower operating frequencies lead to a lower energy consumption at the cost of longer processing time per packet. For the five operating frequencies, Table I shows the packet dequeueing time, defined as the time it takes to process the packet and send it to the RF circuitry for transmission, assuming that the channel is available.

We consider the following performance metrics for our simulations:

- The average energy consumed per packet: The average energy per packet is computed as the total energy consumed by a node divided by the number of packets handled by the node (including lost packets). To keep the results independent of the parameters such as device capacitance, we show the results for the energy consumption in terms of the average operating frequency. From Section III-C, the rate of energy consumption of the wireless network interface card is directly dependent on the operating frequency. Thus the reported results in terms of the operating frequency can be directly interpreted in terms of energy in Joules.
- The packet loss rate: The packet loss rate is defined as the number of packets lost divided by the total number of packets delivered to the MAC layer by the upper layer.

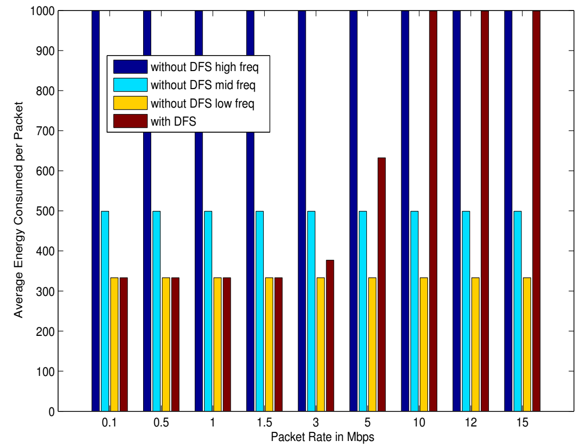

Fig. 3. CBR traffic: Average energy consumed per packet.

Finally, for comparison, we also show the results for wireless devices without DFS. A device without DFS always operates at a single clock frequency. We consider three types of devices without DFS. The first type are the ones that use the frequency corresponding to the lowest frequency in our set of five frequencies used for simulations with DFS. Similarly, the second and third types use the middle and highest frequency from the set of five frequencies, respectively.

A. Results for CBR Traffic

This section presents the results for the case where the application sends CBR traffic. In our simulations, we consider CBR traffic rates in the range 100Kbps to 15Mbps. Figures 3 and 4 show the average energy consumption and packet loss rate for various arrival rates. While the energy consumption per packet of schemes without DFS do not vary with the offered load, we observe that the proposed scheme achieves load proportional energy use. At loads greater than 10Mbps, the network saturates and the proposed DFS scheme always selects the highest frequency. As a result, the energy consumption stays constant for arrival rates greater than 10Mbps. Also, we note that the non-DFS case with low operating frequency has the lowest energy consumption at high loads. However, this comes at the cost of higher loss rates as can be seen from Figure 4. In contrast, the proposed scheme provides a balance between the loss rates and the energy consumed. When the traffic arrival rate is high, the proposed scheme operates at the highest frequency in order to reduce the queue length and packet loss.

B. Results for VBR Traffic

The simulations for VBR traffic were conducted by setting the traffic arrival rate as a random variable following an uniform distribution between a given range, as specified by the maximum and minimum rates. The simulations considered six non-overlapping pairs of maximum and minimum rates, as shown in Table II.

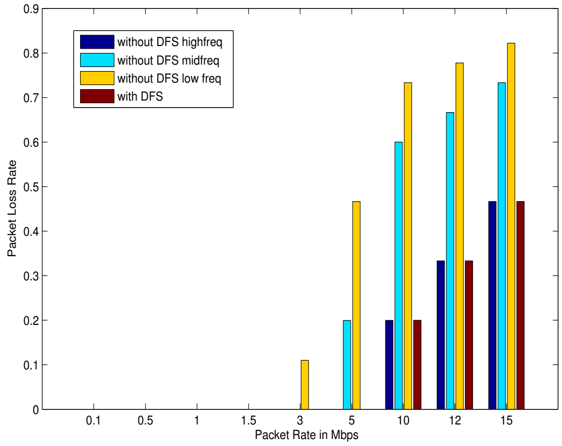

Fig. 4. CBR traffic: Packet loss rate.

TABLE II
VBR PACKET RATES

Index	MinRate	MaxRate
0	0.1mb	0.5mb
1	0.5mb	1.0mb
2	1.0mb	1.5mb
3	1.5mb	2.0mb
4	2.0mb	2.5mb
5	2.5mb	3.0mb

Figures 5 and 6 show the average energy consumption per packet and packet loss rate for the five classes of VBR traffic described in Table II. Again we observe that the proposed scheme provides load proportional energy use. As in the CBR case, when the traffic arrival rate is low, the energy consumption of the proposed scheme equals that of the non-DFS case with low frequency. As the traffic arrival rate increases, the proposed DFS scheme chooses higher operating frequencies and at high loads, the highest operating frequency is chosen.

From Figure 6 we observe that the losses encountered with the proposed scheme are lower than the non-DFS cases with low and medium operating frequencies. Also, the packet loss rate with the DFS scheme is no higher than that of the non-DFS scheme running at the highest frequency.

V. CONCLUSIONS

This paper presented a scheme based on dynamic frequency scaling for achieving load proportional energy consumption in wireless networks. The proposed mechanism is based on observing the transmission queue length and then selecting an operating clock frequency that minimizes the rate of energy consumption while maintaining stability of the queues. Our simulation results show that the proposed scheme achieves a balance between the energy consumption rates and the packet loss rates, while ensuring load proportional energy use.

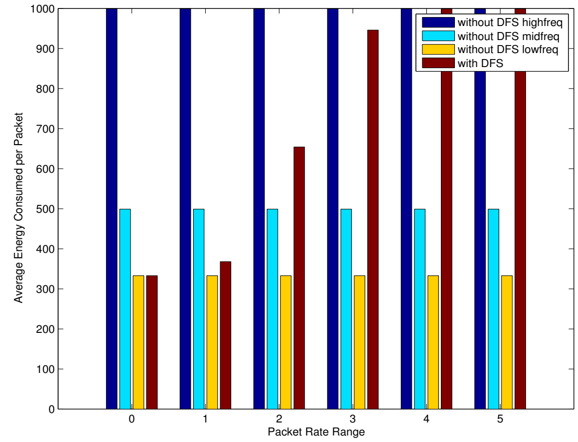

Fig. 5. VBR traffic: Average energy consumed per packet.

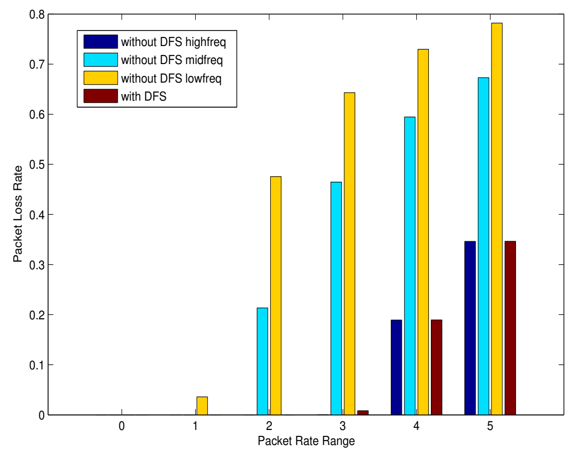

Fig. 6. VBR traffic: Packet loss rate.

REFERENCES

[1] B. Sikdar, "A study of the environmental impact of wired and wireless local area network access," *IEEE Transactions on Consumer Electronics,* vol. 59, no. 1, February 2013.

[2] C. Han et al., "Green radio: radio techniques to enable energy-efficient wireless networks," *IEEE Communications Magazine,* vol. 49, no. 6, pp. 46-54, June 2011.

[3] Y. Chen, S. Zhang, S. Xu and G. Li, "Fundamental trade-offs on green wireless networks," *IEEE Communications Magazine,* vol. 49, no. 6, pp. 30-37, June 2011.

[4] F. Yu, X. Zhang and V. Leung, *Green Communications and Networking,* CRC Press, 2012.

[5] O. Arnold, F. Richter, G. Fettweis and O. Blume, "Power consumption modeling of different base station types in heterogeneous cellular networks," *Proceedings of Future Network and Mobile Summit,* Florence, Italy, June 2010.

[6] M. E. Salehi, M. Samadi, A. Afzali-Kusha, M. Pedram and S. M. Fakhrale, "Dynamic voltage and frequency scheduling for embedded processors considering power/performance tradeoffs," *IEEE Transactions on VLSI Systems,* vol. 19, no.10 pp. 1931-1935, October 2011.

[7] K. Nowka et al., "A 32-bit PowerPC system-on-a-chip with support for dynamic voltage scaling and dynamic frequency scaling," *IEEE Journal of Solid-State Circuits,* vol. 37, no. 11, pp. 1441-1447, November 2002.

[8] H-S. Jung, A. Hwang and M. Pedram, "Predictive-flow-queue based energy optimization for gigabit Ethernet controllers," *IEEE Transactions on VLSI Systems,* vol. 17, no. 8, pp. 1113-1126, August 2009.

[9] K. Choi, R. Soma and M. Pedram, "Dynamic voltage and frequency scaling based on workload decomposition," *Proceedings of Symposium on Low Power Electronics and Design,* pp. 174-179, Newport Beach, CA, August 2004.

[10] K. Choi, W-B. Lee, R. Soma and M. Pedram, "Dynamic Voltage and Frequency Scaling Under a Precise Energy Model Considering Variable and Fixed Components of the System Power Dissipation," *Proceedings of the International Conference on Computer Aided Design,* pp. 29-34, November 2004.

[11] L. Shang, L-S. Peh and N. Jha, "Dynamic Voltage Scaling with Links for Power Optimization of Interconnection Networks," *Proceedings of Symposium on High-Performance Computer Architecture,* pp. 91-102, Anaheim, CA, February 2003.

[12] M. Gupta, S. Grover and S. Singh, "A feasibility study for power management in LAN switches," *Proceedings of ICNP,* pp. 361-371, Berlin, Germany, October 2004.

[13] A. Mishra, R. Das, S. Eachempati, K. Ravishankar, N. Vijaykrishnan and C. Das, "A case for dynamic frequency tuning in on-chip networks," *Proceedings of IEEE/ACM International Symposium on Microarchitecture,* pp 292-303, New York, NY, December 2009.

[14] R. Bolla, R. Bruschi and C. Lombardo, "Dynamic Voltage and Frequency Scaling in Parallel Network Processors," *Proceedings of IEEE HPSR,* pp. 242-249, Belgrade, Serbia, June 2012.

[15] M. Lim,and V. Freeh, "Determining the Minimum Energy Consumption using Dynamic Voltage and Frequency Scaling," *Proceedings of IEEE IPDPS,* pp. 1-8, Long Beach, CA, March 2007.