

# PEARS: A Power-Aware End-to-End Mobility Management Scheme

Surendra Sivagurunathan and Mohammed Atiquzzaman

Telecommunications and Networks Research Lab

School of Computer Science, University of Oklahoma, Norman, OK 73019-6151, USA.

{surain, atiq}@ou.edu

**Abstract**—SIGMA is an IP diversity-based mobility management scheme that aims to reduce handoff latency and packet loss of Mobile IP. In this paper, we propose a power-aware version of SIGMA (called PEARs) which reduces power consumption by intelligently switching off network devices, and improves handoff performance by reducing the number of unnecessary handoffs. The effectiveness of PEARs and SIGMA have been evaluated on a Linux-based experimental testbed. PEARs has been found to require considerably less power than SIGMA. The schemes used in PEARs can also be used to reduce power consumption of multihomed mobile nodes and routers.

## I. INTRODUCTION

The last two decades have seen considerable growth in wireless networks which allow users to access the Internet access without being tied down to one location. The current Internet infrastructure, however, was not initially designed for mobility. Mobile IP (MIP) [1] from the Internet Engineering Task Force addresses the issue of mobility at the network layer, and extends the existing Internet protocol to support host mobility, including handoff, by introducing two network infrastructure entities: Home Agent (HA) and Foreign Agent (FA). A Correspondent Node (CN) communicates with the mobile node (MN) via its HA in the home network, even though the MN may have moved out of its home network. For seamless communication between CN and MN, the MN should be able to handoff quickly between networks.

Base Mobile IP suffers from handoff latency and packet loss – two of the most important indicators of handoff performance in mobile networks. Large handoff latency degrades performance of real-time applications. For example, a large handoff latency will introduce interruption in a video conference session due to breaks in both audio and video data transmission during handoffs. Mobile IP also requires changes in the Internet infrastructure due to the addition of the HA and FA. To address the limitations of Mobile IP, an end-to-end handoff scheme, called Seamless IP diversity based Generalized Mobility Architecture (SIGMA), which utilizes multi-homing and IP diversity to achieve seamless handoff between networks was proposed [2]. Its performance has been thoroughly analyzed [2], [3], [4] and its suitability for use in space networks has also been studied [5], [6].

Power limitations in space networks and mobile handheld devices require mobility management schemes to be power efficient. Inability to consider power consumption of mobile devices is a limitation of most handoff protocols. Most work on power-awareness have been limited to switching off the display and optimizing ad hoc network routing protocols; power awareness of handoff schemes have not been considered.

The research reported in this paper was funded by NASA Grant NNX06AE44G.

The performance of SIGMA has been thoroughly studied in the literature. However, its effect on power consumption of the mobile node has not been considered in previous work. The *objective* of this paper is to develop a power-aware version of SIGMA for use in space and terrestrial networks. We propose a new algorithm called Power-Aware SIGMA (PEARs) that reduces power consumption by optimally switching off network interfaces with insignificantly contribution to the performance of SIGMA. The optimality is achieved by allowing a fast handoff while maintaining stability of the system.

PEARs has been evaluated using a laboratory prototype based on Linux machines and off-the-shelf wireless networking components. Results show a significantly lower power consumption of PEARs when compared to SIGMA. The novelty of this work lies in determining the optimal handoff point between networks to reduce the power consumption of the mobile node by reducing the duration of multiple powered interfaces. The *authors are not aware* of any previous effort to reduce power consumption of end to end mobility management schemes employing multihomed mobile nodes. The *contributions* of this paper are as follows: (i) *Developed* PEARs to improve the performance of SIGMA; and (ii) *Implemented* and *evaluated* the performance of PEARs on a Linux-based experimental prototype.

The rest of the paper is organized as follows: Sec. II gives a brief introduction to SIGMA and its limitations. PEARs is described in Sec. III, and details of the experimental setup can be found in Sec. IV. Results obtained from the experimental setup are described in Sec. V, followed by concluding remarks in Sec. VI.

## II. INTRODUCTION TO SIGMA

To aid the reader in better understanding SIGMA, we briefly describe the various steps of SIGMA handoff in this section (detailed description can be found in [2]). SIGMA is based on exploiting IP diversity, i.e. multiple IP addresses associated with multiple interfaces of a mobile node, to achieve a seamless handoff. We use Stream Control Transmission Protocol [7], a new emerging IETF standard transport layer protocol with multihoming capabilities, to illustrate SIGMA.

### A. Handover Procedure of SIGMA

A typical SIGMA handover (using SCTP as the transport protocol) is shown in Fig. 1, where a multihomed MN (host) connects to the Internet through two wireless access networks. Correspondent node (CN) sends data to MN, which corresponds to services like file downloading or web browsing by mobile users. The handover process of SIGMA consists of the following steps [3]:

**STEP 1: Obtain new IP address:** The handover procedure starts when the MN moves into the overlapping radio coverage

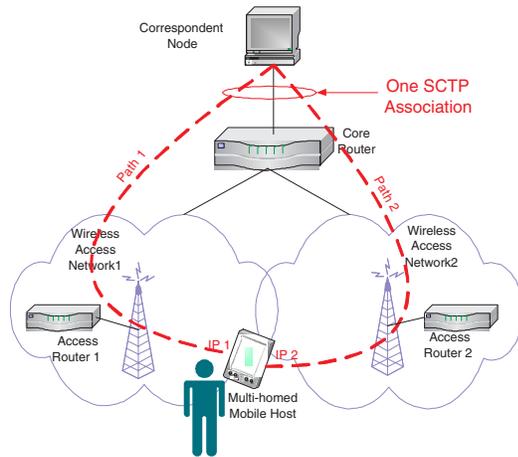


Fig. 1. An SCTP association featuring multi-homing.

area of two adjacent access networks. Once the MN receives a router advertisement from the new access network (Access Router 2 in Fig. 1), MN obtains a new IP address and binds it to the second network interface.

**STEP 2: Add IP addresses to the association:** After the MN obtains a new IP address in STEP 1, it notifies CN about the availability of the new IP address through SCTP's Dynamic Address Reconfiguration option.

**STEP 3: Redirect data packets to new IP address:** When the MN moves further into the coverage area of the new access network, and the signal strength of the new access network is *greater* than that of the current network, a handoff is initiated. MN informs the CN of the handoff, and CN updates its *routing table* and redirects data to the new IP address to increase the possibility of data being successfully delivered to MN. Although not shown in Fig. 1, a location manager is also updated with the new IP address.

**STEP 4: Delete or deactivate obsolete IP address:** When the MN moves out of the coverage of current access network, new or retransmitted data should not be sent to the old IP address of MN. MN uses dynamic address reconfiguration to notify CN that the old IP address is unavailable for data transmission.

### B. SIGMA Handoff Decision and Limitations

Making handoff decisions based on comparison of signal strength of two access points (step 3 in Sec. II-A) suffers from the following disadvantages:

- **Unnecessary handoffs:** Unnecessary handoffs (Ping Pong) result from signal strength fluctuations in overlapping area.
- **Unnecessary modification of Routing Table in MN:** Each handoff requires time to update the routing table (step 3 in Sec. II-A) of the MN<sup>1</sup> and increases the handoff latency.  $N$  unnecessary handoffs due to Ping Pong results in  $N$  modifications of the routing table, resulting in large handoff latency.
- **Unnecessary updating of the Location Manager:** Each handoff requires updating the Location Manager with the new IP address.  $N$  unnecessary handoffs require  $N$  irrelevant update of the Location Manager.
- **Unnecessary power consumption:** The use of multiple interface cards, but use of only one card except in the overlapping region, results in unnecessary power consumption by multiple interface cards.

<sup>1</sup>A Pentium 4-based Linux box requires 54 milliseconds per update.

The MN is in one of two states during handoff [4]: (i) STABLE state where the MN receives data and sends SACK through the same IP address (interface); (ii) UNSTABLE state where the MN receives data through one IP address and sends SACK through a *different* IP address. This is called UNSTABLE because, the MN should use only one IP address at any instant of time. The SIGMA handoff scheme causes MN to be in UNSTABLE state during handoff. The instability of SIGMA results from two factors:

- 1) *Fluctuation of signal strength*, which increases the number of handoffs due to the Ping Pong effect.
- 2) *Route cache effect*, where the operating system kernel of the MN first searches the "route cache" for determining the outgoing interface of a packet, followed by a search in the main "routing table" (also called Forwarding Information Base (FIB)). If the kernel finds a matching entry during route cache lookup, it immediately forwards the packet and stops traversing the routing tables. Because the routing cache and "routing table" and maintained separately by the kernel, updating of the routing table by SIGMA during handoff may not have an immediate effect on the path of a given packet.

This UNSTABLE state for SIGMA increases handoff latency, which is measured by the time taken to "completely" switch from between networks. The MN is completely under the new network, i.e., uses the new IP address for *both* data and SACKs, only after end of the UNSTABLE state. Our *aim* is to reduce the duration of the UNSTABLE state, and thus reduce the handoff latency. It is important to reduce the UNSTABLE state because packets will be lost if an access point becomes unavailable while the MN is using both the interfaces. We propose to remove this UNSTABLE state by using an efficient handoff scheme discussed in Sec. III.

### III. PEARS: POWER-AWARE SIGMA

To overcome the limitation of SIGMA (Sec. II), we propose a new power-aware mobility management scheme called PEARS. The power consumption of MN is reduced by optimally switching off the network interfaces that do not significantly contribute to the handover performance. The optimality is achieved by reducing the handoff latency while still maintaining stability of the system by using a finely tuned Signal to Noise ratio as the handoff trigger and flushing the operating systems route cache following handoffs.

#### A. Handoff Scheme

Commonly used metrics for Link Layer handoff trigger are [8]: (i) Signal strength, which is suitable in the absence of interference or noise; (ii) Signal-to-noise Ratio, which is suitable in the presence of noise; (iii) Signal-to-interference Ratio, which can be used in the presence of both interference and noise; (iv) Bit error rate, which can be used in erroneous channels with interference, although error control is needed to detect and correct errors; (v) Frame error rate, which can also be used in erroneous channels but not recommended due to its computing complexity. Our experimental laboratory environment inside buildings consist of both noise and interference; we thus use Signal-to-Noise Ratio (SNR) for handoff decision, and implement hysteresis to avoid any Ping-Pong effect.

Hysteresis is the minimum value by which the two signals that are compared during handoff differ from one another to trigger a handoff. MN is handed off to the stronger signal if the difference between the two signals is greater than the hysteresis. Depending on the fluctuation of the signal strengths, the hysteresis value may have to be adjusted. For large fluctuations, the hysteresis has to be higher. The

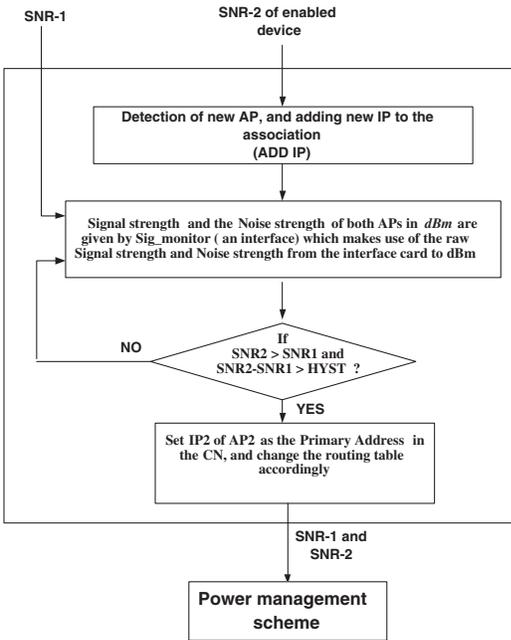


Fig. 2. Handoff scheme of PEARS.

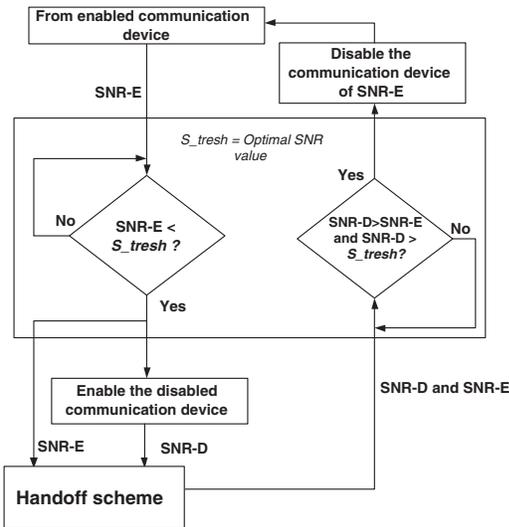


Fig. 3. Power management scheme of PEARS.

value of the hysteresis is thus dynamic and depends on the fluctuations of the signal strengths. This raises the following question: What value of hysteresis should be used in a certain environment? Our handoff scheme using hysteresis is shown in Fig. 2.

### B. Power Management Scheme

SIGMA uses two interface cards, but only one is used for communication at any time. The other card searches for alternate networks, in case the signal from the current access point degrades. Our proposed power management scheme is based on minimizing the time for which both the interface cards are in powered up state.

Fig. 3 shows the flow chart of the power management scheme having three variables:  $SNR - D$ ,  $SNR - E$  and  $S_{thresh}$ .  $SNR - D$  is the Signal-to-Noise Ratio of the signal received from the recently enabled interface card, and  $SNR - E$  is the Signal-to-Noise Ratio from the previously

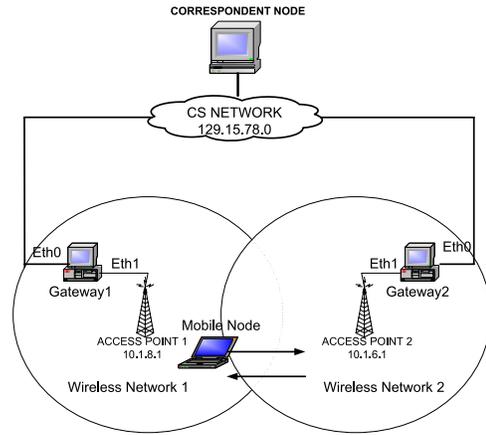


Fig. 4. Topology of the SIGMA experimental testbed.

enabled interface card. At any time,  $SNR - E$  is the Signal-to-Noise Ratio of the primary interface card. When an interface card is disabled, the Signal-to-Noise Ratio from the enabled card becomes  $SNR - E$ .  $S_{thresh}$  is a threshold value that determines when to enable or disable an interface card.

As shown in Fig. 3, the power management scheme initially uses  $SNR - E$  but keeps on checking periodically. A low  $SNR - E$  indicates possible handoff; the secondary interface is enabled when it drops below  $S_{thresh}$ . The two SNRs ( $SNR - E$  and  $SNR - D$ ) are fed to the handoff scheme, where a handoff decision is made based on these values. The two values are then fed back to the power management scheme where a decision to disable one of the interfaces is made. Disabling an interface requires the following two conditions to be satisfied:

- The SNR of the primary interface should be greater than that of the secondary interface.
- The SNR of the primary interface should be greater than  $S_{thresh}$ .

The first condition ensures that the MN is still in the new network and is making use of the recently enabled interface. The second condition ensures that even if the SNR of the new enabled interface is greater than the already enabled interface, it should be greater than the  $S_{thresh}$  to make a meaningful communication. If the second condition is not true, then the MN is away from both the available wireless networks. None of the interfaces is disabled if the second condition is false.

## IV. EXPERIMENTAL TESTBED

Fig. 4 shows the testbed (similar to those previously used in literature [9], [10]) that has been used to evaluate the performance of PEARS. It consists of four entities: Correspondent Node (CN), Gateway 1, Gateway 2, and Mobile Node (MN). The Gateways serve as Access Points to two separate private 802.11b wireless networks. The CN and Gateways are directly connected to the public Computer Science network of the University of Oklahoma.

Traditional file transfer programs (such as, ftp) are based on TCP sockets, and are not available for the SCTP protocol used in PEARS. We, therefore, wrote our own SCTP socket-based file transfer program to transfer data from the CN to the MN. To obtain access to the SCTP socket, we used Linux 2.6.2 kernel with Linux Kernel SCTP (lksctp) [11] version 2.6.2-0.9.0 on both CN and MN. Ethereal was used on both CN and MN to capture packets during handoff. Results from the analysis of the packets are given in Sec. V.

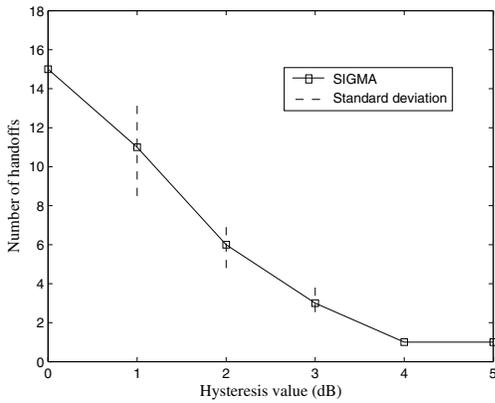


Fig. 5. Relation between hysteresis value and number of handoffs.

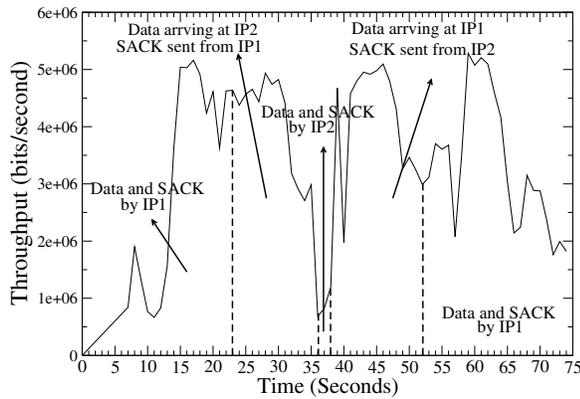


Fig. 6. Throughput without using hysteresis and route caching flush.

## V. RESULTS

In this section, we use the experimental testbed (see Sec. IV) to measure the performance of the handoff and power management schemes of PEARS.

### A. Effect of Handoff Scheme

The effectiveness of the handoff scheme of PEARS has been studied by looking into the effect of hysteresis and route cache flushing on the end to end throughput between CN and MN.

1) *Effect of Hysteresis*: To analyze the effect of hysteresis, the throughput was measured with and without hysteresis. Fig. 5 shows the reduction in the number of handoffs with an increase in the hysteresis value. Figs. 6 shows the throughput as a function of time, where the two UNSTABLE region are each 14 secs. (between 22 to 36 secs. and 38 to 52 secs.) in duration. The regions were determined by analyzing the packets captured by the Ethernet packet sniffer. The large duration of the UNSTABLE region in the absence of hysteresis arises from excessive number of handoffs when the MN is in the overlapping region.

Fig. 7 shows the throughput with *hysteresis* where we can see three STABLE regions spanning the following durations: 0 to 19 secs. (both data and SACKS go through IP1 (10.1.8.100)), 19 to 39 seconds (IP2 (10.1.6.100) is used for data and SACK), and 39 seconds onwards (data and SACKS use IP1). MN always uses the same interface for both data and SACKS, and thus is in the STABLE state when hysteresis is used.

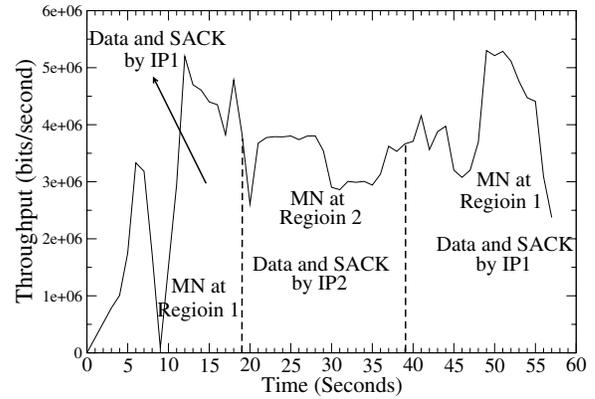


Fig. 7. Throughput with hysteresis and route cache flush.

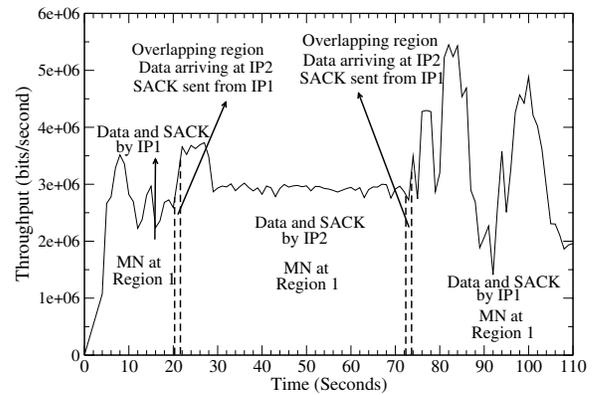


Fig. 8. Throughput with hysteresis but without route cache flush.

2) *Effect of Route Cache Flushing*: To show the effect of route caching, we present throughput with and without flushing the route cache. The throughput *without route cache flush* is shown in Fig. 8 which is divided into five regions of alternating STABLE and UNSTABLE states: STABLE (0 to 20.57 secs.), UNSTABLE (20.57 to 22.18 secs.), STABLE (22.18 to 74.97 secs.), UNSTABLE (74.97 to 75.50 secs.), and STABLE (75.5 secs. onwards). There are two UNSTABLE regions, although hysteresis has been used to ensure a single handoffs in the overlapping region. As discussed in Sec. II-B, in the UNSTABLE state which partially results from the caching effect of the routing table, the MN uses different interfaces for sending data and receiving SACKs. Earlier, we showed in Fig. 7 that the MN is always in the STABLE state when hysteresis and *route cache flush* are used.

With a view to determining the critical components contributing to the handoff latency, we show the measured handoff latency of PEARS (i.e. with route cache flush and hysteresis) in Fig. 9. Route cache flush reduces the time required for the routing table change to only a few milliseconds; separate measurements showed this delay to be around two seconds when route cache flush was not used. The routing table change contributes only a few milliseconds to the handoff latency between MN and CN when compared to the RTT of the Set\_Primary, which varies between 100 to 200 milliseconds. A major part of the handoff latency is due to the RTT, even when the MN and CN are in the same subnet of the Computer Science department at the University of Oklahoma.

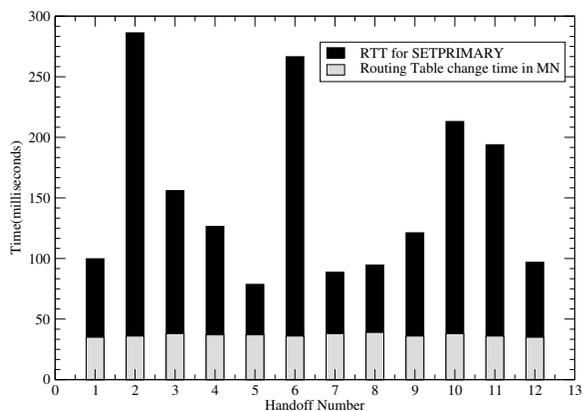


Fig. 9. Various components of the handoff latency.

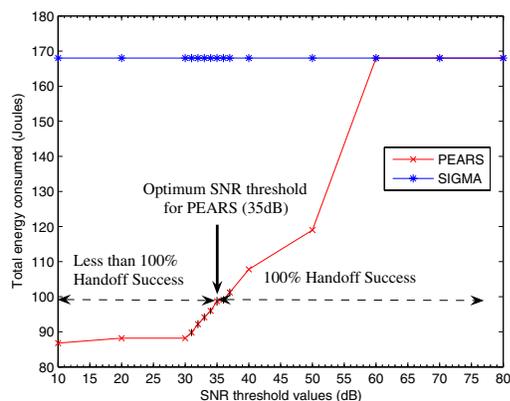


Fig. 10. Energy consumed by SIGMA and PEARS for different SNR threshold values (Critical region 30 dB to 37 dB repeated 10 times to find success rate and standard deviation).

### B. Effect of Power Management Scheme

In this section, we show the power efficiency of PEARS (Sec. III-B) by comparing with SIGMA. To increase the effectiveness of PEARS by using an optimal threshold value,  $S_{thresh}$  (see Sec. III-B), power consumed by PEARS was measured for different values of  $S_{thresh}$ . The maximum SNR recognizable by the Avaya PCMCIA card is 70, and the minimum SNR required for communication is 10. So we

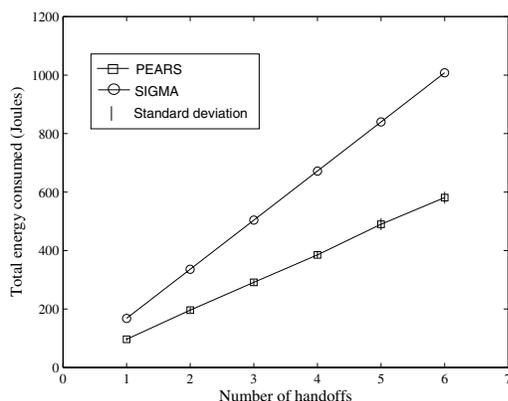


Fig. 11. Energy consumed by SIGMA and PEARS for different number of handoffs (showing standard deviation for PEARS).

collected results for  $S_{thresh}$  in the range 10 to 80.

Fig. 10 shows energy consumed by SIGMA and PEARS for different values of  $S_{thresh}$ . A lower SNR threshold value is found to reduce the energy consumption. However,  $SNR < 35$  results in unsuccessful handoff due to insufficient time between adding the new IP address to the existing association and setting the new IP address as primary (Sec. II-A).

An SNR threshold of 35 has the minimum energy consumption with successful handoff; we therefore, set the threshold to 35 in PEARS to obtain the energy consumption of SIGMA and PEARS for different number of handoffs in Fig. 11. The power consumption of both SIGMA and PEARS increases with increased number of handoffs, but still PEARS consumes less energy than SIGMA for a given number of handoffs.

We used the optimum SNR threshold value to obtain the energy consumption of SIGMA and PEARS for different number of handoffs in Fig. 11. Although the power consumption of both SIGMA and PEARS increases as the number of handoff increases, PEARS still consumes less energy than SIGMA for any number of handoffs.

## VI. CONCLUSIONS

We have proposed PEARS, a power-aware IP-diversity based end-to-end mobility management scheme for data networks. PEARS reduces power consumption by optimally switching off network interfaces having insignificant contribution to handover performance. The optimality is achieved by fast handoff while maintaining stability of the system. We have studied the effect of hysteresis on triggering handoff and the effect of route cache flush on the performance of handoff. Results from Linux-based experimental testbed show that PEARS reduces power consumption when compared to previous schemes.

## REFERENCES

- [1] C.E. Perkins (editor), "IP Mobility Support." IETF RFC 3344, August 2002.
- [2] S. Fu, L. Ma, M. Atiquzzaman, and Y. Lee, "Architecture and performance of SIGMA: A seamless mobility architecture for data networks," *40th IEEE International Conference on Communications (ICC)*, Seoul, Korea, May 2005.
- [3] S. Fu, M. Atiquzzaman, L. Ma, and Y. Lee, "Signaling cost and performance of SIGMA: A seamless handover scheme for data networks," *Journal of Wireless Communications and Mobile Computing*, vol. 5, no. 7, pp. 825–845, Nov 2005.
- [4] S. Sivagurunathan, M. Atiquzzaman, and W. Ivancic, "Improving stability of SIGMA handoff," *IEEE 62nd Semiannual Vehicular Technology Conference*, Dallas, TX, pp. 836–840, Sep 22-25, 2005.
- [5] P. Chowdhury, M. Atiquzzaman, and W. Ivancic, "SIGMA for seamless handover in space," *Sixth Annual NASA Earth Science Technology Conference*, College Park, MD, June 27-29, 2006.
- [6] M. Atiquzzaman and W. Ivancic, "SIGMA: An end-to-end mobility management scheme for space networks," *Fifth Space Internetworking Workshop*, Baltimore, MD, Sep 12-14, 2006.
- [7] S. Fu and M. Atiquzzaman, "SCTP: State of the art in research, products, and technical challenges," *IEEE Communications Magazine*, vol. 42, no. 4, pp. 64–76, April 2004.
- [8] A. Festag, "Optimization of handover performance by link layer triggers in IP-based networks: Parameters, protocol extensions and APIs for implementation," tech. rep., Telecommunication Networks Group, Technische University, Berlin, July 2002.
- [9] S. Seol, M. Kim, C. Yu, and J.H. Lee., "Experiments and analysis of voice over Mobile IP," *13th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, Portugal, pp. 977 – 981, 15 - 18 September 2002.
- [10] W. Wu, N. Banerjee, K. Basu, and S.K. Das, "Network assisted IP Mobility support in wireless LANs," *Second IEEE International Symposium on Network Computing and Applications*, Massachusetts, USA, pp. 257 – 264, 16- 18 April 2003.
- [11] "http://lksctp.sourceforge.net."