Autonomic Network Configuration for Networkable Digital Appliances

Yasuharu Katsuno and Toru Aihara, Member, IEEE

Abstract —*This paper proposes an autonomic network configuration (ANC) technology, which uses autonomic technologies for the network configuration in order for electronics consumer appliances to improve usability and security. ANC provides networkable digital appliances (NDAs) with a network environment detection service and enables them to automatically configure themselves as soon as they are connected to networks. The evaluation shows that ANC's network environment detection is so quick, safe, and accurate that it significantly contributes to the NDA's security and privacy*[1].

Index Terms — location awareness, modular computer design concepts, network appliance, network environment detection service

I. INTRODUCTION

Recently, digital consumer devices are quite similar to personal computers (PC) in terms of system architectures and electronic components. They now have fast CPUs, mass storage, and, most importantly, network connectivity [1], as well. In spite of their rich functionalities and high capabilities, they still suffer from inferior usability, especially for setting up their network connectivity and security, mainly due to the cost and size constraints. They usually cannot accommodate complex settings for securely configuring a LAN or WLAN, regardless of their mobility.

Autonomic technologies [2][3] are receiving increasing attention for addressing usability issues, because these devices should configure themselves with no or minimal user interaction by automatically recognizing their situations and environments. They are especially useful for scenes of configurations, where complex settings are requested using limited user interfaces.

We use autonomic technology for network configurations for consumer electronics appliances, and propose an autonomic network configuration (ANC) technology that performs network environment detection and configures the appliances quickly, accurately, and safely. We implemented the ANC technology on a networkable digital appliance (NDA). ANC supplies the NDA with location awareness, and performs complex network configuration on behalf of its user, who is usually not familiar with the network.

In the next section, we introduce the NDA, which is the platform we implemented ANC on. Section III describes the ANC technology in detail. Section IV evaluates the network environment detection that ANC provides for the NDA. Finally we conclude the paper in Section V.

II. NETWORKABLE DIGITAL APPLIANCE

The NDA was developed to accommodate users who wish to use their computing environments in several different locations, but who do not want to carry their computers. The typical target users are office workers and students, who use their computers throughout their daily activities, such as taking notes, sending mails, making presentations, and preparing reports. Figure 1 and Table I show the NDA photograph and its hardware specifications, respectively. The NDA is based on the IBM PC/AT architecture, and is capable of hosting a variety of operating systems (OSs) such as the Microsoft Windows family and Linux.

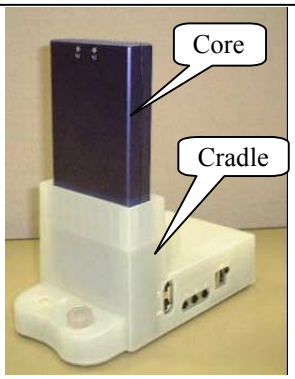

Figure 1: Networked Digital Appliance

TABLE I
Specification of Networked Digital Appliance

	Core	Cradle
CPU	TM5800 933 MHz	
Memory	SDR-256 Mbytes	
HDD	2.5" HDD 20GB	
VGA	Analog Out	X1 HD-DSUB 15-Pin
USB	USB 1.1	X2 Port
Network	10/100baseT	X1 RJ45
KBD	PS2	1 pair
FAN		50x50x10.2 mm
PCCard Slot		X1 PCMCIA Type-2
MiniPCI Slot		X1 (for Ethernet Card)
Power		X1 (for AC Adapter)
Size	83x185x23.4 mm	106x97x106 mm
Weight	350 g	800 g

[1] Yasuharu Katsuno and Toru Aihara are with IBM Research, Tokyo Research Laboratory, 1623-14, Shimotsuruma, Yamato-shi, 242-8502, Japan (e-mail: katsuno@jp.ibm.com, aihara@jp.ibm.com) .

Contributed Paper
Manuscript received November 7, 2004

The NDA was developed with the modular computer design concept [4]. In this concept, the device is divided into two parts: a core and a cradle. The core has the CPU, memory, and disk drive, while the cradle has external connectors to the power, display, keyboard, mouse, and network. The concept enables users to carry their own environments (applications and data) and to access their environment at any places where a cradle is available. A typical user will have several cradles set up in different locations, each with a display, keyboard, mouse, and network connection. Users can move around only with just their personalized cores and still continue to work with their familiar computing environments as if they were carrying their entire computer systems. At the same time, they can share the cradles (as well as the connected display, keyboard, and so on) with each other. The cradle will serve for any core to recover its computing environment.

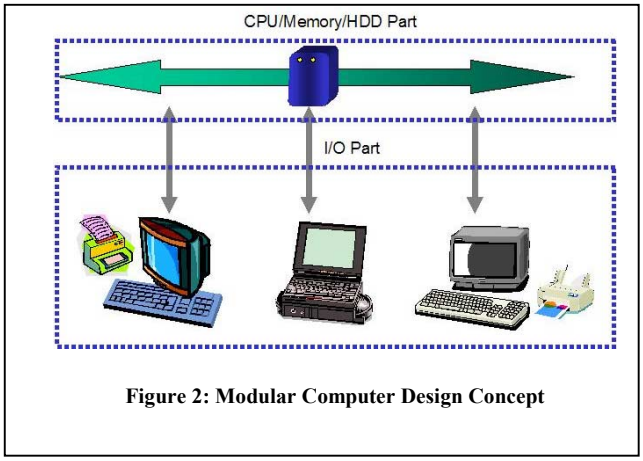

Figure 2: Modular Computer Design Concept

There are two alternative solutions to the NDAs: portable storage and a laptop computer. In the case of portable storage such as a USB key, users can carry their own data, but have to tolerate different computing environments or even have to install their favorite applications in various locations, which could violate the license agreements of the applications. In the case of the laptop computer, users can carry all their data and environment with them, but suffer inconvenience due to the size and weight of the computer, especially due to its LCD, keyboard, and battery.

The NDA is an ideal solution for an organization where people would use computers installed in many locations. For example, in universities, students use computers in classrooms, in workshops, in laboratories, and even at their homes. With NDAs, students only need to carry small, compact cores, but they can still access their own computing environments everywhere.

III. ANC Technology

A. Connectivity problems with NDAs

Although the NDAs can maintain good portability of computing environments, they are still not free from setting up network environments. When users move from one location to another, they usually have to change network settings appropriate to that location, such as the IP address, the proxy server, the printer, the file sharing conditions, and the security configuration. Currently, manual operations are required, because the device cannot detect the changes of the network environments. However automatic settings are definitely desired for consumer devices such as NDAs so that the devices can automatically adapt to the specific settings and so that users don't need to be aware of which network environment the devices are connected to.

Among the network settings, there are two types of important settings. The IP parameter settings such as the IP addresses of the device, the router, and the domain name service (DNS) [5] servers, are indispensable for the device to be connected to the network. Security settings, such as the file sharing, printer sharing, and firewalls, on the other hand, are also important for the device to balance between performance and security or privacy. It is also important to note that the security settings should complete before the network settings to prevent any security exposure.

As far as the network settings concerned, dynamic host configuration protocol (DHCP) [6] is a well designed system for devices to configure IP settings dynamically. Users are no longer required to manually configure their IP addresses, net masks, local gateway IP addresses, and DNS IP addresses. Instead, their OS will request such information from a DHCP server (a DHCP request packet), and automatically configure itself with the parameters returned in the response (a DHCP ACK packet).

However there are still two timing-related security problems when NDAs utilize these DHCP parameters. First, it may be too late to configure security settings, once an IP address is assigned and configured. Second, an IP address is not always unique among the networks, especially when it is a private address.

B. Architecture

ANC is designed to autonomically configure network settings. ANC is equipped with a packet monitor and offers a service to detect the current network environment efficiently by holistically monitoring the link status and network packets directly.

The monitor is based on a kernel-space approach. There are two approaches for monitoring packets: the kernel-space approach and the user-space approach (Figure 3). The kernel-space approach monitors real-time packets directly, while the user-space approach indirectly monitors the network status through user-space APIs. The kernel-space approach is more precise and provides more up-to-date information than the user-space approach. For example, timeouts, window sizes, retransmissions, and error rates can be monitored by the kernel-space approach, but not by the user-space approach.

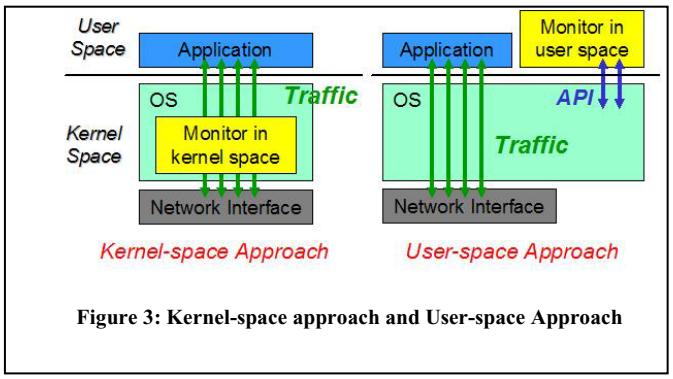

Figure 3: Kernel-space approach and User-space Approach

The packet monitor detects unusual states, such as increases in error rates, drop rates, and collision rates, by constantly monitoring network communications at various levels of the network architecture hierarchy. For example, it looks at link status at the physical layer, Ethernet headers are in the network layer, IP headers in the Internet layer, and TCP headers and UDP headers in the transport layer. The monitor can report on fatal errors (such as drops and bit errors) as well as many sensitive symptoms (such as increasing delay). There are two types of monitor engines: a passive monitor and an active monitor.

The passive monitor is completely silent, and just tracks the packets initiated by applications or the OS. It can monitor the normal network traffic and detects many kinds of errors and unusual behaviors. The active monitor, on the other hand, is a kind of explorer, and creates its own network packets to check the state of network devices. This is useful in order to produce network packets for diagnostics when the passive monitor detects evidence of certain kinds of problems that only produce indirect symptoms, such as delays in network transmissions. It is also necessary in order to probe the alternative paths (such as another network interface or gateway), because they may otherwise not be activated or be used by the OS.

The passive and active monitors are complimentary to each other. For example, the system checks the availability of the local gateway in the following way. The passive monitor constantly just watches for any network packets from or to the local gateways, such as result from normal operations of the user. The active monitor periodically sends an address resolution protocol (ARP) [7] request packet to the local gateway in order to check its availability or responsiveness, which might be necessary when there are no application-driven packets for a long time.

C. Network Environment Detection Algorithm

ANC defines access profile that keeps any information about the network connection at one particular location, such as network, security, Internet, and application settings. ANC identifies the network environment, and tries to match it with an appropriate access profile.

ANC uses the media access control (MAC) addresses of key network devices instead of the network addresses or IP addresses to distinguish between local network environments. The reason is because the MAC addresses do not change and are guaranteed by the IEEE to be unique throughout the world.

By holistically monitoring link status and network packets, ANC tries to identify different key network devices depending on the network configurations. Some network devices are conveniently broadcasting their MAC addresses, but others do not reveal their MAC addresses until they are requested to respond. In addition, some are not always present in all the network environments. Although a local gateway server and a DNS server are always present in the network, ANC is also interested in layer-2 switches, DHCP servers, and local gateways due to their behavior of providing their MAC addresses.

1) Layer-2 Switch

A layer-2 switch is a network device that efficiently connects a number of devices to the same subnet of the LAN as shown in Figure 4. It is commonly used on large or middle size Ethernet networks. It periodically (usually every two seconds as a default) sends a spanning tree protocol (STP) [8] packet using layer-2 multicast to maintain a loop free network. The STP produces frequent control packets, and carries the unique MAC address of the local layer-2 switch that the device is connected to, even though the device itself is usually not interested in STP packets, but just drops them. Without regard to the original purpose of STP, the device can acquire the MAC address of the layer-2 switch within two seconds using the ANC passive monitor, as soon as it connects to a network environment equipped with a layer-2 switch. In addition, the layer-2 switch can be used to precisely locate the position in the same subnet like an access point can.

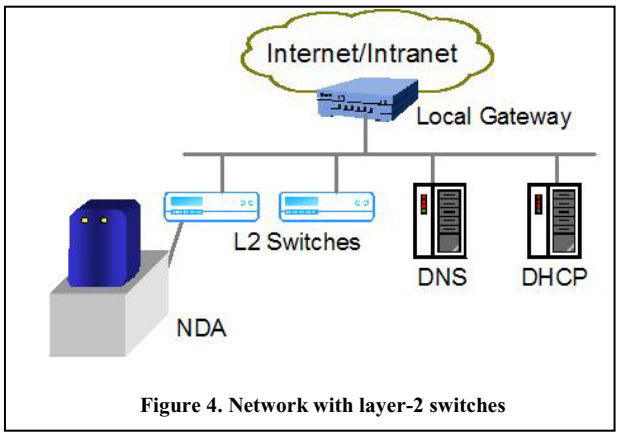

Figure 4. Network with layer-2 switches

2) DHCP Server

As shown in Figure 5, a DHCP server is usually present even in a small network where layer-2 switches are not used. DHCP is quite widely used, and it also transmits a unique MAC address for its local DHCP server. The MAC address that the DHCP packet carries identifies the network that the device is connected to. The DHCP server sends a DHCP offer packet that carries its unique MAC address, in response to a

DHCP discover packet that is sent by a device. Note that the device can send the DHCP discover packets, even when it does not yet have an IP address, because the DHCP discover packet needs no IP address. When the device sends the DHCP discover packet, it receives a DHCP offer packet, and acquires the MAC address of the DHCP server, within a few seconds or as soon as it connects to a network equipped with a DHCP server. However, DHCP offer packets are not sent automatically or periodically, but only in response to a DHCP discover packet from the device. By monitoring a DHCP offer packet in response to the OS's sending a DHCP discover packet, the ANC passive monitor can acquire the MAC address of the DHCP server just before the IP configuration.

Some devices with knowledge of their previous IP address allocation may skip a DHCP discover message, and start with a DHCP request message. Regardless of its response from the DHCP server (either a DHCP ACK or NACK message), the MAC address of the DHCP server is still available before the IP configuration. Devices with the ANC passive monitor can use any message from the DHCP server, while the IP configuration has to wait for the successful reception of a DHCP ACK message.

In some network environments, a DHCP packet carries the MAC address of the local gateway instead of that of the DHCP server, because the DHCP server is not on the same LAN segment, and the local gateway works as a DHCP relay agent. But the MAC address of the local gateway represents the LAN segment in the same way as the MAC address of the local DHCP server does. In either case, the MAC address of the DHCP uniquely identifies the network environment.

We are more interested in the DHCP server than the local gateway. The reason is because their MAC addresses are available even before the device initializes its IP stack. Note that the MAC address of the local gateway can be obtained only after the device gets its IP address, because ARP requires the IP address of the local gateway, which is available only after receiving a DHCP ACK packet that authorizes the device to use the allocated IP address.

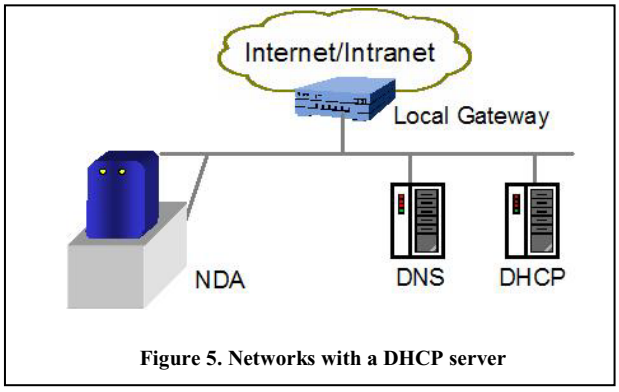

Figure 5. Networks with a DHCP server

3) Local Gateway

In fixed IP environments as shown in Figure 6, a DHCP server is not available, and we have to rely on the local gateway. ANC can use ARP to recognize such environments. The access profiles remember a list of pairs of the IP and MAC addresses of all of the local gateways ever used as fixed IP environments. The ANC active monitor will send an ARP request packet to all the IP addresses in the list until it receives an ARP response that matches the MAC address of the corresponding IP address.

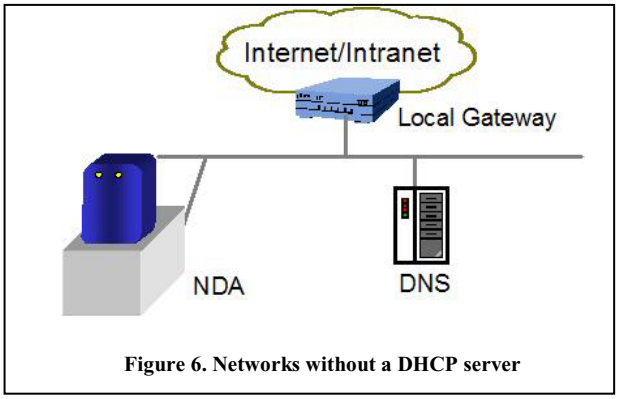

Figure 6. Networks without a DHCP server

With these three steps described above, ANC can distinguish network environments quickly, accurately, and safely. For the sake of quickness and reliability, ANC tries to acquire the MAC addresses of both the layer-2 switch and the DHCP server, giving slight preference to the STP approach due to its periodic advertising. If no such information is available, ANC will use its active monitor to check if there is a local gateway with the known pair of the IP and MAC addresses.

Access profiles will keep the MAC addresses for layer-2 switches, DHCP servers, and local gateways. In some advanced environments, there may be more than one local gateway (to improve reliability and availability) and the access profile should be ready to remember all of them. In addition, since there can sometimes be hardware replacements in the network, access profiles should be ready to replace old MAC addresses with new ones.

D. Example

The ANC technology gives the NDAs location awareness features and supports busy mobile users actively moving from one location to another. The user doesn't have to worry about the settings of the network parameters, and the system can even start location-specific applications. ANC takes care of the settings and starts programs by detecting the network environment the device is connected with.

We implemented a demonstration system for location-based services on the NDAs as shown in Figure 7, which shows the location-awareness without requiring any special hardware, but still offers the similar capabilities as described by PARCTab [9][10], the Olivetti Active Badge system [11][12], and the InfoPad project [13]. In the demo, there are three locations, a classroom, a home, and a laboratory. The mail application launches when the core is docked to the cradle in the classroom, the account book application at home, and the CAD application in the laboratory with the high-resolution display.

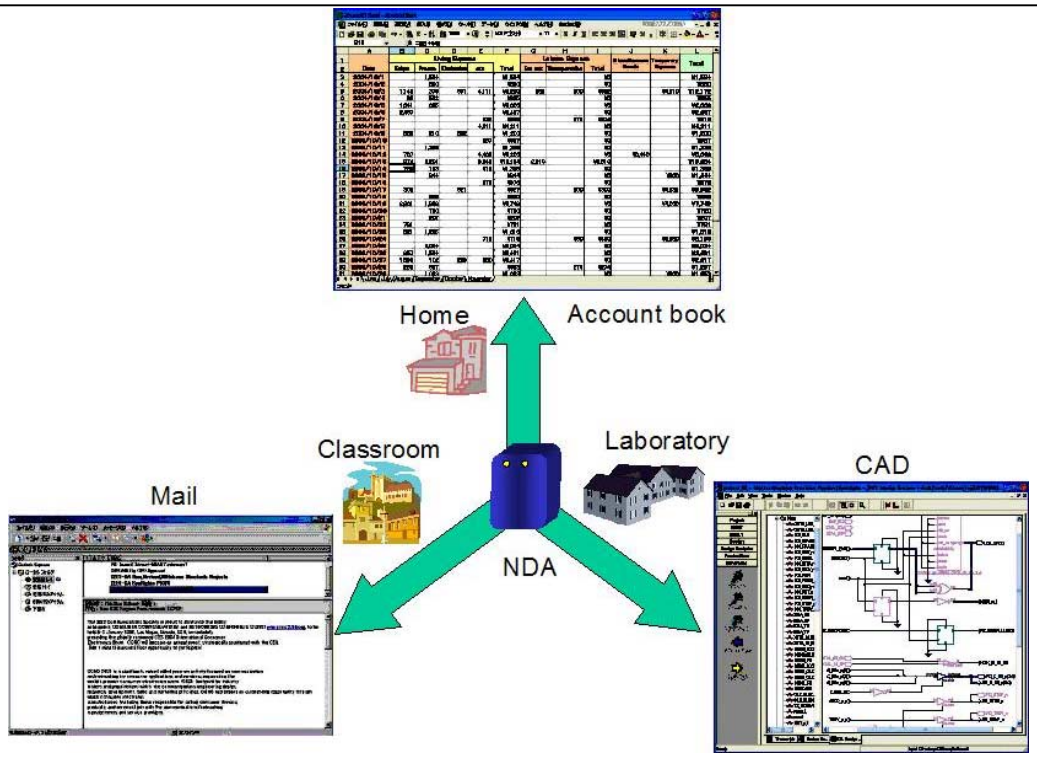

Figure 7: Location-based Service

TABLE II
ANC's Network Detection and Microsoft Windows Network Configuration

(a) Reconnection to the same network after disconnection

Network Environment	With ANC and access profiles		Without ANC	Time difference between ANC detection and OS's IP configuration (sec)
	ANC detection time (sec)	IP configuration time (sec)	IP configuration time (sec)	
Layer-2 Switch	2.4	7.2	7.3	4.9
DHCP Server	5.6	5.6	5.6	0.0
Local Gateway	0.0	0.0	0.0	0.0

(b) Connection to different network after disconnection

Network Environment	With ANC and access profiles		Without ANC	Time difference between ANC detection and OS's IP configuration (sec)
	ANC detection time (sec)	IP configuration time (sec)	IP configuration time (sec)	
Layer-2 Switch	2.3	10.4	10.3 (i)	8.0
DHCP Server	6.1	8.1	8.0 (i)	2.0
Local Gateway	0.0	5.2	Never (0.0 with bad IP settings)	0.0 (ii)

- Note (i) Potential security exposure in cases with mismatched security settings (such as file sharing)
- Note (ii) Security exposure of 5.2 seconds before ANC access profile applied

IV. Evaluation

We repeatedly moved a NDA core from one location to another, and measured the time necessary for ANC to identify the network environment, and compared it with the time necessary for the OS to complete the IP configuration with and without using ANC's access profiles. In each network environment described in Section III-C, we tested 10 times with two scenarios: reconnecting with the same network, and connecting with a different network. We selected Microsoft Windows XP Professional as the OS running on the NDA, and prepared the IP parameter settings in each access profile in advance. We kept the cores disconnected from the network for about one minute before we connected them to the network again. Table II shows the summary of the results.

The ANC's network detection works quickly and accurately in most of the cases. Note that it works completely correctly in all of the connections to a different network. Two reconnection cases without the layer-2 switch are the exceptions, but with minor differences.

ANC detects the network environment more than two seconds faster than Microsoft Windows XP completes its IP configuration, and has enough time to apply the appropriate access profile defined for that location. The core is interested in the location detection in order to set its network and security parameters appropriately, especially when it moves to a different network. Microsoft Windows XP takes two to five seconds more before completing the IP configuration when it is connected to a different network. However, ANC detects quickly and in almost in the same manner, both when it is connected to the same network or to a different one.

ANC can monitor the MAC addresses of the layer-2 switches at an early stage of the connection setup, as described in the previous section. ANC is quite efficient for detecting the environments with layer-2 switches. Since the multicast packets from the layer-2 switch are very frequent, ANC can monitor them as soon as it detects the network connection. There are virtually no differences between the reconnection to the same network and the connection to a different one.

In the DHCP environments, Microsoft Windows XP finishes its IP configuration only after it receives a DHCP ACK packet from the DHCP server. In our test environments, Microsoft Windows XP tries to reuse the previous IP address settings. It skips a DHCP discover packet and starts with a DHCP request packet. Microsoft Windows XP takes about five seconds more when it is connected to a different network than when it is reconnected to the same network. Upon the reception of a DHCP NACK packet, it gives up its old IP settings and reinitializes its IP configuration with its new IP settings. ANC, on the contrary, can detect its DHCP server even with a DHCP NACK packet. Therefore, in the case of connection to a different network, ANC is triggered by a DHCP NACK packet, with which Microsoft Windows XP has to give up the previous IP settings. It can report the network environment about two seconds earlier, while Microsoft Windows XP is still waiting for a DHCP ACK packet in order to reinitialize its IP configuration. In the case of reconnection to the same network, ANC and Microsoft Windows XP are triggered by the same DHCP ACK packet. Microsoft Windows XP completes its IP configuration almost at the same time ANC reports the network environment.

In the local gateway environment, Microsoft Windows XP just enables the network with the fixed IP settings, while ANC broadcasts some ARP request packets. In the case of the reconnection to the same network, Microsoft Windows XP has already finished its IP configuration when ANC detects the network environment. However, in the case of the connection to a different network, ANC can detect the network environment, even when Microsoft Windows XP erroneously finishes its IP configuration. ANC can still suggest that Microsoft Windows XP settings should be corrected to match the current environment.

ANC's promptness in the network environment detection is more dependent on the network to be connected to than the OSs or the previous IP configurations. ANC always detects the network environment efficiently as long as it is connected to an Ethernet with a layer-2 switch. ANC also works reasonably well with a network with a DHCP server, because it can report the changes of the network environment quite early, before the OS receives the necessary DHCP information to configure its IP settings. As long as the computer is reconnected to the same network environment with the same IP parameter settings, it can keep its previous network and security settings, and it is not so interested in the event of the reconnection. For fixed IP settings, the network detection of ANC is still slow even with little chance of security exposure. ANC should request the network packet filtering functions to protect such environments, until it confirms that the current fixed IP settings are appropriate with the current default gateway.

V. Conclusion

We proposed the ANC technology in this paper in order to increase the network-friendliness of the consumer devices, such as the NDAs. ANC monitors link status and network packets, identify key network devices by their MAC addresses, provide network environment detection service for the NDAs, and help them quickly, safely, and accurately reconfigure as soon as they are connected to networks.

We evaluated ANC's capability for network environment detection, and found that ANC detected the network environments quickly and accurately in most of the cases, and contributed to the NDA's security and privacy by protecting confidential or personal information stored in the NDA. The detection time was mostly dependent on the target network environments, but it was not dependent on the previous network connections or the operating systems.

In order to promote the popularity of the NDA approach, we are enhancing the ANC technology so that it can deal with network performance and cost requirements in addition to the basic connectivity maintenance services and so that it can be used with innovative and advanced applications in addition to the basic configuration tools.

REFERENCES

[1] Y. Kushiki, "Digital consumer electronics evolution in the multimedia and network age," *IEEE Symposium on VLSI Technology*, 1999.

[2] S. White, J. Hanson, I. Whalley, D. Chess, and J. Kephart, "An Architectural Approach to Autonomic Computing, " *IEEE International Conference on Autonomic Computing*, 2004

[3] J. Kephart and D. Chess, "The Vision of Autonomic Computing, " *IEEE Computer*, Jan 2003, pp 41-50.

[4] T. Staudter, "The Core of Computing: Meta Pad slices and dices pervasive computing obstacles," *IBM Think Research*, 2002.

[5] J. Postel, "Domain Name System Structure and Delegation," *RFC 1591*, IETF, 1994.

[6] R. Droms, "Dynamic Host Configuration Protocol," *RFC 2131*, IETF, 1997.

[7] D. C. Plummer, "An Ethernet Address Resolution Protocol," *RFC826*, IETF, 1982.

[8] IEEE, *ANSI/IEEE Std 802.1D: Media Access Control Bridges*, 1998.

[9] R. Want, B. N. Schilit, N. I. Adams, R. Gold, K. Peterson, D. Goldberg, J. R. Ellis, and M. Weiser, "The ParcTab Ubiquitous Computing Experiment," *Mobile Computing*, 1995, pp. 45-102.

[10] N. Adams, R. Gold, B. Schilit, M. Tso, and R. Want, "An Infrared Network for Mobile Computers," *Usenix Symposium on Mobile and Location-Independent Computing*, Summer 1993, pp. 41 -51.

[11] R. Want and A. Hopper, "Active badges and personal interactive computing objects", *IEEE Transactions on Information Systems* 38(1), 1992, pp. 10-20.

[12] A. Harter and A. Hopper, "A Distributed Location System for the Active Office," *IEEE Network Special Issue on Distributed Systems for Telecommunications*, Vol. 8, No. 1, January 1994, pp. 62-70.

[13] T. E. Truman, T. Pering, R. Doering and R. W. Brodersen, "InfoPad Multimedia Terminal: A Portable Device For Wireless Information Access," *IEEE Transactions on Computers*, Volume 47, No. 10, October 1998, pp. 1073-1087.

Yasuharu Katsuno received a B.S. degree in electronics engineering in 1996 and an M.S. degree in computer science in 1998, both from Keio University. He joined the Tokyo Research Laboratory in 1998. He was involved in research on wireless communication and computer network..

Toru Aihara (M'83) received a B.S. degree in electronics engineering in 1983 and an M.S. degree in electrical engineering in 1985, both from Tokyo University. He joined the Japan Science Institute of IBM Japan (now the Tokyo Research Laboratory) in 1985. He was involved in research on low-power systems and wireless communication, and also in standardization of the Bluetooth technology.