# Oversized Subnet and Shared NAT: A Practical Approach to Keep Private and Public IP Addresses Together

*Il Hwan Kim and Heon Young Yeom*

School of Computer Science and Engineering, Seoul National University

**Summary**
The growth of IPv4 Internet has been facing the infamous IP address depletion barrier. In practice, many IPv4 Internet edge networks are forced to be expanded by incorporating private addresses and NAT devices. In this paper, major limitations of NAT-expanded private networks are identified. Furthermore, a solution is proposed to encourage the mixed usage of private and public IP addresses in a single edge network domain. The solution comprises of two key ideas: oversized subnet and shared NAT. The oversized subnet removes the routing boundary between private and public hosts. Shared NAT saves public IP address resources by sharing them among several private networks. These ideas not only encourage the coexistence of heterogeneous address classes, but also lead to efficient sharing of global IP addresses

***Key words:***
*Network Address Translator, IPv4 Address Reuse, Residential Network*

## 1. Introduction

Internet has grown up into a worldwide and common infrastructure for data communication. For example, in Korea, xDSL and HFC (Hybrid optical Fiber and Coaxial cable) connections have been supplied at reasonable costs, with aggressive marketing campaigns starting from high-density metropolitan areas. At 2002, the number of domestic Internet subscribers reached over 10 millions, which covers virtually every household in Korea [2]. Until 2003, the cumulative number of public IPv4 addresses allocated to KRNIC amounts to 30 millions. Most of these addresses are divided into subnetworks smaller than the C class network, with /24 or longer netmask. DHCP is commonly used to maintain the dynamic IP address pool. Furthermore, DHCP is also utilized to authenticate and inform the client host with regard to the minimal network configuration such as edge router and DNS server addresses.

As such a prevailing and cost-effective data communication medium, Internet is now recognized as the infrastructure for communication. For instance, all-IP based home network, triple play, and nation-wide telematics are carefully explored as the most promising services despite of its not-totally-reliable nature. These application services need terminal devices to be connected to Internet; these devices need two to ten times more IP addresses than those actually in use today.

Although the global IP address depletion is a menace to development of new data services, Internet service providers tend to regard IP address problem as what should be resolved by the address authorities, not as something that can be solved with their own ingenuity within the given constraints.

IPv6 may be the ideal solution to the shortage problem. But the global transition is deterred by several obstacles such as inter-operability and investment protection. In practice, many IPv4 Internet edge networks are individually expanded by subscriber-initiated deployment of private addresses and NAT (Network Address Translator) [12] devices.

## 2. Nat-Expanded Private Networks

With an increasing number of TCP/IP devices, individual Internet service subscribers build their own LANs to facilitate the inter-communication of their equipments. Under an Internet line subscription with fixed charge policy, a personally installed NAT and switch are the cheapest solution to provide several TCP/IP devices with global Internet accessibility. In this paper, such a form of network will be called as NAT-expanded private network.

A kind of NAT, known as PAT (Port-Address Translator) [17] or NAPT (Network Address-Port Translator), is commonly used to share a single IP address among several hosts. The typical implementation of PAT consists of two phases comprising the stateful packet filter and the header manipulator. On each incoming packet, the first phase is to classify them and identify the session based upon its source and destination IP address and TCP or UDP port number. Once the session is established, header fields of the next packets can be efficiently manipulated according to the cached session-tracking information before being forwarded into the next routing hop.

One of the most popular implementation of NAT is known as Netfilter [1] which is an integrated part of Linux TCP/IP stack. Because virtually every NAT appliances

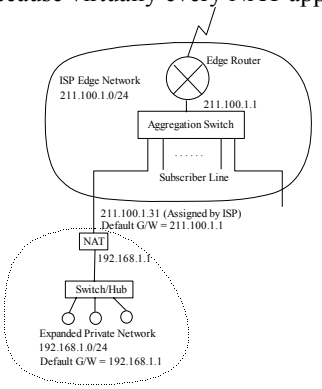


Fig. 1. A typical edge network example, expanded with NAT. The dotted line denotes a premises network expanded with a privately installed NAT device.

available in the market are capable of PAT/NAPT function out of the box, NAT indicates PAT/NAPT through the rest of this paper.

The most common usage of personally installed NAT and expanded private network can be found in small offices as well as at homes, where rather small number of hosts require Internet connection. A typical combination of NAT and private network is shown in Fig. 1. The Internet subscriber is provided with a single line and a public, globally routable IP address. The public IP address is often allocated dynamically by DHCP. The IP address is assigned to WAN interface of NAT, and arbitrary private network addresses are assigned to internal hosts. Those internal hosts share the line and public IP address of NAT to reach public Internet hosts. The NAT device, commonly referred to as an IP-sharing router, is often integrated with L2 switch, hub, wireless AP, or residential gateway.

## 2.1 Legacy Mixed NAT-expanded Networks

In a NAT-expanded private network, private hosts are normally invisible to and not accessible by the outside public hosts beyond the NAT. Internal hosts can initiate an outbound session through the NAT, but from the foreign hosts' point of view, they are identified with the public address assigned to the NAT. The NAT maps a unique TCP/UDP port number to each session in order to distinguish one from the other.

In order to serve inbound requests on a private host, a globally accessible TCP/UDP port number on the NAT should be bound to the listening port of a server process running on the private host. The incoming packets arriving on the reserved port of the NAT are forwarded to the private server process, by translating the destination address and port number fields to the appropriate values. This is called DNAT (Destination-NAT) in Linux/Netfilter implementation. Another way to offer free and unrestricted accesses to an internal host is the DMZ (De-Militarized Zone) configuration. DMZ can be implemented in two flavors, which are static mapping or bypassing. The static mapping approach is rather simpler than the other one, in which a set of public addresses are mapped to a set of local servers statically and exclusively.

In the bypassing approach, an incoming packet is forwarded without any header manipulation, as long as the port is not reserved to another private host. The NAT acts as a usual router and forwards the packet to the local server host. In this approach, a publicly routable IP address should be assigned to the local server host, and the NAT should distinguish the packets for the DMZ server from those for the private hosts. To reduce the public address occupation, many NATs are designed to serve the private hosts with the same public address as the special local server host. An exemplary configuration of such a network is shown in Fig. 2.

Fig. 2 illustrates the bypassing NAT and DMZ [13] configuration. When a packet arrives at the WAN interface of the NAT, it is classified according to predefined filtering rules. If it is destined to the DMZ host, they are forwarded without header manipulation.

Some NATs allow the built-in L2 switch ports to be divided into several partitions by VLAN tagging. In this case, one of the partitions can be declared as DMZ segment, in which all the traffic is bypassed through the WAN interface. As a consequence, the hosts on the DMZ segment are considered the same as the outside public hosts and treated accordingly. Because those hosts directly reach the edge router through an L2 connection, each host in that segment occupies at least one public address. Furthermore, the private hosts and the local servers in the same premises can communicate with each other only through the inefficient network address translator.

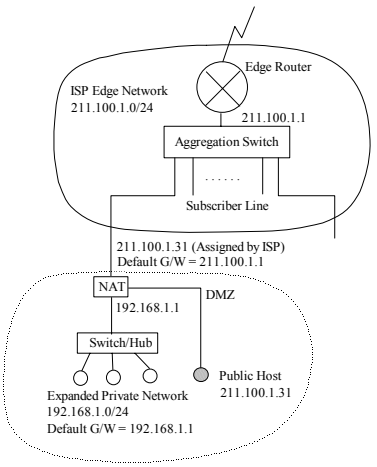

Fig. 2. A legacy NAT-expanded private network. A public host is mixed in the DMZ port.

## 2.2 Limits of Legacy Mixed NAT-Expanded Networks

In any legacy configurations, mixing private and public hosts can incur cumbersome problems with regard to arrangement and inter-operability.

i) Longer occupation time of public addresses:
The NAT consumes at least one public address per one premises network and the occupation tends to be long. Thus, the ISPs must prepare much larger pool of public IP addresses than NAT-free subscriber networks. Under a pricing scheme with fixed-rate charge per dynamic IP address, more and more subscribers want to use NAT. The longer and higher IP address occupation will be a menace to ISPs.

It is noticeable that only a few public addresses are sufficient to provide all of the individual private hosts in a edge network with the outbound one-way NAT service.

ii) The lack of direct routes between the private and public hosts:
The private and public hosts have different subnets from each other. Even if they are at the same physical location, they cannot communicate directly through the local L2 switch. Instead, the NAT is in charge of the communication between them. The packets should be routed or translated properly according to the predefined rules.

The lack of direct routes may pose serious problems in some circumstances. Suppose a case where a residential network consists of an IPTV set-top-box, a home automation server, and a few PCs. It would be enough to assign private addresses to the PCs, because they rarely serve any requests from the outside of the premises; the set-top-box may require a public address since the contents are delivered through multicast protocols or downstream push mechanisms; the home automation server also requires public address to respond the calls from the outside. Because the public hosts and private hosts cannot communicate with each other directly, the premises NAT will get loaded with internal traffic as well as WAN traffic. Also, the overall performance and availability of the local network is limited by the NAT.

In order to avoid this problem and to let the hosts recognize each other as neighbors, static routes may be added to their routing tables. In this way, they can reach each other without the intervention of a router or a NAT. However, this work-around is not applicable in the following cases. First, dynamic public addresses acquired by DHCP change, which imposes the impossible job of changing the static routing table of each host every time; Second, some devices such as SIP phone or IP set-top-box may not accept the direct routes due to their limited TCP/IP implementations.

iii) Inherent limitation of performance:
NAT is as complicated as a router or firewall with stateful packet inspection ability. In addition to generic IP router, NAT performs the following operations upon packet arrival: packet classification, session identification, maintenance of the port-address mapping table, header manipulation, and so on.

Generally speaking, the complexity and hardware cost of NAT implementation is comparable to that of L4 switches rather than L2 switches. Its performance can be hardly better than that of L3 switches under the same level of hardware cost constraints.

iv) Application layer gateways and proxies:
A NAT is loaded with various ALGs (Application Layer Gateway) and proxies to support NAT-unfriendly protocols. It is clear that a NAT platform must be flexible enough to run these software modules and that they are maintained regularly to be updated. Most NAT developers, however, pay little attention to the maintenance of the software. In addition, the software nature of ALGs and proxies poses even more penalties on the performance of NAT.

Although there are a few novel approaches like STUN [10] to the NAT-traversal problem of UDP-based protocols, they cannot remove the necessities of ALGs and proxies completely.

In spite of these drawbacks, NAT is considered as one of crucial components to some parts of Internet, such as residential network. NAT technologies have developed through years, they are not costly, and personal NAT-expanded private networks functions well with small numbers of internal hosts.

With a few fundamental improvements, the mixed NAT-expanded networks can achieve wider deployment with much higher level of services. At first, the public addresses can be saved by deploying a shared NAT server for the private hosts in one whole edge network. Secondly, the direct routes within the same premises ensure much smoother inter-operations between the private and the public hosts. And finally, better overall performance can be achieved as more packets avoid the NAT bottleneck as possible.

## 3. Oversized subnet and Shared NAT

### 3.1 Oversized Subnet

It has been discussed that the lack of direct routes between local hosts is one of the major drawbacks of the legacy mixed NAT-expanded network.

The most straightforward solution to this problem is to assign an oversized subnet mask to local hosts. If a network address is large enough to cover all the private and public addresses in the premises, such an oversized network address is called an oversized subnet. As an example of an oversized subnet, 128.x.x.x/1 network embraces a public network such as 211.x.x.x/24, as well as a private network such as 192.168.x.x/24. Thus, /1 is an oversized subnet mask for 211.x.x.x/24 and 192.168.x.x/24 network. All the public hosts within 211.x.x.x/24 range and the private hosts within 192.168.x.x/24 range belong to the oversized subnet 128.x.x.x/1. If these hosts are assigned with the oversized subnet mask /1, they consider each other as neighbors having the same broadcast domain. Therefore, they will try to reach directly by resolving their MAC addresses through ARP. If they are placed in the same L2 segment, they can directly route to each other without any intervention by other routers or NATs.

An example of an oversized subnet organization is illustrated in Fig. 3. In the oversized subnet organization, the edge router is designated as the default router for the public hosts. And the NAT is assigned as the default gateway for the private hosts. No additional configurations are required in an oversized subnet.

Under an oversized subnet organization, a host divides its destinations into two categories. One of them is in the same broadcast domain, and the other is the outside of its broadcast domain. The routing decision is very simple: try ARP or rely on the default router.

In Fig. 3, A and B consider 128.0.0.0/1 network is local, and 0.0.0.0/1 is outside. Because their view is erroneous and mismatches the real world, an ARP proxy must guide them to an appropriate router. The private host A must forward its outbound packets to a NAT, and the public host B must forward them to the edge router.

### 3.2 Shared NAT

When a private host A in Fig. 3 wants to communicate with the global Internet hosts, its address should be translated into a public one by a NAT. Because the oversized subnet organization removes the logical boundary between each premises network, a single NAT server is sufficient to serve all the private hosts within the edge network.

A shared NAT serves the private hosts with a route to the global Internet, in the similar way as an edge router serving the local public hosts. A large number of individual NATs can be replaced by only one shared NAT. Hence, the total number of public IP addresses occupied by NATs can be reduced significantly across an edge network.

### 3.3 Oversized subnet ARP proxy

In an oversized subnet, the public and private hosts in the edge can communicate with each other directly through ARP. On the other hand, the oversized subnet mask imposes an erroneous view with regard to the network topology. The local hosts will broadcast an ARP query if the destination IP address is within their broadcast domain.

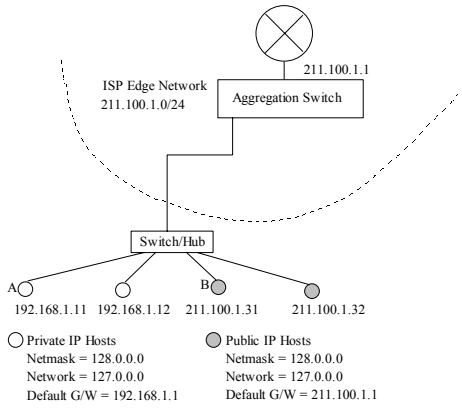


Fig. 3. An oversized subnet organization of public and private IP addresses within the same edge network.

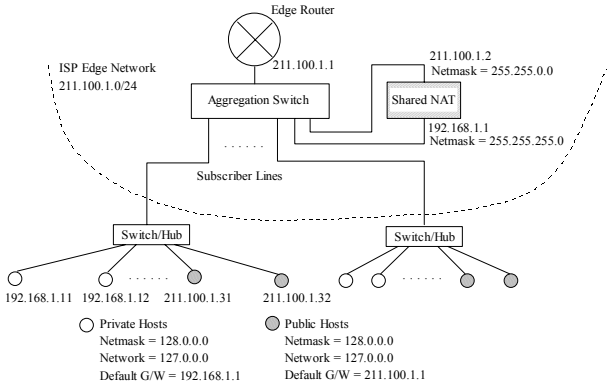


Fig. 4. A shared NAT in the oversized subnet organization.

If the destination host is actually located outside of the edge network or on a different L2 segment, they cannot reach the destination. Thus, an ARP proxy should respond to an ARP request for those hosts which belong to the same oversized subnet but which, in actuality, exist outside of the edge network in order to direct the packets to an appropriate forwarding host, such as an edge router or a shared NAT.

The special version of an ARP proxy is called an oversized subnet ARP proxy. Its responsibility is to listen to all the ARP requests flowing through the edge network, and provides answers to them if necessary. The oversized subnet ARP can be placed anywhere in the edge network as far as it can listen to the broadcasted ARP requests from the local hosts. Note that they can be replicated and co-work with each other to reduce the response time. The ARP proxies can be strategically placed to minimize the risk of ARP storm and the impact of misbehaving clients.

Table 1: Address resolution in oversized subnet

| Source | Destination | | |
|---|---|---|---|
| | **Local Public** | **Local Private** | **Foreign** |
| **Local Public** | N/A | 1) | 2), 3) |
| **Local Private** | 1) | N/A | 5), 6) |
| **Foreign** | 4) | 7) | N/A |

The ARP proxy resolves the oversized subnet routes according to Table 1. The local hosts can communicate with any type of global Internet hosts only by forwarding packets to the MAC address replied by the proxy. The completeness of this table is analyzed below.

1) Local private host (A) ↔ Local public host (B):
If A and B have the same oversized subnet address 128.0.0.0, they consider each other in the same broadcast domain. They acquire the peer MAC addresses via a usual ARP transaction. The packets reach each other without any intervention of the ARP proxy. The ARP proxy must keep silent in this case.

2) Local public host (B)→ Foreign hosts beyond the oversized subnet range (not matching to 128.0.0.0/1):
If a public host B 210.100.1.10 tries to reach a foreign host 65.1.1.1, it computes the subnet prefix as follows: 65.1.1.1/1 = 0.0.0.0. Because it means that the destination is beyond the local network, B forwards its packets to the default gateway. The MAC address of the default gateway is resolved via a usual ARP transaction. In this case, the ARP proxy must keep silent so that the default gateway answers to the query. The default gateway takes care of all the packets according to the usual IP forwarding procedure.

3) Local public host (B) → Foreign hosts within the oversized subnet range (matching to 128.0.0.0/1):
If a public host B 210.100.1.10 tries to reach a foreign host 230.1.1.1, it computes the subnet prefix as follows: 230.1.1.1/1 = 128.0.0.0.Because this means that the destination is inside the local network, B tries to acquire the MAC address of the destination host. But this decision is false, owing to the illusion of oversized subnet. In this case, the ARP proxy must answer the MAC address of the default gateway in response to the query. Then B believes that it can reach the destination by forwarding the packets to the designated MAC address. The default gateway then forwards them according to the usual IP forwarding procedure.

4) Foreign hosts → Local public host (B):
If a packet from a foreign host arrives at the edge router, it looks up the routing table to determine the forward path. If the packet is destined to a local public host B, it is forwarded following the normal IP forwarding procedure. The MAC address of the public host can be resolved via a usual ARP transaction because the edge router and B can directly talk to each other with a layer 2 protocol. In this case, the ARP proxy must keep silent.

5) Local private host (A) ↔ Foreign hosts beyond the oversized subnet range (not matching to 128.0.0.0/1):
If a local private host A 192.168.1.20 tries to reach a foreign host 65.1.1.1, it computes the subnet prefix as follows: 65.1.1.1/1 = 0.0.0.0. Because it means that the destination is beyond the local network, A tries to forward the packet to its default gateway address, a nearby shared NAT. The MAC address of the shared NAT is resolved via a usual ARP transaction. In this case, the ARP proxy must keep silent so that the shared NAT answers to the query. The shared NAT manipulates the packet headers following the usual NAT procedure, and then forwards it to the edge router.

The incoming packet which belongs to the same session will be handled vice versa. The edge router forwards it to the shared NAT, the shared NAT manipulates the headers as usual de-NAT procedure, and finally forwards it to the local private host. It is obvious that the packet is a directly transmitted between the private host and the shared NAT, without the ARP proxy stays out of all these procedure.

6) Local private host (A) ↔ Foreign hosts inside the oversized subnet range (matching to 128.0.0.0/1):
If a private host A 192.168.1.20 tries to reach a foreign host 230.1.1.1, it computes the subnet prefix as follows: 230.1.1.1/1 = 128.0.0.0. Because this means that the destination is inside the local network, A tries to acquire

the MAC address of the destination host. But this decision is false, owing to the illusion of oversized subnet mask. As in case 3, the ARP proxy must answer the MAC address of the shared NAT in response to the query. Then A believes that it can reach the destination by forwarding the packets to the designated MAC address. The packet is forwarded to the shared NAT, according to the ARP result. The rest of the communication process is the same as Case 5.

7) Foreign hosts → Local private host (A):
The private address of local private host A is unknown to the foreign hosts, the inbound connections are impossible in general. Instead, a port on the shared NAT can be reserved to enable the inbound connections. When a new packet arrives on this port, it is translated and forwarded to the corresponding port of the private host. In this way, as known as DNAT [1] or port-forwarding, the local private host can become a server to the foreign hosts. The ARP proxy is independent of the port-forwarding procedure; it must remain silent.

Now it is evident that hosts with oversized subnet can communicate with each other as well as with foreign hosts, thanks to a simple ARP proxy.

## 4. Implementation Considerations

A few remarkable considerations on the implementations of oversized subnet ARP proxy and shared NAT are as follows.

### 4.1 ARP Proxy on Residential Gateway

By nature, an ARP proxy can be located anywhere as long as it is provided with direct layer 2 connections to the local hosts. One of the most attractive insertion points for ARP proxy is the residential gateway.

If an RG (residential gateway) is to be installed in front of each residential network, the ARP proxy can be integrated into it. The RG may also include agents for management, a DHCP relay, a DNS proxy, and some security modules. In cases, an IP address may be required to access the RG remotely. Under the oversized subnet organization, one can assign a private address to the RG without wasting any public addresses.

To be friendlier to the oversized subnet organization, an RG may include a simple packet filter which protects the edge network from broadcast storms and misbehaving hosts. Although it may also be utilized as a security gateway to protect the residential hosts, it is an orthogonal issue for the current concerns.

Unlike NAT-based RGs, an oversized subnet RG can be implemented as a pure L2 device. This is the biggest advantage over NAT-based RGs. The typical hardware constitution resembles that of a L2 switch. Because the ARP packets can be easily classified with a simple bit-comparison circuitry, a very low-powered CPU is sufficient to employ the ARP proxy and broadcast barrier. The ARP proxy is obligated to collect the MAC addresses of some special hosts like edge routers and shared NATs. In necessity, the activities of the ARP proxy can be explicitly controlled through the management agent.

### 4.2 Shared NAT and Local Servers

The primary purpose of a shared NAT is the IP-sharing service for private hosts, and the second one is the port-forwarding service. A shared NAT can be implemented in a very similar way to generic NATs. Since there are many types of NATs available, each ISP can choose one version and integrate rich functionalities such as UPnP IGD and SOCKS server, if needed.

Shared NATs are expected to reduce the permanent occupation rate of public addresses, just like as the web proxies reduce the waste caused by the static and redundant web traffic. Because the Internet availability and performance of private hosts depend on the shared NAT, ISPs should consider building up a high performance cluster of shared NATs on their edges.

In the oversized subnet architecture, the barrier between local private hosts and public hosts is removed. They share the same broadcast domain. With this freedom given, various local servers can run on a par with the shared NAT. For example, a video-on-demand server may run over efficient L2 multicast protocols.

## 5. Discussions

The advantage of oversized subnet can be summed up as follows.

Compared to the legacy NAT-expanded networks, public IP addresses are less wasted and need not to be occupied all the time. The public IP resource can be saved for more precious servers instead of wasted by individual NATs.

No topological changes are required to legacy edge networks. The ARP proxy and shared NAT are sufficient to removing the redundant usage of individual NATs. To protect the edge network from anomalous broadcast packets, it is desirable to make use of a simple filtering bridge along with the ARP proxy.

All hosts within the oversized subnet can directly communicate via L2 links. The inefficiency caused by individual NATs is ameliorated by direct routing between local hosts.

High performance bridge-based residential gateways can be implemented with lower costs compared to NAT-based ones. Contrary to that NAT-based RGs should be loaded with powerful processor and complicated software, bridge-based RGs can be implemented with a few simple extensions to the L2 switches. They are more competitive to the simple L2 switches than NAT-based ones, with respect to the performance.

The better L2 inter-operability leads to easier QoS assurance. Moreover, the full advantages of various L2 technologies are available to the local network, without the interference of the L3 barriers induced by individually placed NATs.

Some heavy duty servers can be migrated into the edge network, along with the shared NAT. Especially, real-time applications and multimedia servers are good candidates to be localized. The high performance RG efficiently delivers the traffic between the localized servers and clients.

As the future works, a new framework for the home network service can be developed upon the oversized subnet architecture. For an example, an IP address allocation scheme with improved DHCP services should be established. The impact of the oversized subnet scheme on the QoS and security must also be analyzed and addressed. The last, but not the least significant challenge is a management platform which would be required to harmonize the components such as ARP proxies, residential gateways, shared NATs, and local servers.

## 6. Related Works

There are some related works with regards to address reuse and NAT extensions.

The restrictions on subnetting scheme and classful IP routing mechanism are alleviated by classless inter-domain routing [14], which is also published as RFC 1519 [7].

A technique to multiplex a public IP address into multiple private addresses is proposed in [17]. PAT is a natural extension to the classic 1:1 NAT techniques. Various aspects of NAT technologies are aggregated in RFCs during years. The most recent one is RFC 3022 [12] at the time of this writing. The Netfilter [1] in Linux kernel, featuring powerful session tracking engine and stateful packet filter, is the most widespread implementation among users and developers. RFC 1918 [9] suggests guidelines on the address allocation for private networks.

In [11], it is pointed out that the penetration rate of residential networks depends on the abundance of IP address resource and the cost price of Internet subscriber lines. The economy around the IP address and cost is one of the main motivations to our research.

MobileNAT [5] applied the NAT techniques to facilitate the mobility of wireless IP networks. The extensive use of DHCP is also applicable to other NAT related works.

As a method for address sharing that exhibits more transparency than NAT, RSIP [4] is proposed and published as RFC 3103 [3]. RSIP requires hosts to be modified in order to interact with RSIP gateways.

In [6], B. Ford made some notices on the ad-hoc network routing crisis that may arise around the edge network, and suggested a self-managing routing protocol to alleviate the problems.

Another point of view on the NAT problem is that it makes the peers confused about their identities by implicitly transforming the address headers. As a consequent, new routing and tunneling protocols that play the similar role as NAT are proposed. A few explicit identity-based routing mechanisms are known so far [8] [15] [16].

Protocols based on UDP datagram are inherently inept to be friendly with NATs. A rather transparent and universal solution, if not panacea, to the NAT-traversal problem of UDP-based protocols is STUN. The STUN is also proposed as an RFC 3489 [10].

## References

[1] Linux 2.4.x netfilter homepage. http://www.netfilter.org
[2] “2003 Korea internet statistics yearbook”, Korea Network Information Center, 2003.
[3] M. Borella, D. Grabelsky, J. Lo, and K. Taniguchi, “Realm Specific IP: Protocol Specification”, IETF RFC 3103, October 2001.
[4] M. Borella and G. Montenegro, “RSIP: Address sharing with end-to-end security”, Special Workshop on Intelligence at the Network Edge, San Francisco, 2000
[5] Milind Buddhikot, Adiseshu Hari, Kundan Singh, and Scott Miller, ”Mobilenat: a new technique for mobility across heterogeneous address spaces”, In WMASH ’03: Proceedings of the 1st ACM Int. workshop on Wireless mobile applications and services on WLAN hotspots, pp. 75-84, New York, 2003
[6] Bryan Ford, “Unmanaged internet protocol: taming the edge network management crisis”, SIGCOMM Comput. Commun. Rev., 34(1):93-98, 2004
[7] V. Fuller, T. Li, J. Yu, and K. Varadhan, “Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy”, IETF RFC 1519, September 1993
[8] P. F. Ramakrishna, “IPNL: A nat-extended internet architecture”, In SIGCOMM '01: Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 69-80, New York, 2001
[9] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, “Address allocation for private Internets”, IETF RFC 1918, February 1996
[10] J. Rosenberg, J. Weinberger, C. Huitema, and R. Mahy, “STUN - simple traversal of user datagram protocol (UDP)

through network address translators (NATs)", IETF RFC 3489, March 2003
[11] P. Savola, "Migration and co-existence of ipv4 and ipv6 in residential networks", Seminar on Internetworking, Helsinki University of Technology ,2002
[12] P. Srisuresh and K. Egevang, "Traditional IP network address translator (Traditional NAT)", IETF RFC 3022, January 2001.
[13] P. Srisuresh and M. Holdrege, "IP network address translator (NAT) terminology and tonsiderations", IETF RFC 2663, March 1999.
[14] P. F. Tsuchiya, "The landmark hierarchy: a new hierarchy for routing in very large networks", In SIGCOMM '88: Symposium proceedings on Communications architectures and protocols, pp. 35-42, New York, 1988
[15] Z. Turányi, A. Valkó, and A. T. Campbell, "4+4: an architecture for evolving the internet address space back toward transparency", SIGCOMM Comput. Commun. Rev. 33(5):43-54, 2003
[16] M. Walfish, J. Stribling, M. Krohn, H. Balakrishnan, R. Morris, and S. Shenker, "Middleboxes No Longer Considered Harmful", USENIX OSDI 2004 San Francisco, December 2004
[17] H. Y. Yeom, J. S. Ha, and I. H. Kim, " IP multiplexing by transparent port-address translator", In Proceedings of the 10th Systems Administration Conference (LISA'96), pp. 113-121, 1996.

**Il Hwan Kim** is a PhD candidate in Distributed Computing Systems Laboratory of Seoul National University. He received his BS and MS degrees in Computer Science from Seoul National University, Korea, in 1994 and 1996. His research interests include computer networks and distributed systems, and his current focus is on residential networks and peer-to-peer systems.

**Heon Y. Yeom** is a professor in the Department of Computer Science and Engineering, Seoul National University, South Korea. He received a BS from Seoul National University majoring in Computer Science in 1984 and an MS and a PhD in Computer Science from Texas A&M University in 1986 and 1992, respectively. His research interests include multimedia systems, distributed systems and fault-tolerant systems.