マルチメディア通信と分散処理　62－1
（1993. 9. 20）

構成発見の世界

村山 優子
WIDE プロジェクト

〒182 東京都調布市調布ヶ丘 1-5-1
電気通信大学情報工学科砂原研

あらまし

様々な環境下において, 興味の対象となる物の存在とそれらの位置について知りたいという要求が発生する. 本論文では, これを構成発見と呼び, そのコンピュータ通信の網間網, すなわちインターネット環境下における問題点を探る. 構成発見は, 何らかの形で, 現在の環境構成についての情報を得ることで達成できる. この１０年の間にインターネットはめざましい規模増大を遂げてきた. 大規模環境についての情報管理は動的学習や管理域分割により解決できる. しかし, どちらの手法においても情報と現状の一貫性や情報の不正確さの問題がおきる. この２つの観点から正確な情報の必要性を説く.

和文キーワード　コンピュータ通信網, 構成管理, 構成発見, 網管理, 安全性, 情報管理

Configuration Detection and Its Problems

Yuko Murayama

WIDE Project

c/o Hideki Sunahara
Dept. of Computer Science
The Univ. of Electro-Comm.
1-5-1, Chofugaoka, Chofu City, Tokyo 182

Abstract

In any environment, one would need to know which objects of interest exist and where they are located. We call the acquisition of this knowledge Configuration Detection. We are interested in hosts and routers on an internetwork. The growth in the number of network objects has been substantial and dynamic for the last decade. The scale problem has been solved either by facilitating dynamic learning or by dividing the management domain into subdomains of a manageable size. In either solution, we face the problems of inconsistency and invalidity. We look into the serious consequences from these two problems.

英文 key words　Computer Networks, Configuration Management, Configuration Detection, Network Management, Security, Information Management

1 Overview and Motivation

Just as the workings of our society are based on information, so also are the operations of computer networks. We call a system whose operation is based on information an information-based system. How well the system works depends on the *credibility* of the information used. Information has long been treated as data in computer networks, and much research has been carried out on how well networks can carry data from one end to another. Shannon identified this level of research as the engineering level [19]. This paper is concerned with one quality of the semantic level of such information; is the information being transmitted actually correct?

The most significant aspect of computer network evolution in the past decade has been its growth. The open architecture of networks [10] [4] and packet switching technology [1] have contributed not only to the development of single networks, but also to their interconnection; the latter is called internetworking. Earlier work on layering concepts for networking [17] provided a great contribution to the success of interconnection of networks. As a result, the size of networks has grown so rapidly in an unexpected way that existing management mechanisms now have difficulty coping, being originally designed to manage only a limited number of objects. One of the problems is knowing what objects exist and where they are located. We term this novel problem *configuration detection* in this paper.

In this paper, we look into the problems of configuration detection in computer networks [14]. We are concerned particularly with network objects such as hosts and routers, and the configuration information relating to their location, such as names, addresses, and routing information.

The paper organisation is as follows. The next section describes the problems of configuration detection. Section 3 describes serious consequences. and Section 4 gives some conclusions.

2 Problem Description

In Shoch's terms [20] one needs to have the knowledge of the following for computer networking activities;

what objects exist,
where they are, and
how one can reach them.

We refer to this knowledge *configuration information* in the rest of this paper. Network monitoring systems have to know which objects to monitor, where they are, and how they can be perceived from the monitors [13], [8]. Management systems require the same knowledge but more crucially, as they need to exercise control over objects. For an electronic mail system, when one sends a message, the least one has to know is who could be the receiver and its address. The routing function needs to know what intermediate systems are available to what destinations. In order to make our discussion simpler, we presume that the objects of interest are hosts and routers in a specific environment, the Internet Protocol (IP) over subnetworks such as Ethernets belonging to one organisation, and that their IP network addresses are the information of concern. In the rest of the paper, we mean a link address by a subnet-level address, whilst we mean an IP address by just an address or a network address.

Traditionally there are two types of systems which are responsible for the maintenance of configuration information: network management systems and information systems. A network management system maintains the location information which is used for management activities of the network environment, such as fault detection and isolation, the installation and removal of a nodal system, and status monitoring. The location information in the management systems has hitherto been managed manually by skillful and knowledgeable operators. This scheme works so long as the size of the network is limited. However, with the growth experienced in wide area and local area networks in the last decade, one cannot hope to continue to cope with the volume and complexity of the information by such manual means.

An information system has been responsible for pro-

viding each participant of network operations with the location information. Name servers and directory services are examples of the static type of information systems, whereas the routing information exchange system [18], the address resolution system [16], and the distributed bridge learning system [9] are of the dynamic type. The former have coped with the growth-in-size problem by introducing hierarchically decentralised management [12], [3], [5], so that the size of local object tables remains manageable — this is the direction that network management systems are following. The latter have adjusted to the new conditions by introducing dynamic learning mechanisms.

Whichever approach is taken, existing systems face two problems: *inconsistency* and *invalidity*. *Inconsistency* is the state of difference existing between registered information and actual configuration; it is caused by the fact that installation and removal of hosts can be done without registration, or vice versa. *Invalidity* includes the incorrectness of registered information and the misconfiguration of a host; it is caused by the fact that there is no verification mechanism.

These shortcomings could lead to various consequences and threats at various levels. The application level threats in such circumstances are unauthorised access to other hosts, masquerade, unauthorised access to servers, and unauthorised disclosure of information; the network level threats are unauthorised tampering with routing and control data, unauthorised use of resources, unauthorised traffic generation, and unauthorised disclosure of information. The former would be best countered at the application level by authentication and access control. In this paper, we are concerned particularly with the network level threats. We are also interested in a serious consequence from misconfiguration of a subnet-level address possibly by inadvertent errors. We describe both the network level threats and a serious consequence of misconfiguration in the following section.

3 Serious Consequences

3.1 overview

In this section, we describe serious consequences from inconsistency and invalidity. Firstly, we present a serious consequence from misconfiguration of a subnet address. Secondly, we discuss the network level threats. Finally, we examine them and identify the most serious ones to be dealt with.

3.2 A consequence from misconfiguration

A serious consequence from misconfiguration of a subnet address is as follows. Another type of serious misconfiguration is that host is configured with both the router's or server's subnet level address, (i.e. a duplicate subnet level address), and the forwarding capability which is usually for routers. The amount of the traffic destined to the authentic router or server will be doubled; the misconfigured host will receive all the packets destined to the authentic host, and forward them to the authentic one. If the misconfigured address is that of a router, packets may travel beyond the organisational networks causing unnecessary traffic over the inter-organisational networks; this could be serious as the resources of other organisations are consumed in vain. Likewise, the packets destined to the misconfigured host will be forwarded by the authentic router. If the address is that of a server, the traffic towards the server may cause congestion at the server's interface.

A network storm could be created when three or more routers, or misconfigured hosts with the forwarding capability, are configured with the same subnet level address with different network level addresses. The following sequence was identified by Perlman [15]. Suppose there are three routers, A, B, and C; note that this could happen also by having hosts misconfigured with a forwarding capability. A packet is sent out with the destination X which should be forwarded to the router C. The packet is also received by A and B as they have the same subnet level address. As the destination network level address is neither A nor B,

they forward to C; the packet forwarded by A is also picked up by B as well as C, and B then forwards it, it will be picked up by A and forwarded again. In this way, each packet sent to one of these three routers will be multiplied by two others until they finally expire at the timeout specified in the time to live (TTL) field of each packet.

Generally, if there are n routers and misconfigured hosts with the same subnet level address, a packet with a TTL value, t, sent to one of them will be regenerated by the other $n-1$ routers and hosts for the first round; one of the n routers is a correct router for this packet and would forward it off the subnet. From the second round on, each of those $n-1$ regenerated packets will generate $n-2$ others. In this way, packets are generated exponentially. Eventually, the network is flooded with messages, and the only way to halt the traffic is to remove the misconfigured routers and hosts. The number of regenerations, S, is expressed as follows:

$$S = (n-1) + (n-1)(n-2) + \ldots + (n-1)(n-2)^{t-2}$$

$$= (n-1)\sum_{i=1}^{t-1}(n-2)^{i-1}$$

which may be expressed as follows:

$$S = \frac{(n-1)\{(n-2)^{t-1}-1\}}{n-3}$$

Either a large n or t would be enough to cause the network to be flooded; the TTL could be specified as large as 255 — the maximum TTL, as the TTL field is one octet long. Surprisingly, some systems on a network are sending and receiving the Network File Store maintenance packets with the TTL value of 255.

This effect could be invoked easily by the misconfiguration of a host with a subnet level broadcast address, because all routers take packets sent with the subnet level broadcast address, and try to forward them. This is called "a chain reaction" coined by Manber [11] to characterise the phenomenon that a design failure in a computer or in a network protocol will propagate throughout a network, causing a breakdown of the entire network.

3.3 Threats

In this section, we describe the threats at the network level. They are listed as follows:

1. **Unauthorised tampering**
2. **Unauthorised use of the resource (resource stealing)**
3. **Denial of services by unauthorised traffic generation**
   - **a. a storm by a chain reaction**
   - **b. victimising a host by unsolicited messages**
4. **Unauthorised disclosure of information**

In the network level attack 1, **unauthorised tampering**, network packets can be redirected to a bogus host, giving it control over the network packets, by misuse of a routing information exchange protocol including an unsolicited management message like an ICMP Redirect message. A bogus host might be configured with an unassigned address. It initiates a routing information exchange protocol, and other routers start forwarding packets to this bogus host by mistaking it as a new router. Alternatively two bogus routers can be used in such a way that one is configured with the address of the existing local router, and the other with an unassigned address. These bogus routers never appear in the local subnet operations, so that it will not be detected by the authentic one. The first bogus router sends an unsolicited network level management message, an ICMP Redirect message, to the victim hosts indicating that further messages should be sent to the second bogus router. The other hosts correct their routing tables by replacing the authentic router with the second bogus router. In both instances, the deceiver gains control over traffic which should have gone through the authentic router. It will drop some of the packets, and forward the rest to the authentic one.

Another way to redirect packets is possible if there exists inconsistency, e.g. a router is removed without notifying the hosts which use the router. A bogus

host pretends to be the authentic router by answering ARP Requests from victim hosts; the victims then forward packets to the deceiver. Alternatively its source address is changed to the deceiver's address, and the packet is resent to another possible router if any; in this way the deceiver can intercept the transaction between the victim and the destination hosts.

For **unauthorised use of the network resource**, 2, a deceiver may steal network bandwidth without being noticed, by using an unassigned local address. A deceiver may gain unauthorised access to network services, by communicating with another bogus host.

Unauthorised traffic generation has two aspects; one is to provoke a chain reaction in a network; the other to produce unsolicited packets. A 3-a attack uses a chain reaction. A deceiver configures a bogus host X with a network address which is not in use, and with a subnet level broadcast address. The deceiver produces an IP packet destined for itself, and broadcasts it out to the subnet. The packet is regenerated by the routers and misconfigured hosts with the forwarding capability described in the previous section. The deceiver could make use of a subnet level address other than the broadcast address; however, in this case, he needs to have more than three hosts configured with a forwarding capability, all of which have the same subnet level address.

For a 3-b traffic generation attack, a traffic generator can be configured by a deceiver with any address, foreign or local, and then it can be used for attacking a site, or a group of sites. A bogus host will never be detected and go on attacking from time to time by changing its own address. The generator can produce hundreds of request packets to victim hosts. The victim hosts will be unable to process the quantity of packets, and consequently the users will be denied network services due to overloading. This could lead to denial of distributed services such as the Network File Store, due to congestion.

Unauthorised disclosure of information, 4, is learning of some of the registered addresses by unauthorised observation of traffic.

3.4 Evaluation of the consequences

Once a network storm has started, we can cease it only by stopping all the misconfigured or even healthy routers on the subnet. The misconfiguration of a subnet address is also serious when it produces the extra traffic but not enough to cause a network storm. The legal packets would be doubled without being detected, and could go beyond an organisation. One would argue that each gateway checks the packet ID, however, those gateways are basically state-less, i.e. they have no record on which packets they forwarded.

Unauthorised tampering by unauthorised redirection at the network level is serious because it causes denial of services at the application level, no matter how strong the protection mechanisms are provided at the application level. Unauthorised redirection should be prevented by having counter measures either at the network level or the subnetwork level. Network level protection, however, does not protect the network from a deceit which makes use of inconsistency. The subnetwork level protection counters resource stealing as well. It also provides protection from unauthorised traffic generation caused by a chain reaction; however, it could be prevented more easily by having a filter of incoming ARP packets at each router.

Unauthorised use of the network resource is a serious problem in the inter-organisational network environment in particular. This would be countered by access control such as policy routing [6] [7]. However, nothing can be done for the resource stealing at the subnet level by policy routing, and an appropriate control would be required for unauthorised use of a subnet.

Unauthorised traffic generation by unsolicited messages is serious; it may cause denial of services at the application level, and it cannot be prevented. Detection of the location of the attacker is possible but only to the extent determining the subnet on which the deceiver is operating. It is difficult to identify the host being used by the deceiver, as the generated packets may not include the real address of the host.

Authentication and/or access control at the application level would solve many problems which are ob-

stacles for configuration detection; there would then be little point in pursuing a deceit at the application level as a deceiver can benefit little.

Traffic generation by a chain reaction setting the subnet broadcast address could solved by filtering the ARP packets, so that no ARP packet with the subnet broadcast address would not be accepted by routers. Traffic generation by unsolicited messages is difficult to protect against, but it may be possible to determine the subnet on which deceiver is operating.

Information disclosure at the network level can be a serious problem because a network configuration can be learned by anyone.

Those serious consequences could be countered individually according to the environmental needs. For instance, in the intra-organisation network environment, unauthorised redirection, unauthorised use of resources, and information disclosure would be more serious than the others. Whereas in the intra-organisational networks, misconfiguration problems and unauthorised redirection would be serious.

The fundamental problem, however, is that the current networks operate with information such as addresses which have not been verified.

4 Conclusion

The primary design goal of internet protocols has been emphasised to be connectivity and reachability [4], [17] over the various types of subnetworks rather than access control. The Internet Protocol is a successful example in this respect. The consequence is that the internet protocol environment is insecure. In a traditional network information systems such as name server and routing information exchangers, an information item has been claimed without verification. There are many pitfalls where anyone can break in [2]. There are many chances that network operators can make hazard mistakes — e.g. a typing mistake in setting up a network address. Incidents have been almost avoided, partly because of the implied mutual trust in the research internetwork community, and partly because of the size of the internetwork has been manageable for an operator to input information items correctly.

As the size of an internetwork grows and the diversity of nodal system architecture becomes significant throughout both inside and between single organisations, some control is required to protect the network from both malicious and innocent incidents. The original Core Gateways of the ARPANET and MILNET were operated by BBN, and access control prevented manipulation of routing tables. As the Internet becomes more diverse, more general techniques are needed. Policy routing is based on the idea that the control would be exercised at a border gateway on an inter-organisational network environment.

In this paper, we have described some serious problems from using incorrect information in networking from the viewpoint of configuration detection. The very reason for those network incidents is that a network operation is reliant on some information such as addresses. The fundamental problem is that there is no way to know how trustworthy the information in use is.

Acknowledgements

The topic presented in this paper was researched during the author's stay at University College London. I would like to thank Peter T. Kirstein for his support. I would also like to thank many people at University College London, Peter Williams, Gordon Joly, John Andrews, Ping Hu, and the others.

The discussion with Radia Perlman was useful for investigation of the consequences from misconfiguration.

References

[1] V. G. Cerf andP. T. Kirstein. Issues in packet-network interconnection. *Proceedings of the IEEE*, Vol. 66, No. 11, pp. 1386–1408, November 1978.

[2] S. M. Bellovin. Security problems in the tcp/ip protocol suite. *ACM Computer Communication Review*, Vol. 19, No. 2, April 1989.

[3] CCITT and ISO. Recommendations x.500 series; the directory - x.500 (iso 9594-1) overview of concepts, models and services; x.501 (iso 9594-2) models; x.509 (iso 9594-8) authentication framework; x.511 (iso 9594-3) abstract service definition; x.518 (iso 9594-4) procedures for distributed operation; x.519 (iso 9594-5) protocol specifications; x.520 (iso 9594-6) selected attribute types; x.521 (iso 9594-7) selected object classes. International Standard X.500, March 1988.

[4] V. G. Cerf and E. Cain. The dod internet architecture model. *Computer Networks*, pp. pp.307–318, July 1983. North-Holland.

[5] D. R. Cheriton and T. P. Mann. Decentralizing a global naming service for improved performance and fault tolerance. *ACM Transactions on Computer Systems*, Vol. 7, No. 2, pp. 147–183, May 1989.

[6] D. Clark. Policy routing in internet protocols. RFC 1102, May 1989.

[7] D. Estrin. Policy requirements for inter-administrative domain routing. *Computer Networks and ISDN Systems*, No. 22, pp. 179–191, 1991.

[8] R. Hinden, J. Haverty, and A. Sheltzer. The darpa internet: Interconnecting heterogeneous computer networks with gateways. *COMPUTER (IEEE)*, Vol. 16, No. 9, pp. pp.38–48, September 1983.

[9] IEEE. Draft ieee standard 802.1: Part d, mac bridges. International Standard Revision C, IEEE Project 802 Local and Metropolitan Area Network Standards, August 1987.

[10] ISO. Iso 7498 information processing systems - open systems interconnection - basic reference model. International Standard ISO 7498, 1984.

[11] U. Manber. Chain reactions in networks. *IEEE COMPUTER*, Vol. 23, No. 10, pp. 57–63, October 1990.

[12] P. Mockapetris. Domain names - concepts and facilities. RFC 1034, November 1987.

[13] Y. Murayama. An introduction to the status alarm system. Internal Note 1667, Dept. of Computer Scienece, University Colege London, January 1985.

[14] Y. Murayama. *Configuration Detection and Verification in Computer Networks.* PhD thesis, University of London, December 1991.

[15] R. Perlman. Personal message, the discussion about the doubled traffic caused by the misconfiguration of a host, August 1991.

[16] D. Plummer. An ethernet address resolution protocol. RFC 826, November 1982.

[17] L. Pouzin. Internetworking. In W. Chou, editor, *Computer Communications*, Vol. Volume 2: systems and applications, chapter 15. Prentice-Hall, 1985.

[18] E. C. Rosen. Exterior gateway protocol (egp). RFC 827, October 1982.

[19] C. Shannon. *The Mathematical Theory Of Communication.* The University Of Illinois Press: Urbana, 1949.

[20] J. Shoch. Inter-network naming, addressing, and routing. *Conf. Proc. of IEEE COMPCON Fall 1978*, pp. 72–79, 1978.