

PIM WG  
Internet-Draft  
Intended status: Standards Track  
Obsoletes: [RFC1112](#)  
Expires: January 13, 2022

S. Deering  
Retired  
T. Eckert (Ed.)  
Futurewei USA  
July 12, 2021

Host Extensions IP Multicasting - Any Source Multicast (ASM)  
draft-eckert-pim-rfc1112bis-00

## ABSTRACT

This memo specifies the extensions required of a host implementation of the Internet Protocol (IP) to support Any Source Multicast (ASM) IP Multicasting or abbreviated IP Multicast.  
Distribution of this memo is unlimited.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."  
This Internet-Draft will expire on January 13, 2022.

## Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## [2](#). INTRODUCTION

The host extensions defined in this memo are called Any Source Multicast (ASM) IP multicast or abbreviated IP multicast. The term Any Source Multicast is used to distinguish these extensions from Source Specific Multicast (SSM) IP multicast as defined by [\[RFC4607\]](#). The abbreviation IP multicast always refers to this memo's extensions.

This memo applies to both IPv4 and IPv6. When it uses the term IP it implies either or both version of the IP protocol. It uses the terms IPv4 and/or IPv6 explicitly when referring to functions applicable to only a specific version of the IP protocol.

This document is a revision of [\[RFC1112\]](#). See [Appendix II](#). for a detailed list of changes from that memo.

IP multicasting is the transmission of an IP datagram to a "host group", a set of zero or more hosts identified by a single IP destination address. A multicast datagram is delivered to all members of its destination host group with the same "best-efforts" reliability as regular unicast IP datagrams, i.e., the datagram is not guaranteed to arrive intact at all members of the destination group or in the same order relative to other datagrams.

The membership of a host group is dynamic; that is, hosts may join and leave groups at any time. There is no restriction on the location or number of members in a host group. A host may be a member of more than one group at a time. A host need not be a member of a group to send datagrams to it.

A host group may be permanent or transient. A permanent group has a well-known, administratively assigned IP address. It is the address, not the membership of the group, that is permanent; at any time a permanent group may have any number of members, even zero. Those IP multicast addresses that are not reserved for permanent groups are available for dynamic assignment to transient groups which exist only as long as they have members.

Internetwork forwarding of IP multicast datagrams is handled by "multicast routers" which may be co-resident with, or separate from, internet gateways. A host transmits an IP multicast datagram as a local network multicast which reaches all immediately-neighborhood members of the destination host group. If the datagram has an IP time-to-live greater than 1, the multicast router(s) attached to the local network take responsibility for forwarding it towards all other networks that have members of the destination group. On those other member networks that are reachable within the IP time-to-live, an attached multicast router completes delivery by transmitting the datagram as a local multicast.

This memo specifies the extensions required of a host IP implementation to support IP multicasting, where a "host" is any

internet host or gateway other than those acting as multicast routers. The algorithms and protocols used within and between multicast routers are transparent to hosts and will be specified in separate documents. This memo also does not specify how local network multicasting is accomplished for all types of network, although it does specify the required service interface to an arbitrary local network and gives an Ethernet specification as an example. Specifications for other types of network will be the subject of future memos.

### 3. LEVELS OF CONFORMANCE

There are three levels of conformance to this specification:

Level 0: no support for IP multicasting.

There is, at this time, no requirement that all IPv4 implementations support IP multicasting. Level 0 hosts will, in general, be unaffected by multicast activity. The only exception arises on some types of local network, where the presence of level 1 or 2 hosts may cause misdelivery of multicast IP datagrams to level 0 hosts. Such datagrams can easily be identified by the presence of a class D IP address in their destination address field; they should be quietly discarded by hosts that do not support IP multicasting. Class D addresses are described in [section 4](#) of this memo.

Level 1: support for sending but not receiving multicast IP datagrams.

Level 1 allows a host to partake of some multicast-based services, such as resource location or status reporting, but it does not allow a host to join any host groups. An IP implementation may be upgraded from level 0 to level 1 very easily and with little new code. Only sections [4](#), [5](#), and [6](#) of this memo are applicable to level 1 implementations.

Level 2: full support for IP multicasting.

Level 2 allows a host to join and leave host groups, as well as send IP datagrams to host groups. Most IPv6 hosts require Level 2 support because IPv6 Neighbor Discovery ([\[RFC4861\]](#), as used on most link types) depends on multicast and requires that nodes join Solicited Node multicast a

Level 2 requires implementation of the Internet Group Management Protocol (IGMP) for IPv4 and the equivalent Multicast Listener Discovery Protocol (MLDv) for IPv6, and extension of the IP and local network service interfaces within the host

The current protocol versions are IGMPv3 [\[RFC3376\]](#) and MLDv2 [\[RFC3810\]](#) or later versions of either protocol [\[RFC5790\]](#).

All of the following sections of this memo are applicable to level 2

implementations.

#### 4. HOST GROUP ADDRESSES

IPv4 Host groups are identified by class D IP addresses, i.e., those with "1110" as their high-order four bits. Class E IP addresses, i.e., those with "1111" as their high-order four bits, are reserved for future addressing modes.

In Internet standard "dotted decimal" notation, host group addresses range from 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is guaranteed not to be assigned to any group, and 224.0.0.1 is assigned to the permanent group of all IPv4 hosts (including gateways). This is used to address all IPv4 multicast hosts on the directly connected network. There is no multicast address (or any other IP address) for all hosts on the total Internet. The addresses of other well-known, permanent groups are to be published in "Assigned Numbers".

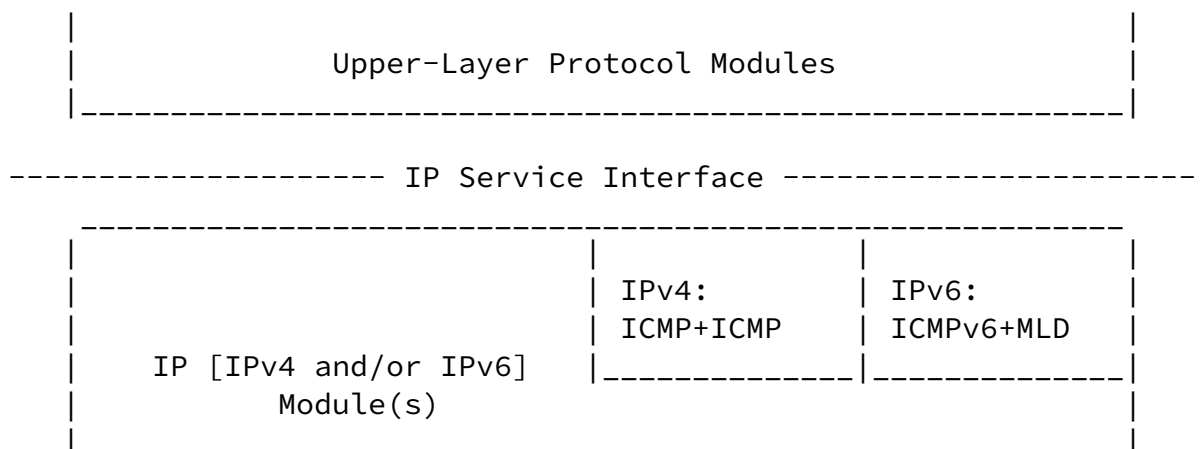
IPv6 Host groups are identified by IPv6 addresses as defined in [\[RFC4291\]](#) and updated by [\[RFC7346\]](#), [\[RFC7371\]](#).

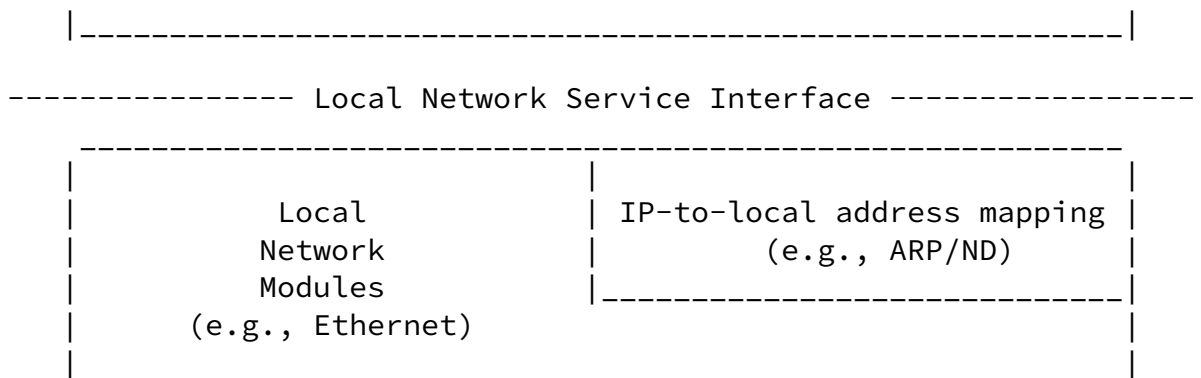
IPv4 and IPv6 addresses as specified in [\[RFC4607\]](#) are not used for ASM IP multicast and are not considered IP host groups. They are instead only the destination address part G of Source Specific Multicast (SSM) IP multicast (S,G) channels.

[Appendix I](#) contains some background discussion of several issues related to host group addresses.

#### 5. MODEL OF A HOST IP IMPLEMENTATION

The multicast extensions to a host IP implementation are specified in terms of the layered model illustrated below. In this model, ICMP/ICMPv6 and (for level 2 hosts) IGMP/MLD are considered to be implemented within the IP module, and the mapping of IP addresses to local network addresses is considered to be the responsibility of local network modules. This model is for expository purposes only, and should not be construed as constraining an actual implementation.





To provide level 1 multicasting, a host IP implementation must support the transmission of multicast IP datagrams. To provide level 2 multicasting, a host must also support the reception of multicast IP datagrams. Each of these two new services is described in a separate section, below. For each service, extensions are specified for the IP service interface, the IP module, the local network service interface, and an Ethernet local network module. Extensions to local network modules other than Ethernet are mentioned briefly, but are not specified in detail.

## [6. SENDING MULTICAST IP DATAGRAMS](#)

### [6.1. Extensions to the IP Service Interface](#)

Multicast IP datagrams are sent using the same "Send IP" operation used to send unicast IP datagrams; an upper-layer protocol module merely specifies an IP host group address, rather than an individual IP address, as the destination. However, a number of extensions may be necessary or desirable.

First, the service interface should provide a way for the upper-layer protocol to specify the IP time-to-live of an outgoing multicast datagram, if such a capability does not already exist. If the upper-layer protocol chooses not to specify a time-to-live, it should default to 1 for all multicast IP datagrams, so that an explicit choice is required to multicast beyond a single network.

Second, for hosts that may be attached to more than one network, the service interface should provide a way for the upper-layer protocol to identify which network interface is to be used for the multicast transmission. Only one interface is used for the initial transmission; multicast routers are responsible for forwarding to any other networks, if necessary. If the upper-layer protocol chooses not to identify an outgoing interface, a default interface should be used, preferably under the control of system management.

Third (level 2 implementations only), for the case in which the host is itself a member of a group to which a datagram is being sent, the service interface should provide a way for the upper-layer protocol to inhibit local delivery of the datagram; by default, a copy of the

datagram is looped back. This is a performance optimization for upper-layer protocols that restrict the membership of a group to one process per host (such as a routing protocol), or that handle loopback of group communication at a higher layer (such as a multicast transport protocol).

IPv6 socket extensions supporting these functions are defined in [\[RFC3493\]](#),

## [6.2.](#) Extensions to the IP Module

To support the sending of multicast IP datagrams, the IP module must be extended to recognize IP host group addresses when routing outgoing datagrams. Most IP implementations include the following logic:

```
if IP-destination is on the same local network,
    send datagram locally to IP-destination
else
    send datagram locally to GatewayTo( IP-destination )
```

To allow multicast transmissions, the routing logic must be changed to:

```
if IP-destination is on the same local network
or IP-destination is a host group,
    send datagram locally to IP-destination
else
    send datagram locally to GatewayTo( IP-destination )
```

If the sending host is itself a member of the destination group on the outgoing interface, a copy of the outgoing datagram must be looped-back for local delivery, unless inhibited by the sender. (Level 2 implementations only.)

The IP source address of the outgoing datagram must be one of the individual addresses corresponding to the outgoing interface.

A host group address must never be placed in the source address field or anywhere in a source route or record route option of an outgoing IP datagram.

## [6.3.](#) Extensions to the Local Network Service Interface

No change to the local network service interface is required to support the sending of multicast IP datagrams. The IP module merely specifies an IP host group destination, rather than an individual IP destination, when it invokes the existing "Send Local" operation.

## [6.4.](#) Extensions to an Ethernet Local Network Module

The Ethernet directly supports the sending of local multicast packets by allowing multicast addresses in the destination field of Ethernet packets. All that is needed to support the sending of multicast IP datagrams is a procedure for mapping IP host group addresses to Ethernet multicast addresses.

An IPv4 host group address is mapped to an Ethernet multicast address by placing the low-order 23-bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01-00-5E-00-00-00 (hex). Because there are 28 significant bits in an IP host group address, more than one host group address may map to the same Ethernet multicast address.

Mapping of IPv6 host group addresses to Ethernet is defined in [\[RFC2464\]](#) and [\[RFC6085\]](#).

## [6.5](#). Extensions to Local Network Modules other than Ethernet

Other networks that directly support multicasting, such as rings or buses conforming to the IEEE 802.2 standard, may be handled the same way as Ethernet for the purpose of sending multicast IP datagrams. For a network that supports broadcast but not multicast, such as the Experimental Ethernet, all IP host group addresses may be mapped to a single local broadcast address (at the cost of increased overhead on all local hosts). For a point-to-point link joining two hosts (or a host and a multicast router), multicasts should be transmitted exactly like unicasts. For a store-and-forward network like the ARPANET or a public X.25 network, all IP host group addresses might be mapped to the well-known local address of an IP multicast router; a router on such a network would take responsibility for completing multicast delivery within the network as well as among networks.

## [7](#). RECEIVING MULTICAST IP DATAGRAMS

### [7.1](#). Extensions to the IP Service Interface

Incoming multicast IP datagrams are received by upper-layer protocol modules using the same "Receive IP" operation as normal, unicast datagrams. Selection of a destination upper-layer protocol is based on the protocol field in the IP header, regardless of the destination IP address. However, before any datagrams destined to a particular group can be received, an upper-layer protocol must ask the IP module to join that group. Thus, the IP service interface must be extended to provide two new operations:

JoinHostGroup ( group-address, interface )

LeaveHostGroup ( group-address, interface )

The JoinHostGroup operation requests that this host become a member

of the host group identified by "group-address" on the given network interface. The LeaveGroup operation requests that this host give up its membership in the host group identified by "group-address" on the given network interface. The interface argument may be omitted on hosts that support only one interface. For hosts that may be attached to more than one network, the upper-layer protocol may choose to leave the interface unspecified, in which case the request will apply to the default interface for sending multicast datagrams (see [section 6.1](#)).

It is permissible to join the same group on more than one interface, in which case duplicate multicast datagrams may be received. It is also permissible for more than one upper-layer protocol to request membership in the same group.

Both operations should return immediately (i.e., they are non-blocking operations), indicating success or failure. Either operation may fail due to an invalid group address or interface identifier. JoinHostGroup may fail due to lack of local resources. LeaveHostGroup may fail because the host does not belong to the given group on the given interface. LeaveHostGroup may succeed, but the membership persist, if more than one upper-layer protocol has requested membership in the same group.

IPv6 socket extensions supporting these functions are defined in [\[RFC3493\]](#), [\[RFC3678\]](#) specifies these functions for IPv4 and IPv6 (as well as for SSM).

## [7.2](#). Extensions to the IP Module

To support the reception of multicast IP datagrams, the IP module must be extended to maintain a list of host group memberships associated with each network interface. An incoming datagram destined to one of those groups is processed exactly the same way as datagrams destined to one of the host's individual addresses.

Incoming datagrams destined to groups to which the host does not belong are discarded without generating any error report or log entry. On hosts with more than one network interface, if a datagram arrives via one interface, destined for a group to which the host belongs only on a different interface, the datagram is quietly discarded. (These cases should occur only as a result of inadequate multicast address filtering in a local network module.)

An incoming datagram is not rejected for having an IP time-to-live of 1 (i.e., the time-to-live should not automatically be decremented on arriving datagrams that are not being forwarded). An incoming datagram with an IP host group address in its source address field is quietly discarded. An ICMP/ICMPv6 error message (Destination Unreachable, Time Exceeded, Parameter Problem, Source Quench, or Redirect) is never generated in response to a datagram destined to an IP host group.



The list of host group memberships is updated in response to JoinHostGroup and LeaveHostGroup requests from upper-layer protocols. Each membership should have an associated reference count or similar mechanism to handle multiple requests to join and leave the same group. On the first request to join and the last request to leave a group on a given interface, the local network module for that interface is notified, so that it may update its multicast reception filter (see [section 7.3](#)).

The IP module must also be extended to implement the IGMP protocol for IPv4 and the MLD protocol for IPv6. IGMP/MLD are used to keep neighboring routers informed of the host group memberships present on a particular local network.

### [7.3](#). Extensions to the Local Network Service Interface

Incoming local network multicast packets are delivered to the IP module using the same "Receive Local" operation as local network unicast packets. To allow the IP module to tell the local network module which multicast packets to accept, the local network service interface is extended to provide two new operations:

JoinLocalGroup ( group-address )

LeaveLocalGroup ( group-address )

where "group-address" is an IP host group address. The JoinLocalGroup operation requests the local network module to accept and deliver up subsequently arriving packets destined to the given IP host group address. The LeaveLocalGroup operation requests the local network module to stop delivering up packets destined to the given IP host group address. The local network module is expected to map the IP host group addresses to local network addresses as required to update its multicast reception filter. Any local network module is free to ignore LeaveLocalGroup requests, and may deliver up packets destined to more addresses than just those specified in JoinLocalGroup requests, if it is unable to filter incoming packets adequately.

The local network module must not deliver up any multicast packets that were transmitted from that module; loopback of multicasts is handled at the IP layer or higher.

### [7.4](#). Extensions to an Ethernet Local Network Module

To support the reception of multicast IP datagrams, an Ethernet module must be able to receive packets addressed to the Ethernet multicast addresses that correspond to the host's IP host group addresses. It is highly desirable to take advantage of any address filtering capabilities that the Ethernet hardware interface may have,

so that the host receives only those packets that are destined to it.

Unfortunately, many current Ethernet interfaces have a small limit on the number of addresses that the hardware can be configured to recognize. Nevertheless, an implementation must be capable of listening on an arbitrary number of Ethernet multicast addresses, which may mean "opening up" the address filter to accept all multicast packets during those periods when the number of addresses exceeds the limit of the filter.

For interfaces with inadequate hardware address filtering, it may be desirable (for performance reasons) to perform Ethernet address filtering within the software of the Ethernet module. This is not mandatory, however, because the IP module performs its own filtering based on IP destination addresses.

#### 7.5. Extensions to Local Network Modules other than Ethernet

Other multicast networks, such as IEEE 802.2 networks, can be handled the same way as Ethernet for the purpose of receiving multicast IP datagrams. For pure broadcast networks, such as the Experimental Ethernet, all incoming broadcast packets can be accepted and passed to the IP module for IP-level filtering. On point-to-point or store-and-forward networks, multicast IP datagrams will arrive as local network unicasts, so no change to the local network module should be necessary.

### APPENDIX I. HOST GROUP ADDRESS ISSUES

This appendix is not part of the IP multicasting specification, but provides background discussion of several issues related to IP host group addresses.

#### Group Address Binding

The binding of IP host group addresses to physical hosts may be considered a generalization of the binding of IP unicast addresses. An IP unicast address is statically bound to a single local network interface on a single IP network. An IP host group address is dynamically bound to a set of local network interfaces on a set of IP networks.

It is important to understand that an IP host group address is NOT bound to a set of IP unicast addresses. The multicast routers do not need to maintain a list of individual members of each host group. For example, a multicast router attached to an Ethernet need associate only a single Ethernet multicast address with each host group having local members, rather than a list of the members' individual IP or Ethernet addresses.

#### Allocation of Transient Host Group Addresses

This memo does not specify how transient group address are allocated. It is anticipated that different portions of the IP transient host group address space will be allocated using different techniques. For example, there may be a number of servers that can be contacted to acquire a new transient group address. Some higher-level protocols (such as VMTP, specified in [RFC-1045](#)) may generate higher-level transient "process group" or "entity group" addresses which are then algorithmically mapped to a subset of the IP transient host group addresses, similarly to the way that IP host group addresses are mapped to Ethernet multicast addresses. A portion of the IP group address space may be set aside for random allocation by applications that can tolerate occasional collisions with other multicast users, perhaps generating new addresses until a suitably "quiet" one is found.

In general, a host cannot assume that datagrams sent to any host group address will reach only the intended hosts, or that datagrams received as a member of a transient host group are intended for the recipient. Misdelivery must be detected at a level above IP, using higher-level identifiers or authentication tokens. Information transmitted to a host group address should be encrypted or governed by administrative routing controls if the sender is concerned about unwanted listeners.

## APPENDIX II. Changes from [RFC1112](#)

This document updates [RFC1112](#) with the following changes:

- o It removes the claim that these host extensions are "... the recommended standard".
- o It is written to apply to both IPv4 and IPv6 by adding equivalent detail for IPv6.
- o It introduces the term "ASM IP multicast" as another term for "Host Extension".
- o It removes the original [appendix I of RFC1112](#) that defined the IGMP version 1.

## APPENDIX III. Discussion and Explanations

[RFC-editor: Please remove this section]

### Intention:

This document is intended to be an update to [RFC1112](#) for the following reasons:

[RFC1112](#) is at the time of this writing the only FULL INTERNET STANDARD describing

[RFC1112](#) includes the specification of IGMPv1. PIM WG would like to make IGMPv1

### Open Issues:

Has any document after [RFC1112](#) (re-)defined the mapping of IPv4 multicast group addresses to ethernet multicast MAC ? If so, then we should include a reference to it and update the appropriate text in this rfc1112bis.

This document uses/defines the term "host group", which really is only a term relevant in this traditional ASM service model, but not in SSM. Therefore new text stating that IP multicast group addresses from [RFC4607](#) are not included in this ASM definition. This is hopefully aligned with the text in [RFC4607](#).

Is it appropriate to have included text to refer to socket API for ASM eg: [rfc3678](#). These socket APIs are primarily about UDP sockets, and only rarely for IP level. This document only specifies an IP Service Interface. Note that [RFC4607](#) (SSM) also refers to the socket interface extensions for SSM under a section called "Extensions to the IP Module Interface".

Discuss: Why would this document still need to be a standards track document ? IETF typically does not assign standard track to pure API/service-interface document. With IGMPv1 removed from the document, what does it still make it standards track, aka: define in a normative fashion interoperability impacting behavior of nodes ?

Answer1: Definition of the IPv4 ASM address space, Definition of the IP Multicast group to ethernet MAC address mapping for IPv4. The document now also contains references to these standards aspects for IPv6, but those are references to prior standards track documents.

Answer2: This document is similar in scope to [RFC4607](#) (SSM) which is standards track. The newly defined interop impacting behavior on the wire are also limited: address ranges for IPv4/IPv6 SSM, and it is referring to [RFC1112](#).

Discuss: Which documents should we claim this document is updating ? Hopefully none other than [rfc1112](#) - [rfc1112](#) itself is referenced by > 60 RFCs.  
si

#### Author's Addresses

Stephen E. Deering  
Retired  
Vancouver, British Columbia  
Canada  
Email: bob.hinden@gmail.com (email secretary)

Toerless Eckert (editor)  
Futurewei Technologies Inc. USA  
2220 Central Expy  
Santa Clara, CA 95050  
United States of America  
Email: tte@cs.fau.de