NS-2 Modeling for Project 25: Inter-RF Subsystem Interface (ISSI)

September 30, 2011
Simulation Model v1.0.0

1 Overview 3
1.1 Design Assumptions 3
1.2 Software Model 4
2 Mobility Management 5
2.1 RFSS 5
2.2 Subscriber Unit (SU) 6
2.3 Subscriber Group (SG) 7
3 Call Control Management 8
3.1 Unit-to-Unit Call 8
3.2 Group Call 10
3.2.1 Handling of Simultaneous Requests 11
3.3 Resource Management 12
3.3.1 RTP Resources 12
3.3.2 RF Resources 13
3.3.3 Priority and Pre-emption 13
4 Media Transmission 13
4.1 PTT Control 14
4.2 Application 15
4.2.1 Statistical Traffic Generation 16
4.2.2 Trace-based Traffic Generation 19
5 Protocol Modeling 19
5.1 ISSI 19
5.2 SIP 19
5.3 RTP 19
6 Known Issues and Limitations 19
7 References 19

1 Overview

The Project 25 (P25) defines several interfaces enabling communication for the Public Safety community. The software described in this document provides a simulation model for the Inter-RF Subsystem Interface (ISSI). This interface is designed to enable multiple RF Subsystems (RFSSs) from different manufacturers to communicate via a standard protocol.

1.1 Design Assumptions

As shown in Figure 1, an RFSS provides services to the mobile users within its coverage area. Those services include location tracking, call setup, and RF allocations for both Subscriber Units (SUs) and Subscriber Groups (SGs). An RFSS is composed of several logical components but the actual implementation is manufacturer dependent. However, two RFSSs communicate via the G interface that connects their RF Subsystem Gateways (RFGs). In the current model, all internal components and the behavior of the RFSS are abstracted and modeled within the RFG. Similarly, the Common Air Interface (CAI) between the Base Radio (BR) and the mobile radios (MR) is not modeled, and the RF resources are abstracted. Therefore, throughout this document, the terms RFSS and RFG are interchangeable.

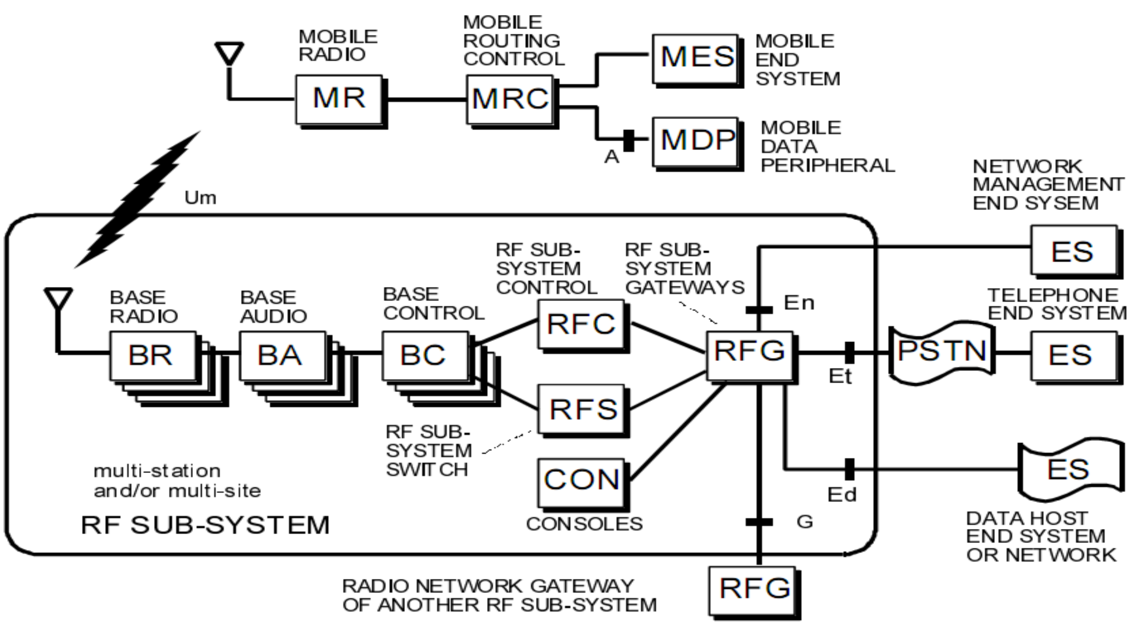

Figure 1: RFSS Architecture

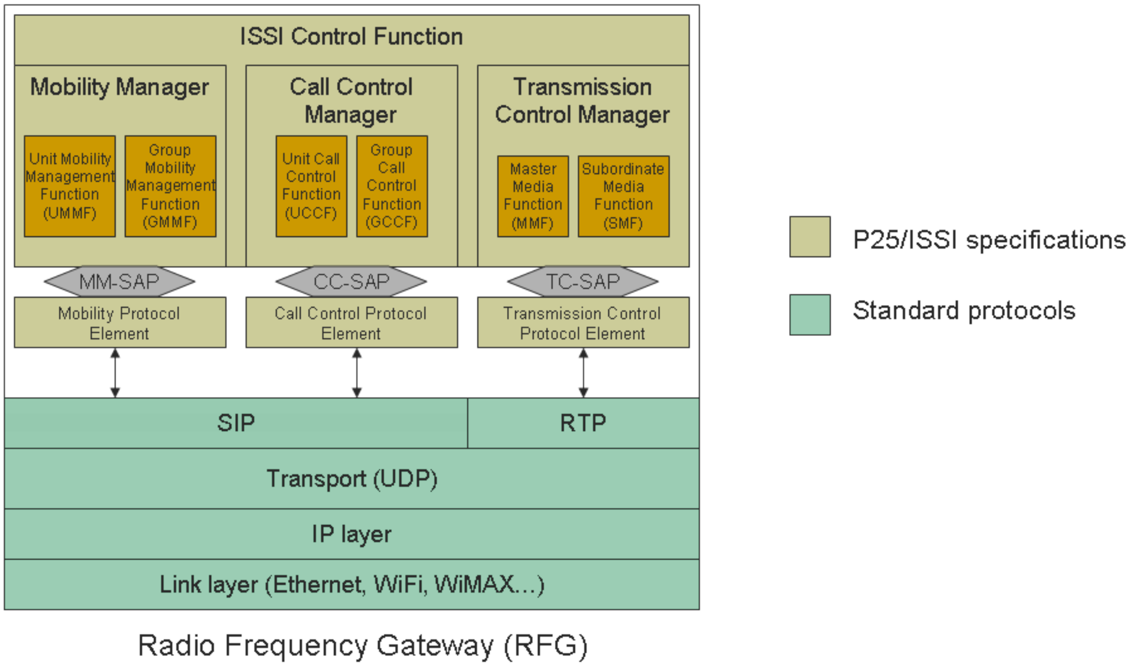

Figure 2: RFG Design

The architecture of the RFG is shown in Figure 2. It provides the services of the RFSS, namely the mobility management (Section 2), call control management (Section 3), and transmission control management (Section 4). The ISSI protocol is built at the Application layer in the TCP/IP layer model; it uses the Session Initiation Protocol (SIP) for signaling and Real-time Transport Protocol (RTP) for carrying data, both described in Section 5.

1.2 Software Model

The ISSI modeling tool developed is based on the ns-2 simulator. This tool, widely used in the research community, is a discreet event simulator for network environment, providing components from the lower layers (e.g., Ethernet, Wi-Fi, WiMAX) to the transport layers (e.g., UDP, TCP, SCTP). Due to its open source, ns-2 can be easily modified to add new features and support additional protocols such as ISSI. Figure 3 shows the internal modules of an RFG as implemented in ns-2 and identifies the components that were reused, as well as the ones created specifically to support the ISSI. The rest of the document describes each module focusing on the implementation-specific mechanisms that can affect the behavior and/or performance of the ISSI specifications.

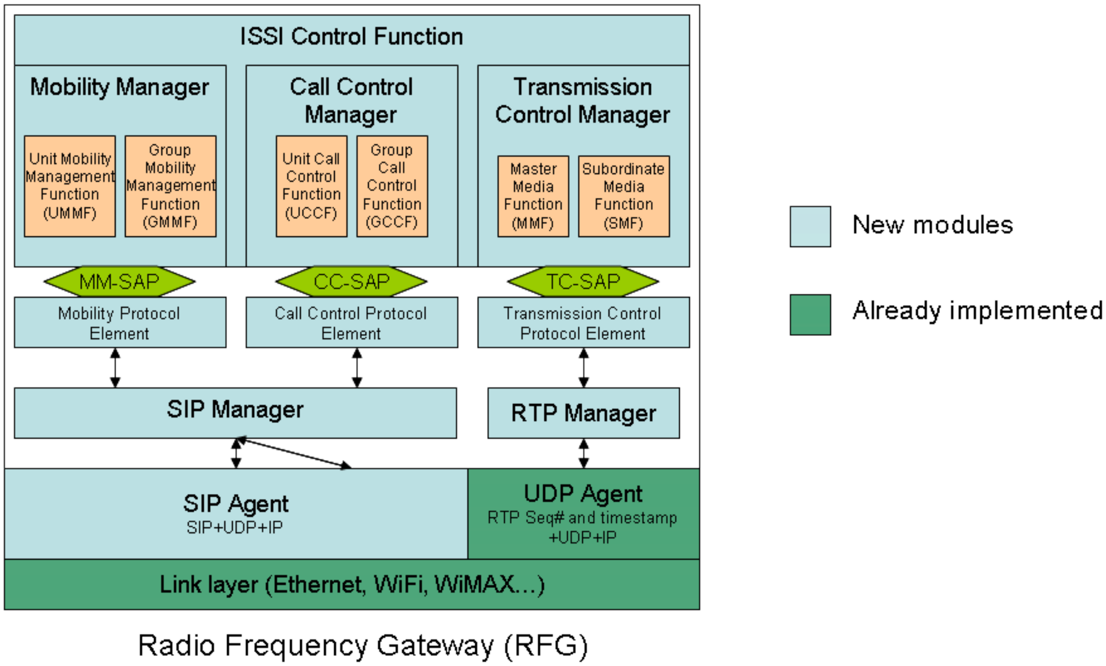

Figure 3: RFG Simulation Model

2 Mobility Management

The Mobility Management is responsible for tracking the location and defining the access permissions of the SUs and SGs throughout the network.

2.1 RFSS

An RFSS can play one of two roles: Home or Serving. A Home RFSS is the entity that tracks the movement of a particular SU or SG. In the current model, the Home RFSS is configured at the beginning of the simulation and cannot be changed. A Serving RFSS represents the current location of the SU or SG. We can note that an SU can only be served by one RFSS at a given time while an SG can be registered at multiple RFSSs if multiple SUs that are part of the group are located on different RFSSs.

From the model point of view, the roles are assigned by the following actions:

- *Home*: by adding a new SU/SG. During this process, the SU/SG is created and assigned its identifier based on the RFSS's configuration (WACN ID and SYSTEM ID).
- *Serving*: by calling the register function on the RFSS. By doing so, the user requests that an SU "moves" to the RFSS, thus updating its location. In turn, the RFSS will try to perform the registration with the SU's Home RFSS and may perform registration with the SG's Home RFSS for which the SU is a member (if not already done). It is possible to register the SU at its Home RFSS, in which case the roles are collocated. A Serving RFSS can lose its role when the user deregisters an SU or when the SU is "moved" to another RFSS (i.e. roaming).

When a new Serving RFSS tries to register an SU, the Home RFSS indicates the previous serving, if it exists, that the SU has left its coverage area by sending a Roamed indication message.

The value of the registration lifetime is defined per RFSS and is identical for both SUs and Groups. The default value is set to 3600 s. The serving RFSSs maintain the registration of the SUs and Groups by sending periodic registration messages prior to the expiration of the registration. Upon receiving a successful registration response, the serving RFSS computes the time for sending the next registration message using the following formula:

*Time before sending next registration = 0.9 * lifetime – registration delay*

For example, assuming that the lifetime value is 50 s, if the serving RFSS sends an initial registration at 10 s and receives a response at 12 s, the next registration message will be sent in 0.9*50-(12-10) = 43 s, which is equivalent to simulation time 12+43=55 s

2.2 Subscriber Unit (SU)

An SU represents a single mobile terminal that can roam among the RFSSs and place calls to other SUs/SGs. Since the CAI is not modeled, the SU is a logical entity stored in the RFSS and there are no network messages exchanged between the SU and the RFSS (only call functions). An instance of the SU is created at the Home RFSS and is used for configuration and tracking purposes. Other instances are created and destroyed as needed at the Serving RFSSs. Once created, the SU's service profile can be configured. Currently, the Home RFSS only stores one profile, therefore the same information is provided during registration regardless of the Serving RFSS. Table 1 lists the various attributes for the service profile and the current support level. Reasons for nor supporting some attributes include software tool limitations and our release timeline.

Table 1 : User Service Profile Attributes Support

Name	Support Level
System Access Permission	Supported. This attribute is used in call admission control.
Duplexity	Not supported. The duplexity seems to impact the communication between the BS and the SU, not the communication between the RFSSs.
Secure Capable	Not supported.
Group Call Capability	Supported. This attribute is used in call admission control.
U2U Call Permission	Supported. This attribute is used in call admission control.
U2U Call Priority	Supported. This attribute is used in call admission control.
Interconnect Permission	Not supported. Simulation tool does not provide PSTN.
Interconnect Priority	Not supported.

	Simulation tool does not provide PSTN
Authentication Type/Capability	Not supported.
Authentication Policy	Not supported.
Authentication Parameters	Not supported.
Availability Check	Supported. This attribute is used in call admission control.
Call Set-Up preference	Supported. This attribute is used in call admission control and impacts call setup delays.
Radio Inhibit	Not supported.
Manufacturer Specific Parameter	Not supported.

The movement of an SU is controlled by registering/deregistering it to the RFSSs and the process of connecting/disconnecting the SU and the Serving RFSS is instantaneous. A roaming condition for an SU is created by configuring two registrations on two different RFSSs without explicit deregistration on the first RFSS.

Once an SU is successfully registered, the serving RFSS sends periodic registration messages to maintain the association.

Note: an SU does not have an initial Serving RFSS. Unless an RFSS is registering the SU, that SU will not be reachable and will not be able to place calls.

2.3 Subscriber Group (SG)

An SG is a logical entity that represents an aggregation of SUs that are part of the group. The implementation is similar to the SU. As such, an instance of the SG is created at the Home RFSS and is used for configuration and tracking purposes while other instances are created and destroyed as needed at the Serving RFSSs. Once created, the SG's service profile can be configured. Currently, the Home RFSS only stores one profile, therefore the same information is provided during registration regardless of the Serving RFSS. Table 2 lists the various attributes for the service profile and the current support level. Reasons for nor supporting some attributes include software tool limitations and our release timeline.

Table 2: Group Service Profile Attributes Support

Name	Support Level
Access Permission	Supported. This attribute is used in call admission control.
Announcement Group	Not supported.
Priority	Supported. This attribute is used in call admission control.
Emergency Capable	Supported. This attribute is used in call admission control.
Emergency Pre-emption	Supported. This attribute is used in call admission control.

Hang Time	Supported. This attribute is used for releasing RF resources.
Confirmed Call Setup Time	Supported. This attribute defines the maximum setup time the home RFSS will wait to confirmed group calls.
Interrupt Mode	Supported/Modified. If the mode selected is priority based, the priority of the call is used not the priority of the SU.
Security	Not supported.
Interconnect Flag	Not supported. Simulation tool does not provide PSTN
Interconnect Security	Not supported. Simulation tool does not provide PSTN
Manufacturer Specific Parameter	Not supported.

The movement of an SG is not directly configurable. Instead, the Home RFSS maintains the list of Serving RFSSs currently registered. The registration/deregistration of an SG is triggered by the movement of the SUs. When an SU moves to a new Serving RFSS, this RFSS analyses the list of groups configured on the SU. For each group not already registered with this Serving RFSS, a new group registration is triggered. This means the registration only occurs when the first SU moves to the Serving RFSS. If additional SUs with the SG arrive at the Serving RFSS, no additional signaling is required. Similarly, the SG's deregistration occurs when the last SU registered for the group leaves (either via explicit deregistration or when receiving a roamed indication).

Once a Group is successfully registered, the serving RFSS sends periodic registration messages to maintain the association.

3 Call Control Management

The Call Control Manager is responsible for the setup, modification, and teardown of Unit-to-Unit calls and group calls. The user configures a call by selecting the calling SU and the called entity (another SU or an SG), the time at which the call must be initiated, the time at which it will end, and application information (i.e. the conversation). While the call configuration is done before running the simulation, the outcome of the call depends on the status of each entity involved in the call when it is initiated. SU registration status, type of call, and RF resources are some examples of parameters that can impact the outcome of a call.

3.1 Unit-to-Unit Call

A Unit-to-Unit call represents is a communication between two SUs as selected by the user. The model supports all configurations regarding the position of the SUs, creating non-collocated and collocated configurations. When two entities are collocated (for example the Calling Serving RFSS and Calling Home RFSS), the RFSSs do not exchange network messages, but the logical steps to establish the calls are still identical. For the case where all the entities are collocated, all the signaling is internal and thus network performance cannot be measured.

The Unit-to-Unit calls support several parameters allowing for different configurations as follows:

- *Initial transmitter*: In the ISSI specifications, the Called Home RFSS is responsible for determining the SU that will initially talk after the call is setup. The current implementation does not have any algorithm for such selection. Instead, the user of the model (or the GUI) is responsible for selecting the SU (calling or called) that will send the first voice spurt once the call is setup.
- *Preference*: It indicates whether the calling SU wishes to check the availability of the called SU during the call setup. Via the GUI, the default value is taken from the Service Profile of the calling SU, but it can be overwritten. This overwriting capability has been added as a convenience to perform different type of calls from the same SU within the same simulation. It is important to note that during the call setup, this value may be changed by the Called Home based on the capabilities of the called SU. Furthermore, since the RFSS is modeled as a single node (the RFG), we cannot perform the signaling required to test the availability. In the model, the Availability Check adds an additional delay to the call setup. The value of the delay is configured per RFSS since it would partially depend on the network configuration in the real system. The default value is set to 200 ms.
- *Emergency call*: It indicates if the call is an emergency. The default value is Non-emergency.
- *Initial teardown*: This attribute is not part of the ISSI specifications. Instead, it has been added for testing additional configurations. It represents the SU that will initiate the teardown, if the call successfully reaches the end as configured by the user. In some special situations (e.g. roaming and pre-emption), the call may be torn down by a separate entity.

The priority of a Unit-to-Unit call is determined by the U2U call priority attribute of the calling SU.

The model also supports call roaming for both the calling SU and the called SU. When this happens, the Home RFSS makes a best effort to re-establish the connection with the unit that has roamed. In the event that the Home RFSS fails to setup the connection with the unit on the new Serving RFSS, the call is torn down.

The call setup can fail for several reasons. Table 3 shows the checks performed during the call setup and the type or error reported when they fail. In addition, the call can end prior to the user configured stop time due to pre-emption or call priority (see Section 3.3). We note that for a Unit-to-Unit call, RTP resources and RF resources are required throughout the duration of the call.

Table 3: Unit-to-Unit Call Setup Failure Cases

Check	Error Reported
Calling/Called unit is registered (performed at both Serving and Home RFSS)	SU not registered
The RFSS starting the call is the current serving RFSS for the calling SU.	SU not registered
Calling unit has permission to place the call (System Access Permission and Emergency level)	Feature not supported
Calling unit is already in a call	SU busy
Called unit is already in a call of higher priority	SU busy
Called unit has permission to receive the call (System Access Permission and Emergency level)	Feature not supported
RTP resources not available (when entities are not collocated)	No RTP resources
RF resources not available at Serving RFSSs	No RF resources

3.2 Group Call

Group calls represent the communication between one or more SUs that have been configured with the same group. To configure a group call, the user selects the calling SU and the group to call. A successful group call involves a minimum of two RFSS: the Serving RFSS where the calling SU is located, and the SG's Home RFSS. If other RFSSs are registered (or register while the call is ongoing), the Home RFSS will invite them to join the group call. Unit-to-Unit calls and Group calls differ in several ways. Firstly, Group calls can succeed even if some Serving RFSSs do not have RTP/RF resources. However, depending on whether the call is confirmed or unconfirmed, the setup delays or Push to Talk (PTT) control will be different. Secondly, the teardown of the call is not initiated by an SU. Instead, it is the Home RFSS that decides when to end the call (based on the Tgchhangtime parameter). In the model, this has an impact in the configuration of multiple calls to the same group. Since we need to track individual calls (as per the user configuration), we cannot overlap calls. If a second call is placed while the first one is active (call stop time + Tgchhangtime), the call will be rejected and considered as a duplicate.

The following attributes can be configured for group calls:

- *Call type*: it indicates whether the call is confirmed or unconfirmed. Confirmed calls will wait for some time until most/all the resources are available before allowing PTT transmissions.
- *Priority*: The desired priority of the call. This attribute is used when handling simultaneous call requests to the same group or when handling simultaneous PTT requests during a call.
- *Emergency call*: It indicates if the call is an emergency. The default value is Non-emergency.

As in the case of Unit-to-Unit calls, the Group call setup is subject to several checks as shown in Table 4.

Table 4: Group Call Setup Failure Cases

Check	Error Reported
Group is registered (performed at both Serving RFSS and Home RFSS)	Group not registered
Group has permission to place or receive the call (Access Permission and Emergency level)	Feature not supported
Calling Serving RFSS already has an active call for the group	Duplicate call
Calling Serving RFSS has no RTP resource	No RTP resources
Incoming call request has lower priority than outgoing request	Duplicate call

3.2.1 Handling of Simultaneous Requests

Each call configured by the user is assigned an internal unique identifier. This identifier is used to match the output information (statistics) with the configuration information. When simultaneous (or quasi-simultaneous) calls to the same group are configured, the request messages can overlap (Home to Serving and Serving to Home). Based on the priority of the calls or the role of the RFSS, one request will be accepted and the other one will be rejected. It is therefore possible to have the final group call be a combination of multiple group calls, as illustrated in Figure 4. In this example, the priority (p2) of call 2 is considered higher than the priority (p1) of call 1. Consequently, the statistics measured for the resulting group call is split among the two calls and the total duration of the call is determined by the longest configured call.

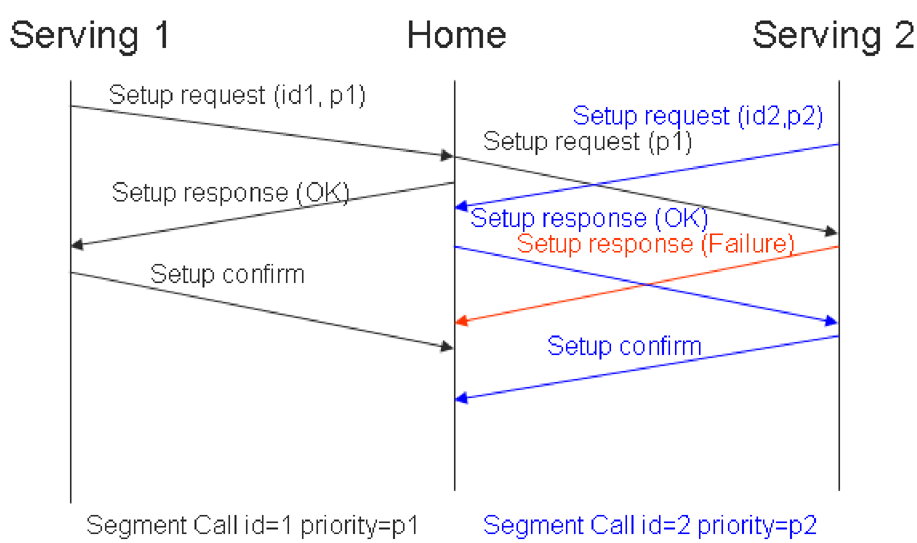

Figure 4: Simultaneous Group Call Requests

3.3 Resource Management

Adequately handling resources has an important impact on the performance and behavior of the RFSS and, in turn, the performance of P25 systems. There are two main limiting resources: the RTP resources to communicate between RFSSs and RF resources to communicate between the RFSSs and the SUs.

3.3.1 RTP Resources

The RTP resources represent the available ports on the RFSS to transmit or receive PTT data. Information about the port allocated by an RFSS is sent via the Session Description Protocol (SDP) and carried in the SIP messages. Per definition, when RFSS A sends its port information to RFSS B, it indicates the port to which data needs to send information. It is then implementation dependent whether RFSS B allocates one port for sending and one port for receiving or if a single duplex port is used. The current model implements the second case (single duplex port). Because of the various call configurations, the number of RTP resources required for Unit-to-Unit calls or Group calls varies. As shown in Figure 5, that number varies from 0 (when all entities are collocated) to 6 for Unit-to-Unit calls, and varies based on the number of Serving RFSSs for Group calls.

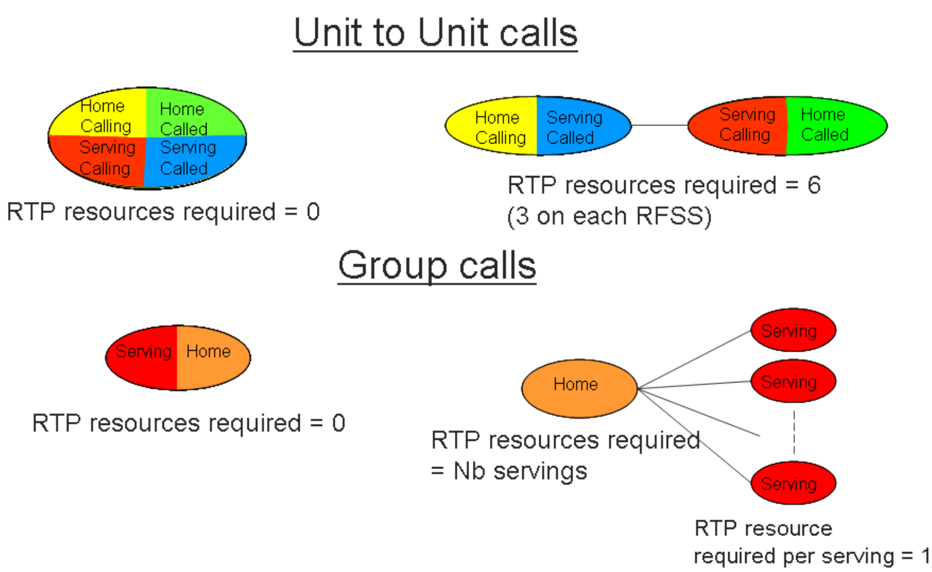

Figure 5: Example of RTP Resource Allocation

In the current implementation, the user can configure the maximum number of RTP resources available at an RFSS (default is 100 ports). When calls are setup, requests are made to the resource handler to obtain an RTP resource. If this fails, the request is queued according to call priority as detailed in Section 3.3.3 and a timer is created. If a resource becomes available before the timer expires, the request is granted and the Call Control Manager is notified to complete the call setup. If the timer expires, the Call Control Manager is also notified and the call fails. The current timer value is set to 30 s based on the assumption that SIP messages are retransmitted for up to 32 s, thus allowing the failure response to come back before the SIP transaction expire (see Section 5.2 for details).

3.3.2 RF Resources

The RF resources are modeled as abstract resources (i.e. they are not mapped to any particular frequency), since the CAI is not supported. They are required at the Serving RFSSs to enable communication between the RFG and the SU. For Unit-to-Unit calls, two RF resources are always required while it varies according to the number of Serving RFSSs for Group calls. The management of RF resources is identical to the RTP resource management. The model assumes that there is always a control channel for exchanging signaling messages between the RFG and the SU. However, the user can configure the number of RF data channels (default is 10 channels).

For Unit-to-Unit calls, the RF resources are allocated for the duration of the call. For Group calls, the allocation is dependent on the type of call and the RF hang time. During the course of a Group call, the Serving RFSS may release the RF resource used by the call due to preemption or RF hang time. If the call is confirmed, the Serving RFSS sends a message to the Home RFSS to indicate a change in the RF resource allocation state.

3.3.3 Priority and Pre-emption

Priority defines the importance of the calls. The ISSI specification defines priorities at several levels. Firstly, a call can be tagged as an emergency or non-emergency call. Secondly, there are Unit-to-Unit calls and Group calls. Finally, there are priority levels within the Unit-to-Unit calls and Groups calls, respectively based on the service profile of the Calling SU and the SG. We categorize the calls in four groups, for higher priority to lower priority: Emergency Group Call, Emergency Unit-to-Unit Call, Non-emergency Group Call, and Non-emergency Unit-to-Unit Call. Within each category, the priority attribute is used to classify the calls.

The different categories are used to determine when a busy SU needs to switch to a new incoming call (and potentially teardown its current call) as well as in the queuing of RTP/RF allocation requests.

Preemption of resources (RTP or RF) occurs when an Emergency Group Call is configured with ruthless preemption. In the event that no resource is available at the RFG, the call with lowest priority has its resource removed. If the preempted call is a Unit-to-Unit call, it is torn down. If the preempted call is a Group Call, the call segment is modified to indicate the loss of RTP resource, and/or if the call is confirmed for the loss of RF resource. In addition, the preempted Group Call is placed in the resource request queue to obtain a new resource when it becomes available.

4 Media Transmission

The media transmission modules are responsible for handling the push-to-talk (PTT) requests and for carrying the voice payloads from one RFSS to another. The behavior of the users is modeled in the Application module and is responsible for starting and ending a call as well as starting and ending PTT spurts.

4.1 PTT Control

Several primitives are available to control the PTT spurts. Table 5 shows the list of primitives and their current support.

Table 5: Supported Transmission Control Functions

Transaction	Support
TC	Supported.
TC_SpurtRqst/TC_Decision	Supported (see below for current algorithm)
TC_DecisionTimeout	Supported. If no answer is received, the local policy is to deny the spurt.
TC_IsUnstoppable	Not supported
TC_SetLosing	Supported. The losing bit is managed by the Transmission Control Manager in group calls. It is set based on the decision to grant/deny new spurts.
TC_SpurtStart	Supported. Used to request a new spurt
TC_SpurtEnd	Supported. Used to indicate the end of a spurt.
TC_Audio	Supported. Used for every audio packet.
TC_Mute	Supported. An RFG configured with the mute capability will request the losing transmissions to be muted.
TC_Error	Supported. This error is generated when an RFG stops receiving PTT Heartbeat messages from a peer.

For Unit-to-Unit calls, the spurts are always granted. The implementation of the application is such that one user will try to avoid starting a new spurt while it is receiving a spurt.

For Group calls, many SUs can trigger PTT requests to transmit audio at the same time (or at very close time). The Home RFSS uses the attributes of the call (call type and interrupt mode) as well as the priority of the spurt to determine if the request is granted, denied, or waiting. The priority of the spurt is currently defined as the priority of the call that was used to setup the call (multiple segments can have different priorities if there were simultaneous requests during the call setup). Figure 6 shows the decision algorithm implemented.

If a spurt request does not receive any response, the local policy on the requesting RFG decides whether to self grant the spurt (and start transmitting) or to deny the spurt. In the current implementation, the local policy is fixed and the behavior is to deny the spurt.

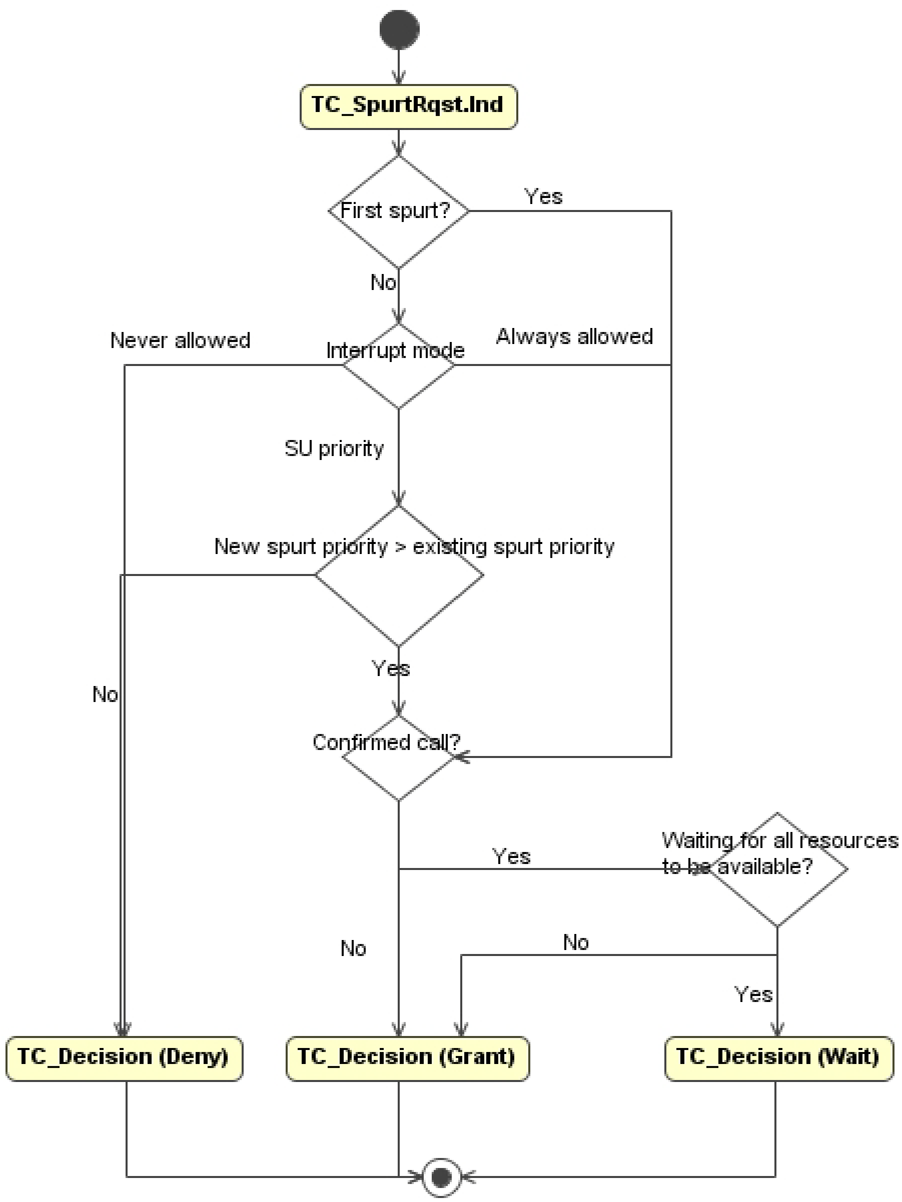

Figure 6 : PTT Decision Algorithm for Group Call

Whenever a new spurt is granted, the losing bit of the updated for the previous active spurt. Only audio packets received with the losing bit set to 0 are transmitted to the SU.

Once an RTP session has been opened between two RFGs, Heartbeat messages are exchanged to monitor the connection (default interval: 10s). In addition, Heartbeat messages are also sent on muted spurts (default interval: 0.5s).

4.2 Application

The Application module aims to model the behavior of users on a PTT system. For a Unit-to-Unit call, it would be similar to two people talking on the phone where generally one user talks after the other with some exceptions when both try to talk at the same time of that a user speaks multiple times because the peer user does not respond quickly.

The Application can be configured in two separate ways:

- Statistically: the interval and the duration between each PTT spurts are based on a normal distribution with a mean value configurable by the user.
- Trace-based: using an audio conversation, spurts can be assigned to different users. In this mode, we are able to record the messages received by each user and play back the audio that it would have received.

The user is also capable of configuring the number of voice samples (1, 2, or 3) that are carried in each packet.

There is one instance of the Application per call per user. When the call is created, the Application on the calling SU is initialized. The Applications for the called SU or the members of the group are created when the SUs effectively join the call.

4.2.1 Statistical Traffic Generation

The state machine for a Unit-to-Unit call is shown in Figure 7. If the call is successful, the first SU to trigger a PTT is determined by the initTx attribute (user configurable). After that, spurts are generated based on a normal distribution. The call ends if the call setup fails or if the stop time of the call is reached. In the later case, the SU configured as the initTeardown by the user will initiate the call teardown.

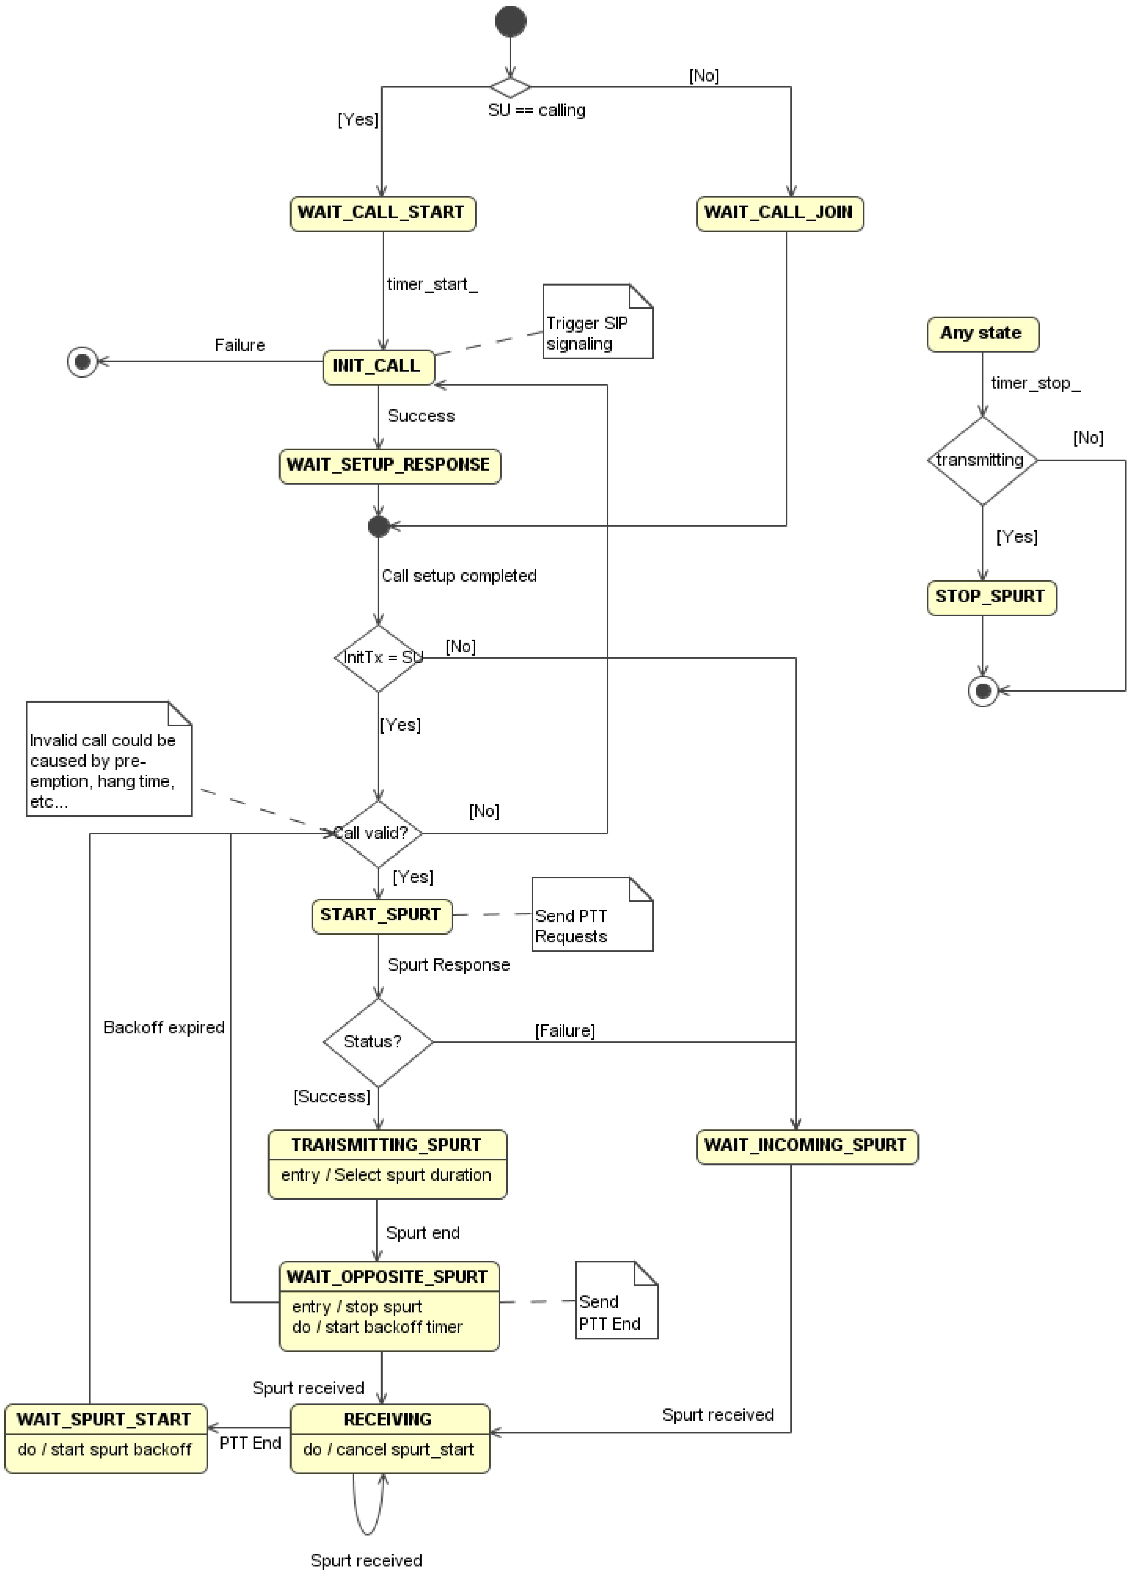

Figure 7: SU Call Application State Machine

The state machine for a Group call, shown in Figure 8, is very similar to the Unit-to-Unit model. However, only the calling SU may send the first PTT spurt.

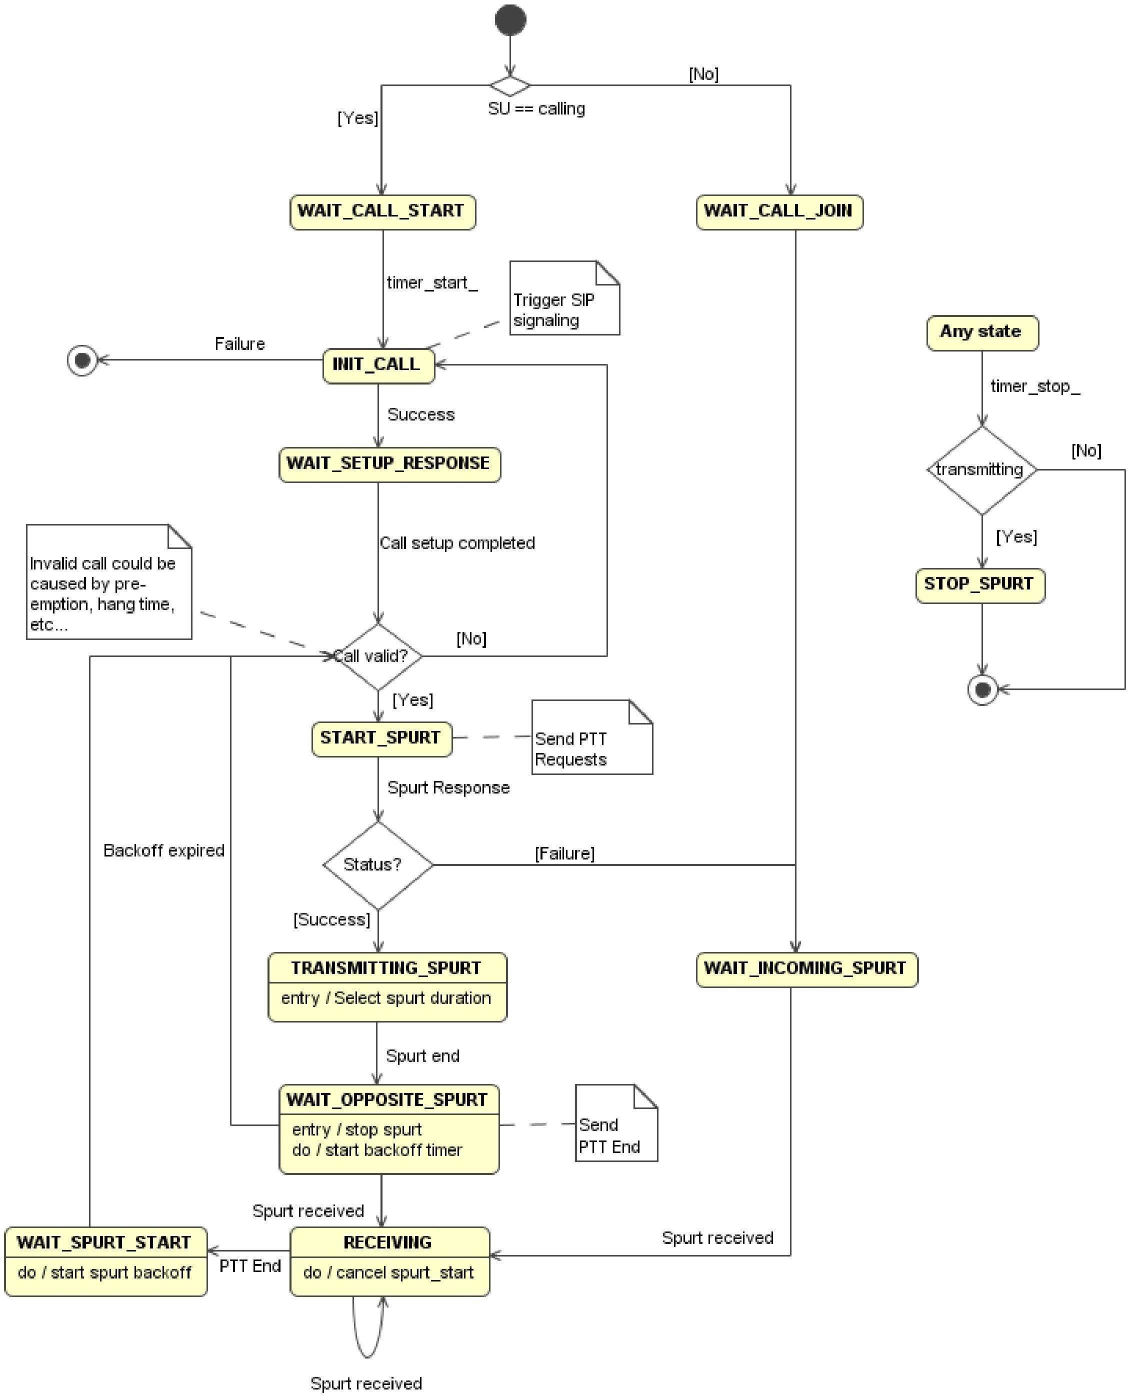

Figure 8: Group Call Application State Machine

The use of statistical traffic generation is recommended for scenarios where the content of the transmission is not relevant and when setting up large scale scenarios.

4.2.2 Trace-based Traffic Generation

An alternative traffic generation is based on the traces derived from an audio file. The time and duration of each spurt is then fixed by the traces. Therefore, there is no backoff mechanism and it is possible to send and receive at the same time if the spurts overlap. By recording the packets received by each SU, we can create an audio file containing all the voice data and listen to it during the playback of the simulation. Due to the overhead needed to create the audio traces, the use of trace-based traffic is recommended for demonstrations.

5 Protocol Modeling

This section describes the elements of the protocols' implementation that may affect the expected behavior of the ISSI.

5.1 ISSI

The SUID and SGID contain the WACN-ID and SYSTEM-ID of the home RFSS. However the serving RFSS still needs to discover the RFSS-ID of the home RFSS. In the model this discovery procedure is instantaneous.

5.2 SIP

To obtain the address of a SIP server based on its URL, a discovery mechanism such as Domain Name System (DNS) is necessary. In the model, this discovery procedure is instantaneous.

5.3 RTP

RTP header fields are used only in computing the packet size.

6 Known Issues and Limitations

- The packet size is approximated since we carry additional metadata in the packets and some fields are not used.
- No support for Fixed Network Element (FNE).
- No support for Supplementary Data.

7 References

- Project 25 Inter-RF Subsystem Interface: Messages and Procedures for Voice and Mobility Management Services (TIA-102.BACA-A January 2009)
- The Network Simulator – ns-2 (http://www.isi.edu/nsnam/ns/)
- RFC 3261 "Session Initiation Protocol"
- RFC 4566 "SDP: Session Description Protocol"