A Congestion Notification Technique for SMDS Networks

P. Havala, H. Kaur, M. Sosa and M. Ulema

Bellcore
331 Newman Springs Road, Red Bank, NJ 07701

Abstract

The proposed congestion notification strategy for SMDS networks presented in this paper is based on explicit notification, which uses control messages to inform other network elements of congestion. Network elements respond to explicit notifications by matching the traffic entering the network to the network loading. For an effective congestion management strategy, the notifications should carry sufficient information about the congested resource so that the notified Switching Systems (SSs) may reduce traffic towards the congested SS. Since the internal architecture of an SS is not visible to other SSs in the network, the congestion notification mechanism maps the SS resource congestion to path congestion, identified by open interfaces such as the inter-SS, inter-carrier, and customer-to-network interfaces, which are visible to all SSs in the network.

1. Introduction

Switched Multi-megabit Data Service (SMDS) is a data transport service that provides LAN-like performance across wide geographical areas. A key feature of SMDS is its *connectionless* method for switching user packets to provide a seamless networking environment for the end users on LANs, which are typically connectionless. LANs and customer premises equipment (CPE) interface with SMDS using the Subscriber-Network Interface (SNI)[1]. Switching Systems (SSs) in a network are connected across Inter-Switching System Interfaces (ISSIs) using the ISSI Protocol (ISSIP)[2]. A network supporting SMDS in a Local Access and Transport Area (LATA) can be connected to other networks across an Inter-Carrier Interface using the ICI Protocol (ICIP)[3]. To support service features and performance objectives published for SMDS[1][3], a network consisting of a number of SSs must provide, among other things, the necessary features to route a packet from its source SS to its destination SS. The Routing Management Protocol (RMP) is based on the Internet standard Open Shortest Path First (OSPF) protocol[4].

A congestion management strategy needs to be in place to alleviate congestion and prevent degradation of network performance. The congestion management strategy for SMDS[2] has three components: measurement of congestion, notification, and discard. SSs in the network monitor their resource usage with the measure of congestion (MOC) and detect congestion when the MOC exceeds congestion thresholds. SSs send explicit congestion notifications (ECNs) to all SSs in the network. SSs receiving ECNs help alleviate congestion by discarding packets directed towards the congested resource. The Congestion Management Protocol (CMP) dictates the sending and receiving of ECNs as well as the contents of the ECN messages[2]. Congestion measurement and discard policies for SMDS networks are described in [2].

For an effective congestion management strategy, the notifications should carry sufficient information about the congested resource so that the notified SSs may reduce traffic towards the congested SS. Since the internal architecture of an SS is not visible to other SSs in the network, the congestion notification mechanism maps the SS resource congestion to path congestion, identified by open interfaces such as the ISSI, ICI or SNI, which are visible to all SSs in the network. Variations of this scheme can be used to identify congested links, link sets or both.

This paper describes a congestion notification technique proposed for SMDS networks. The connectionless nature of SMDS and the limited visibility an SS has of other SS architectures make notification of internal resource congestion to other SSs a challenging task.

The organization of the this paper is as follows: Section 1 provides background for the path-based notification schemes. Section 2 provides major components of the path based schemes. Section 3 discusses the four path-based notification schemes that were analyzed. Section 4 contains recommendations. Finally, Section 5 summarizes the paper.

2. Elements of a notification technique

This section describes the elements of ECN messages based on the path-based notification scheme, the distribution of the messages and determination of the affected traffic from the contents of the ECN message at the notified SS.

SSs in a SMDS network monitor their resources to determine resource availability[3]. A smoothed MOC is used to measure the level of congestion of the resources. (The term Congestion Level is used for this purpose.) The resource congestion is mapped into the affected open interfaces, and all the SSs in the network are notified using ECN messages that are carried in Congestion Management Protocol Data Units (CMPDUs). SSs use the CMP for sending and receiving ECN messages. The CMP defines the ECN message format, the CMPDU format and the procedures for sending and receiving ECN messages.

Structure of ECN message: Congestion in an SS module or Network Element (NE) causes paths through that SS to be congested. Through a series of examples, this section describes the relationship between the congested resource and the paths through the SS that are affected, as well as the elements of the ECN message needed to identify the congested paths.

The examples below refer to the SS shown in Figure 2-1, where SS 1 serves ISSI links to five other SSs. A set of links between two adjacent SSs form a *link set*. For example, Links 1 and 2 between SS 1 and SS 2 form a link set. SS 1 comprises three SS NEs, or nodes. Each node serves one or more ISSI links and one or more supplier-specific intra-SS links (i.e., Proprietary Interfaces); therefore each node has one or more ISSI modules (ISSIMs) and one or more Proprietary Interface Modules (PIMs). In the examples that follow, congestion in a module is assumed to affect outgoing traffic (i.e., those packets leaving a congested module) only.

Example A -- ISSIM congestion; output link set affected: In this case, the ISSIM marked with an "A" in Figure 2-1 is congested at Congestion Level 2 (CL 2). Congestion in this ISSIM will result in degraded performance for all packets that enter SS 1 and are headed to (or through) SS 5. To notify other SSs in the network of this congestion, the ECN messages sent by SS 1 need to convey to other SSs that all paths through SS 1 to SS 5 are congested at CL 2 (see Table 2-1).

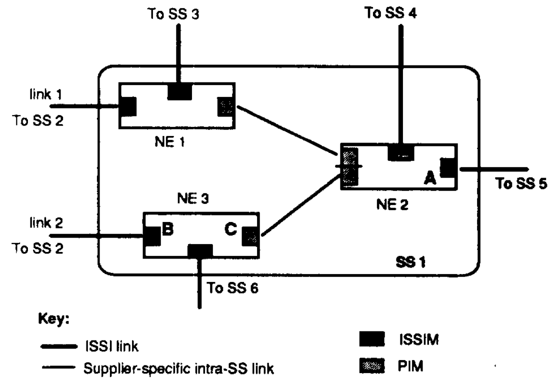

Figure 2-1. Example SS

Table 2-1. Elements of ECN message for Example A

Congested SS	To SS	CL
1	5	2

Example B -- ISSIM congestion; output link affected: In this case, the ISSIM marked with a "B" in Figure 2-1 is congested. Congestion in this ISSIM will result in degraded performance for all packets that enter SS 1 and are headed to (or through) SS 2 *over link 2*. Packets exiting SS 1 over link 1 experience no such degradation.

As in Example A, SS 1 should notify other SSs that paths through SS 1 and SS 2 are congested. However, the ECN should also indicate that only link 2 to SS 2 is affected (see Table 2-2). If the ECNs sent by SS 1 contained only the information in Table 2-1, other SSs in the network might discard *all* packets headed to SS 2 through SS 1. The presence of an output link identifier in an ECN message enables the optimization of packet discard at other SSs when some, but not all, of the links in a link set at the congested SS are affected by the congestion.

Table 2-2. Elements of ECN message for Example B

Congested SS	To SS	Output Link	CL
1	5	2	2

Example C -- PIM congestion: In this example, a PIM marked "C" in Figure 2-1 is congested. Packets that enter

NE 3 (over link 2 from SS 2 or the link from SS 6) and exit SS 1 over the ISSI links served by NE 1 or NE 2 may be lost or delayed as a result of this PIM congestion. The affected paths are as follows:

- The path from SS 2 over link 2, through SS 1, to either SS 3, SS 4, or SS 5.
- The path from SS 6, through SS 1, to either SS 2 over link 1, or SS 3, SS 4, or to SS 5.

Table 2-3 shows the elements of the ECN message in this case.

Table 2-3. Elements of ECN message for Example C

From SS	Input Link	Congested SS	To SS	Output Link	CL
6	—	1	2	1	2
6	—	1	3,4,5	—	2
2	2	1	3,4,5	—	2

All other kinds of SS congestion may be viewed as the superposition of congestion in ISSIMs and PIMs, so that the elements shown in Table 2-3 allow optimization of packet discard at notified SSs for all types of congestion.

Distribution and receipt of messages: When the MOC of a resource crosses the onset or abatement thresholds, ECN messages are generated to inform other SSs of the change in congestion level. Congested SSs flood the CMPDU throughout the network using a mechanism based on Shortest Path Spanning Tree (SPST) to inform other SSs of their congestion.

On receiving an ECN message, an SS determines the traffic that will pass through the congested resource. An SS that receives an ECN message needs to first determine if the congested path segment is on its route to other destinations. If so, the notified SS determines the affected traffic that will traverse the congested paths and then discards a portion of the originating traffic[2].

Procedures to determine if the congested path segment is on the notified SS's routes to other destinations: SSs transport data units using SPSTs created by the RMP. (Each SS generates an Individually Addressed and a Group Addressed Routing Table from the SPSTs, retaining the next hop SS information in the tables.) Figure 2-3 illustrates the topology of an example SMDS network and the SPST rooted at SS 1, assuming all links have equal capacity.

When an SS receives an ECN message, it needs to determine if the path described in the message is on its SPST. If it is not, then the congestion in the path segment does not affect any traffic from the notified SS and thus the notified SS can ignore the ECN message. To determine if a path segment is on its SPST, an SS maintains the complete spanning tree.

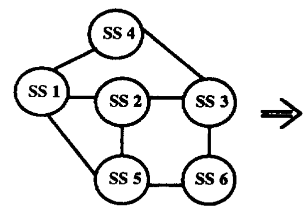

Dest. SS	Previous SS
1	-
2	1
3	2
4	1
5	1
6	5

Network Topology — A Shortest Path Spanning Tree at SS 1

Figure 2-2 An example network topology

The notified SS determines if the specified path is on the SS's SPST as follows:

- If the ECN has (From_SS, Congested_SS, To_SS) and (From_SS, Congested_SS) and (Congested_SS, To_SS) are two hops on the SS's SPST, then the congested path segment is on the notified SS's SPST and may affect the traffic from the notified SS.
- If an SS has ISSIM congestion over the link set (Congested_SS, To_SS), then the affected path segments can be represented as (Congested_SS, Congested_SS, To_SS). In this case, if (Congested_SS, To_SS) is a hop on the notified SS's SPST, then the congested path segment is on the notified SS's SPST and may affect traffic from the notified SS.
- If an SS has a SS-wide congestion then the affected path segments can be represented as (Congested_SS, Congested_SS, Congested_SS). If an SS has SS-wide congestion then that congestion will affect all traffic from the notified SS that will traverse the congested SS.

Therefore if From_SS is the same as Congested_SS, then all incoming path segments to Congested_SS are considered to be a hop on the affected path. Similarly, if To_SS is the same as Congested_SS, then all outgoing path segments from Congested_SS are considered to be a hop on the affected path segment.

Determination of the affected traffic: If an ECN message only carries congested path segment information then all the traffic originating at the notified SS and destined for To_SS and SSs below To_SS in the notified SS's SPST are assumed to be affected by the congestion.

In addition to the path segment information, if the ECN message includes congested link information, then the notified SS needs to determine the traffic that will traverse the congested links. SSs use a common load splitting algorithm to determine the specific link in a link set that is used for transporting data units[3]. The load splitting algorithm is a deterministic algorithm which any SS in the network can use to determine if their traffic will traverse the congested link. If information about individual links is coupled with the load splitting algorithm, then the CMP can use this to determine whether packets will pass through the congested links.

3. Alternative notification schemes

This section discusses path-based notification alternatives. Although these alternatives are all based on the generic technique discussed in Section 2, they differ in the level of granularity each gives about the affected path segment, and thus in the complexity needed to implement the strategies.

3.1 Link-to-link

This notification strategy identifies traffic affected by a congestion event within an SS by listing pairs of ISSI links. Traffic which enters the SS over the first ISSI link in a pair and which is supposed to be transmitted by the SS over the second ISSI link in the pair should be considered as traversing the congestion event.

This alternative can be used to describe PIM as well as ISSIM congestion events. Thus, all the information needed for the Example C case could be encoded in ECN messages. This provides a fine level of granularity for determining the specific traffic which would actually experience congestion, so that only this traffic will be affected by the congestion management scheme. However, even this detailed scheme is not exact, as the routing and load splitting procedures used internally within an SS may not map all traffic between two ISSI links over the same SS-internal path.

ECN Message Format: To support this level of path segment granularity, an ECN message would use the fields shown in Figure 3-1.

SS_ ID	From_ SS	To_ SS	Low_ Key	High_ Key	In_ Link	Level

Figure 3-1. ECN format for link-to-link case

The meaning and usage of these fields is as follows:

- *SS_ID* field gives the SS ID of the SS experiencing congestion (the one which first originates the ECN). SS ID is a six decimal digit binary coded identifier and *SS_ID* is a three octet field.
- *From_SS* and *To_SS* fields each list the SS ID of an SS neighboring the SS experiencing congestion. These SS IDs are each six binary-encoded decimal digits, making each of these fields three octets. The values inserted in the *From_SS* and *To_SS* fields are such that the ECN indicates congestion on some path segment from the SS with SS ID *From_SS*, through the originating SS (the SS with SS ID *SS_ID*), to the SS with SS ID *To_SS*.
- *Low_Key* and *High_Key* fields are used to indicate the links within the link set between the SSs with SS IDs *From_SS* and *SS_ID* and the links within the link set between the SSs with SS IDs *SS_ID* and *To_SS* that are affected by the congestion. These fields rely on the standard load splitting algorithm[2]. The keys assigned to a particular link should be contiguous. This allows the ECN to specify a range of keys to refer to a set of links which are affected by the congestion. *Low_Key* is simply the smallest key in the key range, and *High_Key* is the largest key in the key range. Each of these two fields is 2 octets, based on the size of the load splitting key.
- *In_Link* field is used in the first phase of the distribution scheme to indicate which link from the SS with SS ID *From_SS* is in the congested path. In the example in Figure 3-2, the *In_Link* field should be set to A.

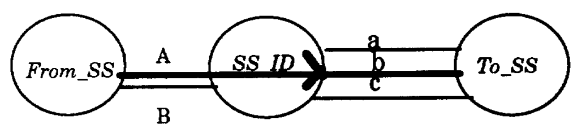

Figure 3-2. Example congestion event

Table 3-1. Load splitting assignment

Link	Low Key	High Key
a	0	21844
b	21845	43689
c	43690	65535

- *Level* is a four bit binary-encoded field. It indicates the level of congestion being experienced by traffic which follows the described path.

ECN message distribution: Once an SS determines that it is congested, and a particular link-to-link path segment is affected, the SS encodes the congestion information about the link-to-link congestion within an ECN message. The *From_SS*, *SS_ID*, and *To_SS* fields are simply encoded with the SS IDs of the SSs on the path segment. However, the identity of the particular links within the link sets which are affected need to be mapped to the *Low_Key* and *High_Key* fields. For the link set from the congested SS to the *To_SS*, the congested SS has all the information needed to do this. The congested SS knows the key range assigned to the affected outgoing link, and places the end values of this range in the *Low_Key* and *High_Key* fields. However, this does not indicate which link from the *From_SS* is really experiencing congestion. In fact, the congested SS cannot directly encode this information into the ECN fields because it does not know the precise division of key ranges to links that the neighboring SS is using over its link set to the congested SS.

Therefore, to support link-to-link notification, the proposed notification scheme uses a two step notification process:

Step 1 -- The congested SS sends a single copy of the ECN to the SS with SS ID *From_SS*. The *In_Link* field is used to indicate a local identifier of the link from *From_SS* to *SS_ID* which is on the congested path. This requires the local link identifiers to be maintained consistently for both ends of a link set. That is, if the congested SS specifies the link with an ID of 3 as congested, the *From_SS* needs to understand to which link this refers. These identifiers only need to be maintained across both ends of a link set, and do not need to be maintained globally.

Step 2 -- *From_SS* can then determine the key range it has assigned to the link referenced in the *In_Link* field of the ECN it receives. This key range is then compared with the key range given by the *Low_Key* and *High_Key* fields. The overlap of the two key ranges is the actual key range that is affected by the congestion. The *From_SS* replaces the values in the *Low_Key* and *High_Key* fields with the corresponding values for the new range. If the new range is empty, the ECN is simply discarded. Otherwise, *From_SS* distributes the ECN along its group addressed spanning tree. Note that the *SS_ID* field still specifies the congested SS.

Processing required at SSs: The congested SS needs to determine the affected path segments in terms of the previous SS, the next hop SS, the local identifier of the incoming link, and the key range of the outgoing link.

When it receives the ECN, *From_SS* needs to determine the key range of the link in the link set to *SS_ID* with local identifier *In_Link*. It needs to then determine the intersection of this key range with the key range specified in the received ECN.

When SSs receive the ECN sent by From_SS, they need to use the SS IDs in the ECN to determine if the path segment is on the shortest path spanning trees used for routing. The affected traffic is determined as described in Section 2.

3.2 Link Set-to-Link

In this notification strategy, ECN messages identify affected traffic within the congested SS by describing an incoming link set and an outgoing link. The affected traffic are those packets that enter the congested SS over any link in the incoming link set, and that are destined to leave the congested SS over the outgoing link. This notification strategy represents a simplification of the link-to-link scheme, since the *In_Link* field of the ECN message (see Figure 3-2) and its associated procedures (e.g., the two-step distribution process described in Section 3.1) are not required.

This alternative may be used to describe ISSIM congestion events, as illustrated in Examples A and B in Section 2. The link set-to-link scheme is also effective in describing PIM congestion events where all links of the incoming link set are affected. In Example C in Section 2, the link set-to-link scheme can describe the congested paths from SS 6, through SS 1, to SS 2, SS 3, SS 4, and SS 5. However, the scheme can not effectively describe the congestion paths from SS 2 through SS 1 (and to SSs C, D, and E), since only one link in the link set from SS 2 to SS 1 is affected.

ECN message format: The ECN message for this alternative has the format shown in Figure 3-3.

SS_ID	From_SS	To_SS	Low_Key	High_Key	Level

Figure 3-3. ECN format for link set-to-link case

ECN message distribution: When an SS determines that it should send an ECN message, it encodes the information about the link set-to-link congestion within the ECN message. As in the link-to-link case, the congested SS encodes the *From_SS*, *SS_ID*, and *To_SS* fields with the SS IDs of the SSs on the path segment. The congested SS indicates the congested link (within the link set) to the *To_SS* by populating the *Low_Key* and *High_Key* fields as described in Section 3.1. The ECN messages are then flooded throughout the network along the congested SS's group addressed spanning tree. Because the *In_Link* field is not present in the link set-to-link alternative, the two-step process described in Section 3.1 is not required.

Processing required at SSs: The congested SS needs to determine the affected path segments in terms of the previous SS, the next hop SS, and the key range(s) of the link(s) to the next hop SS.

When other SSs receive the ECN sent by the congested SS, they need to use the SS IDs in the ECN to determine if the path segment is on their shortest path spanning trees used for routing. The affected traffic is determined as described in Section 2.

3.3 Link Set-to-Key Bin

This notification strategy represents a simplification to the link set-to-link notification scheme. The purpose of the modification is to limit the number of key ranges which have a congestion state associated with them, as well as to provide fixed key ranges to simplify notified SS procedures.

This alternative may be used to describe ISSIM congestion events, as illustrated in Examples A and B in Section 2. It also does as well as the link set-to-link method in describing the congestion in Example C (assuming each link in a link set has an equally-sized key range). However, in general, this alternative will not always provide the same level of granularity that can be achieved with the link set-to-link alternative.

In this notification strategy, ECN messages identify affected traffic within the congested SS by describing an incoming link set and an outgoing key range. Traffic which enters the SS over the first ISSI link set and which has a key within the key range, within the granularity of the key bins (see below), is considered as traversing the congestion event. The key range may not correspond one-to-one with keys assigned to the outgoing links. Rather, the complete key space is divided into fixed bins and these bins are used to determine which traffic streams are affected by the congestion.

ECN message format: To support this level of path segment granularity, an ECN would use the same fields as the link set-to-link notification scheme as shown in Figure 3-3. The key range in the ECN indicates the keys on the outgoing link set which are affected by the congestion. However, the congested SS may, although not required, take into account the key bin divisions used by the notified SSs.

ECN message distribution: When an SS determines that it should send an ECN message, it encodes the information about the congestion within the ECN message. The ECN distribution scheme is the same as for the link set-to-link alternative. The ECN message is flooded throughout the network along the congested SS's group addressed spanning tree.

Processing required at SSs: The congested SS needs to determine the affected path segments in terms of the previous SS, the next hop SS, and the key range on the outgoing link set affected by the congestion. This is the same as for the link set-to-link notification alternative. However, when specifying the key range affected by congestion at a resource, the congested SS may decide to use a rough level of granularity, based on its knowledge of the notified SS procedures. This may limit the state information it needs to store about its congestion events. Note that the link set-to-link scheme would not necessarily prevent this. The real difference comes in the notified SS procedures.

When SSs receive the ECN sent by the congested SS, they need to use the SS IDs in the ECN to determine if the path segment is on their shortest path spanning trees used for routing. If not, the ECN can be discarded.

A notified SS divides the key space for each destination SS into *NumBins* equally sized ranges, where *NumBins* is a network-wide constant and must be a power of two. The proposed value for *NumBins* is four. Rather than maintain state for arbitrary key ranges, a notified SS maintains the congestion state for each destination SS only per each of these fixed *key bins*.

Therefore, when it receives an ECN that does affect its shortest path spanning tree, the notified SS must first determine which of the *NumBins* key bins are affected by the ECN. Since *NumBins* is a power of two, it is a simple procedure to determine which key bins have more than half of their keys contained in the range specified by the ECN. Those key bins are the ones that are considered affected by the ECN.

From there on, the notified SS procedures are the same as for the link set-to-link set notification scheme, except that the state information is maintained per affected key bin per destination SS, rather than just per destination SS.

In addition, when receiving a data unit over an SNI, the SS must determine the key for the data unit before deciding whether the data unit should be dropped.

For small values of *NumBins*, such as four, this method seems quite reasonable, yet provides finer granularity of congestion events than simple link set-to-link set notification. For example, if *NumBins* equals 4 and an outgoing link set has one, two, or four links in it, all of equal capacity, and each link has an equally sized key range (e.g., default key range assignment is used), then the key bin method gives perfect link set-to-link notification. For other link set sizes, the key bin notification will not be exact, but still allows more precision in congestion notification than a link set-to-link set notification scheme.

3.4 Link Set-to-Link Set

In this scheme, congested SSs notify other SSs of a congested path segment identified by the affected link sets. When an SS experiences congestion, it maps its congested entities, PIMs, ICIMs, ISSIMs to affected link sets. The affected path segments are defined in terms of a three SS path segment.

This scheme is useful when the majority of the links in a link set are congested. This alternative may be used to describe ISSIM congestion events, as illustrated in Example A in Section 2. Since the granularity of specification is a link set, this scheme may result in under-specification in cases where congestion affects traffic over specific links only. The procedures at notified SSs are simpler as compared to the notification methods described earlier since it is assumed that all links in a link set are at the same level of congestion. The notified SS therefore needs to only determine if the congested path segment is on its spanning tree. When congested links are specified, the SS also needs to execute the load splitting algorithm on the destination and source addresses of the data units it is transporting, to determine if the traffic also traverses the specified links in a link set. This scheme is appropriate for resource congestion that affects link sets rather than links.

ECN message format: The format of the ECN message used is shown in Figure 3-4.

SS_ID	From_SS	To_SS	Level

Figure 3-4. ECN format for link set-to-link set case

ECN message distribution: When an SS determines that it should send an ECN message, it encodes the information about the link set-to-link set congestion within the ECN message. The SS encodes the *From_SS*, *SS_ID*, and *To_SS* fields with the SS IDs of the SSs on the congested path segment. The ECN messages are then flooded throughout the network along the congested SS's group addressed spanning tree.

Processing required at SSs: The congested SS needs to determine the affected path segments in terms of the previous SS and the next hop SS to populate the ECN message.

When other SSs receive the ECN sent by the congested SS, they need to use the SS IDs in the ECN to determine if the path segment is on their shortest path spanning trees used for routing. Since affected key ranges are not advertised, all the links in the specified link sets are assumed to be congested.

4. Analysis of alternatives

Table 4-1 compares the four path-based notification schemes presented in Section 3. From the analysis in Section 3, it is clear that the distribution scheme needed to support the **link-to-link** notification alternative is the most complex. Besides being complex, the two step notification process adds a large amount of delay between the times when the congested resource first sends an ECN message and when other SSs receive the notification and can respond to it.

Table 4-1. Comparison of the four notification schemes

Features	Link-to-link	Link set-to-link set	Link set-to-link	Link set to-key bin
Congestion events identifiable by ECN messages	ISSIM, PIM	ISSIM	ISSIM, Sub-set of PIM congestion events	ISSIM, Sub-set of PIM congestion events
ECN distribution	two stage process	one stage process	one stage process	one stage process
Processing at congested SS	identifying affected path segments and links	identifying affected path segments	identifying affected path segments and links	identifying affected path segments and bins
Processing at notified SSs	procedures in Sec. 3	procedures in Sec. 3	procedures in Sec. 3	procedures in Sec.3
Max. no. of congestion states/dest/SS	2^{16}	1	2^{16}	*NumBins*

The **link set-to-link set** notification scheme is the simplest of the path based notification schemes. The congested SS need only identify pairs of neighboring SSs affected by congestion at a resource. This level of granularity helps limit the size of tables and the number of ECNs that the congested SS may need to send and maintain. The processing at notified SSs is also simplified as they only need to work at the granularity of link sets (neighboring SSs), and don't need to compute the key for a data unit before determining if it should be dropped. Therefore, the link set-to-link set notification scheme makes sense for congestion events that are accurately described at that level of granularity.

Unfortunately, as discussed earlier, not all congestion events can be accurately described by link set-to-link set notification. The **link set-to-link** notification scheme does seem to capture many congestion events with a suitable level of granularity. It also has a single step notification scheme. The problem is that there are 2^{16} possible keys. Thus, there are theoretically 2^{16} different key ranges for which a notified SS would need to maintain state for each destination SS. Therefore, some rules are clearly needed to limit the number of key ranges for which a notified SS needs to maintain state, as well as to simplify implementations.

The **link set-to-key bin** approach is a modification to the link set-to-link notification scheme. The purpose of the modification is to limit the number of key ranges with which a congestion state is associated, as well as to provide fixed key ranges to simplify notified SS procedures. In this alternative, the congested SS performs the same procedures as in the link set-to-link scheme. However, when specifying the key range affected by congestion at a resource, the SS may decide to use a rough level of granularity, based on its knowledge of the notified SS procedures. The notified SS maintains the congestion state for each destination SS only for each of the fixed key bins. For small values of *NumBins*, such as four, this method seems quite reasonable, yet provides finer granularity of congestion events than simple link set-to-link set notification. For example, if *NumBins* equals four and an outgoing link set has one, two, or four links in it, all of equal capacity, and each link has an equally sized key range (e.g., default assignment is used), then the key bin method gives perfect link set-to-link notification. For other link set sizes, the key bin notification will not be exact, but will allow more precision in congestion notification than the link set-to-link set scheme.

Because of the simplicity and its low processing overhead at the notification sending and receiving SSs, the *link set-to-link set* scheme was selected for SMDS in Reference [2]. However, to more accurately describe congestion events (including some types of PIM congestion), Reference [2] also encourages suppliers to implement the *link set-to-key bin* method of congestion notification (which can be viewed as an extension of the link set-to-link set method). In terms of generic criteria, this means that elements of the CMP associated with the link set-to-link set technique are listed as "Requirements," while elements of the CMP associated with the link set-to-key bin technique are enumerated as "Objectives."

5. Summary

This paper presents the notification aspect of the proposed congestion management strategy for SMDS networks. SMDS Congestion Management uses a path-based scheme to map internal SS resource congestion to open interfaces (e.g., ISSI, ICI) and uses ECNs to notify other SSs. The generic elements of path-based schemes are explained. Four alternatives using path-based scheme are described and evaluated based on their usefulness in identifying congestion, effectiveness for congestion management and complexity at the ECN sending and ECN receiving SSs. Lastly, the paper provides justification for the *link set-to-link set* and the *link set-to-key bin* notification schemes proposed for congestion management in SMDS networks.

References

1. TR-TSV-000772, Generic System Requirements in Support of Switched Multi-megabit Data Service, Bellcore Technical Reference , May 1991.

2. TA-TSV-001059, *Generic Requirements for SMDS Networking*, Bellcore, Technical Advisory, Issue 2, August 1992.

3. TR-TSV-001060, *Switched Multi-megabit Data Service Generic Requirements for Exchange Access and Intercompany Serving Arrangements*, Bellcore Technical Reference, December 1991.

4. J. Moy, *OSPF Version 2*, Internet Engineering Task Force RFC 1247, July 1991.