ARCHITECTURE FOR TDM CIRCUIT EMULATION OVER IP IN TACTICAL NETWORKS

Øivind Kure
Q2S, NTNU O.S. Bragstads plass 2E, N-7491 Trondheim, Norway,
email: okure@item.ntnu.no
and
Ingvild Sorteberg
Baseline Communications as, Havegt. 2, N-2010 Strømmen, Norway,
email: ingvild.sorteberg@baseline.no
and
Knut Øvsthus
Forsvarets forskningsinstitutt (FFI), P.O.box 25, N-2027 Kjeller, Norway,
email: knut.ovsthus@ffi.no

ABSTRACT

TDM circuit emulation over IP (TDMoIP) may be used as a migration step towards a full IP solution. Several different TDMoIP architectures exist offering different degrees of robustness, service quality guarantees and management. The possible architectures are presented and their pros and cons are discussed.

Several tactical TDM based networks use flooding and pruning for routing and resource reservations. We propose a solution based on IP multicast that is efficient and requires no changes to the existing TDM signaling and mobility handling. The same signaling and routing scheme may be used for tactical Voice over IP (VoIP).

In TDM networks, resources are reserved during call setup. Priorities are invoked if there are insufficient resources available, and one or more lower priority calls are released. We discuss different QoS architectures and mechanisms that may be used to support the military priority scheme and the stringent delay and loss requirements of TDMoIP.

INTRODUCTION

Many military tactical networks are based on a TDM infrastructure (e.g. Eurocom) supporting telephony and data communications services. The TDM networks do not offer a very flexible transmission service for data and is difficult to adjust to changes in communication demands.

The motivations for IP based tactical networks are typically cost/performance, use of COTS technology and the desire to support the requirements of Network Centric Warfare (NCW). The main argument against using a pure IP infrastructure is that Voice over IP (VoIP) is still an immature technology and requires extensive investments both in the voice network infrastructure and new terminals, e.g. VoIP telephones. Also existing VoIP solutions do not support the military mobility, priority and robustness requirements. VoIP will eventually be the favored solution for carrying voice traffic, but until these requirements can be met a possible solution is to emulate the existing TDM services over an all-IP infrastructure. This requires small changes to the existing telephony services and capitalizes on existing investments and at the same time offers a flexible utilization of the network resources. TDM over IP (TDMoIP) is therefore a likely migration solution and may be deployed as part of IP based tactical networks.

There is ongoing work within IETF Working Group PWE3 (Pseudo Wire Emulation End-to-End) to standardize solutions for emulation of TDM over packet switched networks [1], [2] and [3]. The standardizations work has concentrated its effort on defining the packet encapsulation format for TDM signals. Little effort has been spent on aspects like call routing, signaling and QoS handling.

The contributions of this paper are in three areas, 1) evaluation of different TDMoIP architectures for tactical deployment, 2) the design of protocols for TDMoIP systems, where flooding is used in the propagation of TDM signaling messages and 3) a discussion of mechanisms supporting QoS and priority handling. The results of the last two areas are also relevant when introducing VoIP in military networks.

First the problems related to the use of TDMoIP in military networks are presented. Then four different architectures offering advantages and disadvantage with respect to support the military service requirements are discussed. A routing protocol mirroring the existing flooding based TDM routing is presented and finally the applicability of different QoS architectures is discussed based on the need to offer the required functionality and robustness.

PROBLEM STATEMENT

The main reason for deploying TDMoIP in military networks is that it offers a cost-effective migration solution towards a pure IP solution.

Commercial products are available offering encapsulation and de-capsulation of TDM trunks or individual channels. However, direct deployment of existing COTS TDMoIP solutions is not recommended. There are significant differences in the service requirements, for example in military networks there is a need to support per call priority levels and efficient re-affiliation solutions. Also there is a huge difference in the operational networking conditions for private enterprise networks and military tactical networks. In tactical networks, it can not be assumed that the network always offers sufficient bandwidth. Therefore, robust mechanisms are required to ensure fast response to congestion situations and avoid wasting network resources.

TDM OVER IP ARCHITECTURES

There are two main architectures, unstructured and structured TDMoIP. In addition, there are three alternative ways to emulate a structured TDM circuit. The different methods are presented and discussed based on their ability to support the requirements of military tactical networks. These requirements are: 1) end-to-end quality of service taking into account the need to support different levels of priority, 2) support for network and user mobility without requiring manual configuration of network elements and 3) efficient network utilization.

The TDMoIP functionality may be implemented in the TDM switch itself or as a separate adapter. We have not made any assumptions regarding the placement of the TDMoIP functionality and view this as an implementation choice that will depend on cost and whether existing TDM switches can be extended with this new functionality.

A. Unstructured TDM over IP

In unstructured TDMoIP, the TDM stream is transparently encapsulated and transported across the IP network, e.g. the IP network is viewed as point-to-point links. The main advantages are that the TDM circuit emulation functions may be implemented without any understanding of the TDM services and signaling and that it requires no changes to the TDM network nodes. Silent suppression may be supported to offer a better utilization of network resources. The disadvantages are that the solution introduces uncontrollable delay since the TDM signal must be packetized and de-packetized at every TDMoIP hop. Also there is no way to control that the TDMoIP packets are not sent over the same links several times resulting in low network utilization. In a large network, the total end-to-end delay may cause poor voice quality. Congestion in the IP network will affect entire trunks, causing degraded service quality for all calls including high priority calls. Therefore, an unstructured TDMoIP solution can not support the military priority levels and the architecture is not discussed further in this paper.

B. Structured TDM over IP

Structured mode refers to the case where TDM channels are individually transported across the IP network or grouped depending on their destination.

The advantage of the structured mode is that only active channels are transported across the IP network ensuring a very efficient utilization of network resources. TDMoIP packetization and de-packetization is only performed once, reducing the end-to-end delay considerably. Handling TDM channels individually allows channel management based on priority and quality of service requirements. The disadvantage is additional packet and processing overheads. If TDM adapters have stringent delay requirements, the number of small packets may become very large. Packet overhead can be reduced by grouping TDM channels towards the same destination into the same packet, irrespective of their priority. The disadvantage of this scheme is that more high priority traffic is generated. There are three alternative methods that might be implemented to support structured TDMoIP.

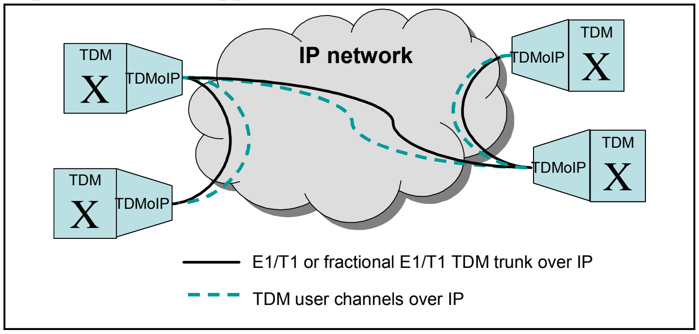

Figure 1: Alternative 1 - structured TDMoIP in trunk mode

In the first alternative, individual channels are transported in the same way as in the trunk mode, i.e. along the route determined by the TDM routing and packets are relayed at each TDMoIP adapter, see Figure 1. The advantages of this alternative are that the existing TDM routing can be utilized and the TDM samples only experiences one packetization delay. The disadvantage is that that the TDMoIP adapter must act like an application layer gateway supporting switching of TDMoIP packets. This will require that the TDMoIP adapter understands the TDM signaling and keeps a forwarding table with the mapping between the next TDM hop and IP hop for every TDM channel. This alternative may also result in larger delay due to many hops and less efficient utilization of the network resources since TDM samples are not necessarily transported along the shortest IP route.

In the second alternative, the signaling channel is set up according to the TDM trunk architecture, but the user channels are routed directly between the source and destination TDMoIP adapter, Figure 2. The advantages are that this minimizes the delay by not having to route the

TDM voice and data traffic via many TDMoIP adapters and it still ensures that existing location and mobility services are supported. There is no need to make extensive changes to the TDM signaling protocol. The disadvantages are that since the user and signaling channels are routed differently, network failures may cause the TDM service to experience an inconsistent network view. Also, the TDM signaling will impose a call admission control scheme based the link capacity seen from the signaling even though the IP network may have available resources, and vise versa the IP network may not have enough resources while the TDM signaling does not see this and admits new calls.

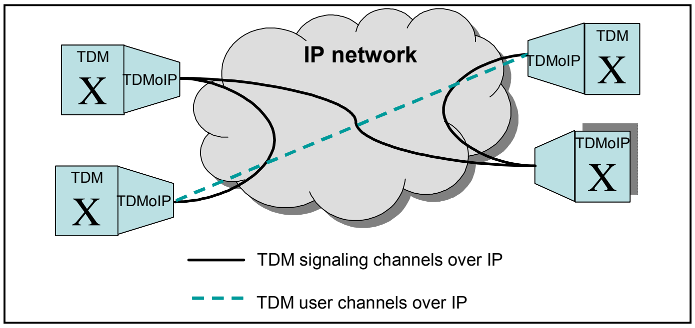

Figure 2: Alternative 2 - structured TDMoIP with trunked signaling and directly routed user channels.

In the third alternative, both the signaling and user channels are routed directly between the source and destination TDMoIP adapters. The advantages are that the TDM system takes full advantage of the IP least cost (shortest path) routing and fault management and traffic analysis is made easier. The disadvantage is that it does not directly support all existing services and requires more extensive changes to the signaling and routing to support location and re-affiliation of users. The rest of the paper discusses solutions for this alternative.

CALL ROUTING

Tactical TDM networks often utilize flooding and reversed spanning tree path for call routing. This allows for robust location and re-affiliation of terminals and users at the cost of overhead.

For TDMoIP, TDM addresses are mapped to IP addresses and routing of TDM signaling messages in the IP network is similar to the telephony routing experienced in VoIP systems. Three call routing solutions are presented; 1) TRIP (Telephony Routing in IP), 2) use of location/ directory servers, and 3) a multicast based solution.

A. Telephony Routing in IP (TRIP)

TRIP [4] is standardized by IETF and is a BGP4-based interdomain call routing protocol used to exchange telephone number ranges that are serviced by the different voice gateways or localization servers (LS). TRIP is intended for interconnection of VoIP and traditional telephony networks. In addition to telephony number ranges, TRIP supports exchange of policy and accounting information across administrative boundaries.

TRIP may be used to exchange information between TDMoIP adapters. This would include telephone number ranges supported and possibly TDM capacity information. The main disadvantage of using TRIP is that mobility handling is costly to support; every re-affiliation would require distribution of explicit routing updates to all TDMoIP adapters. The result is increased network overhead and large routing tables since aggregating of routing information becomes difficult. TRIP requires TCP sessions to be established and maintained. Depending on the number of TDMoIP adapters, substantial overhead may be introduced. The advantage of using TRIP is that it supports a smooth transition towards VoIP, assuming that VoIP routing will be based on TRIP.

B. Directory

In a directory approach, the mapping between IP address and TDM address/number is maintained by updating directories. When a user moves or a TDMoIP adapter changes its IP address, the directory is updated. A user that re-affiliates must register at the local directory and the information is either distributed to all directory servers or to a master directory server if a hierarchical directory structure is supported. When establishing a new call, the originating TDMoIP adapter queries its local directory for the IP address of the TDMoIP adapter serving the terminating subscriber. If the local directory holds the information, it is given to the originating TDMoIP adapter directly. If not, the local directory must send a request either to the master or to all directory servers.

The disadvantage of this solution is that it is costly to provide a reliable and robust directory service. Backup solutions must be implemented in case of server failures and robust and efficient replication schemes are required to support mobility. Replication may result in large network overhead, particularly in the case of many mobile users.

We do not recommend a directory based location service since it requires not only the introduction of a directory system and its related protocols, but also additional functions in the TDMoIP adapters to update the directory and process requests. Finally, we also believe that this solution is less robust than our recommended solution.

C. Multicast based call routing

We recommend a TDMoIP call routing mirroring the existing TDM call routing in order to support the exiting TDM mobility scheme. The underlying idea is to emulate the TDM spanning tree search in the IP network, combined with caching of hints for the location. This can be done using IP multicast. The strength is that each TDMoIP adapter does not need know the address of the other TDMoIP adapters in the IP network. All TDMoIP adapters within an administratively defined area are members of the same multicast group. Standard multicast routing protocols are used to establish a distribution tree.

When a new call is initiated, the originating TDMoIP adapter checks if it has an entry for this telephone number in its routing table. The call routing table holds a list of phone numbers, their associated IP address, outgoing interface, age, and flags for additional information. If a routing entry does not exist, has timed out, or receives no response, the call setup message is sent to the IP multicast address and is routed along the multicast tree. The call setup messages will be processed by all members of the multicast group, and the TDMoIP adapter serving the requested telephone number will respond. The exact procedure depends on the size and structure of the network.

In networks where the source and destination always have direct IP connectivity, the response message contains the terminating TDMoIP adapter's IP unicast address. This address is then added to the originating TDMoIP adapter's routing table. Multicasting is then essentially an effective address resolution mechanism.

In large multi-domain networks, where routing through TDMoIP adapters or call servers is required, the procedure is more complex. The TDMoIP adapters may be members of more than one multicast group offering a bridging function. In this case, an identifier is required to uniquely identify a call setup message. This ensures that the same setup message is not introduced into the same multicast group or the same network several times. Typically, the originator and a local sequence number are used as identifier.

When the TDMoIP adapter receives a call setup message, it checks whether a valid entry already exists in the forward ongoing path table. If not, the address of the source TDMoIP adapter and the identifier are stored in the reversed outgoing path table. The call message is then forwarded to the multicast address for the new area, after the reversed outgoing path table has been checked to avoid flooding loops. This is repeated by every TDMoIP adapter each time area borders are crossed. Eventually the correct TDMoIP adapter receives the call setup message and a response is generated.

The format of the signaling packets must be augmented with a path metric and the address of the forwarding TDMoIP adapter or call server. As multicast signaling packets are received, the outgoing reversed path table must be recalculated to identify the optimal previous hop. A TDMoIP adapter can potentially receive multiple copies of a signaling message, since there can be many potential forwarders. The first packet will be forwarded after a configurable hold time; the rest will only update the reversed tree. The path cost in the signaling message will not necessarily be correct since there is always a probability that a shorter path to the source has been discovered after the message was forwarded. The holding time before the message is forwarded at the boundary between two areas represents a trade-off between delay in the signaling message and the accuracy of the cost metric used. The call setup response message follows the reverse spanning tree generated in the forwarding process.

The preferred multicast protocol is primarily a function of whether all routers are multicast enabled, and the density of TDMoIP adapters to routers. Due to the relatively low amount of signaling traffic and an equal number of sender and receivers, a shared tree will be the optimal choice for distribution tree. Multicasting can use scooping to avoid multicast requests to parts of the network where the bandwidth is very limited. The robustness of the solution will depend on the choice of multicast routing protocol.

QUALITY OF SERVICE ARCHITECTURES

In existing tactical TDM networks, routing of the signaling messages and routing of the call follows the same path. This assures that resources are available at each TDM hop. Military TDM services also support three or four levels of priority. To support this functionality, the IP network must be able to pre-empt low priority calls and assure that high priority calls are granted sufficient resources. This section discusses different methods to guarantee the quality of service (QoS) for the TDM traffic in an IP network.

A. IntServ and DiffServ

IETF has standardized two methods for supporting QoS in IP networks. The Integrated Services (IntServ) architecture [5] offers per-flow resource reservation and admission control, and uses RSVP (Resource reSerVation Protocol) for signaling [6]. RSVP may be extended to support pre-emption and can be used to handle the military priority scheme. The main disadvantage is that every network node is required to hold per-flow state information. This may lead to scalability problems when handling many flows and is the main reason for the lack of commercial interest in IntServ. Another problem with IntServ is that RSVP messages can not pass through IPsec devices directly, but requires functionality as specified in [11].

The Differentiated Services (DiffServ) architecture [7] was defined as a response to the IntServ scalability problems. It is based on flow aggregation. IP packets requiring the same treatment are marked using the Differentiated

Services Code Point (DSCP), a 6-bit encoded field of the DiffServ (DS) byte in the IP header [8]. DiffServ capable routers implement packet forwarding behaviors, Per Hop Behaviors (PHB), for distinct traffic types based on the DSCP-value. Currently two PHB are standardized, Expedite Forwarding (EF) and Assured Forwarding (AF) [9]. The latter supports four classes and within each class packets can be assigned three different levels of drop precedence.

The routers realize the different PHB using Active Queue Management (AQM), a combination of packet scheduling and discard strategies. The discard strategies typically use random packet discard, e.g. Random Early Detection (RED), spreading the packet loss to several flows in order to shed load during periods of congestion.

B. Explicit Congestion Notification (ECN)

ECN [10] is a mechanism used to control TCP flows and might be used either alone or together with DiffServ. ECN offers a method where routers when detecting increasing traffic load marks packets by setting two of the bits in the TOS/DS field to notify the end system about the network congestion. The end system acts upon this by requesting the TCP sender to slow down. If used with TDMoIP, the TDMoIP adapter must be able to read the ECN field and initiate a call termination or block all new calls. ECN may impose a problem with IP security, since it requires that the TOS/DS field is carried transparently across security boundaries. The main advantage of using ECN is that it offers a possibility to notify the end systems of network congestion before packet loss occurs.

PROPOSED QOS MECHANISMS

Since the TDM system has no control of the routing and traffic load in the IP network, mechanisms are needed to ensure that the TDMoIP system responds by reducing the input load and preventing additional calls to be established when the network resources are exhausted.

A. Congestion control mechanism

A standard AF-based DiffServ solution is not suitable since the existing drop mechanisms (e.g. RED) tends to spread the packet loss to the maximum number of flows. Depending on how TDM reacts to packet loss, the result is either a prolonged period of reduced subjective quality for many calls or the termination of a larger than required number of TDM flows. For similar reasons a pure tail-drop mechanism is not an alternative. If the synchronicity of TDM traffic was preserved in the network, a tail-drop marking scheme would focus on the same sequence of packets ensuring fast shutdown of the offending flows. As the routers disturb the relative sequence, the likelihood is that a larger number of flows will be affected.

We propose a DiffServ based architecture with a modified queue management. The Assured Forwarding PHB is used, and the call priority is mapped to an AF drop precedence. The probabilistic drop mechanism is replaced with a combination of tail-drop and ECN marking of a fixed number of packets. The TDMoIP adapter reacts to the ECN marking by closing down the connections. This ensures an adaptive load control mechanism.

There is a delay (Round Trip Time + processing) from the first marking until the queue experiences a reduced load. If the ECN marking is continued in this period, a larger number of calls than necessary will have been marked for termination. By marking only a fixed number of packets, we bound the number of flows that are terminated. After the given number of packets has been marked, the queue discipline reverts to tail-dropping where packets are discarded if the queue length is above the watermark for the drop precedence. Tail-drop is activated for a configurable period of time, determined by the Round Trip Time, to protect the downstream routers from further congestion. The ECN mechanism is rearmed after a configurable time-out.

The network operator can by changing the number of marked packets; make a trade-off between responsiveness to overload and utilization of the network.

Pre-emption of traffic flows is ensured by mapping the TDM priority levels to the drop precedence levels of the AF-class. If a link receives too much traffic, packets belonging to the low priority class are candidates for being ECN marked and dropped first. This mechanism ensures that higher priority calls going through the congested nodes are shielded.

B. Admission control

One inherent problem in the use of DiffServ with AQM and our proposed solution with ECN marking is the lack of stability during severe congestion. In TDM, admission control ensures that new calls will be denied while existing calls are untouched. With the various DiffServ schemes, existing and new calls have the same probability of being terminated, since flow termination is based on random packet drop or marking.

From a fairness view point this is acceptable; long lasting calls should not always be protected. However, there is a substantial difference in the stability. Protecting existing calls ensures that at least established calls are allowed to progress and terminate naturally. Under severe congestion randomly dropping one of the existing calls creates thrashing where nothing is accomplished. Therefore, the proposal must be augmented with an admission control mechanism.

The objective of call admission should be to avoid thrashing. The simplest solution will be to add a trigger function with a refresh timer for blocking new calls in the TDMoIP adapters. Once an adapter experiences a drop frequency above a limit for a given priority class, it will block all new calls for that class. It will remain in blocking state as long as the drop frequency is above the limit plus a configurable interval before it starts accepting new calls again. This is a fair and distributed admission control scheme that will be used only during severe congestion. It is timer based to ensure that an adapter will not remain blocked, while other adapters create more traffic on a congested path. It will only affect adapters that generate traffic on congested paths. The disadvantage is the temporal uneven distribution of blocked adapters; some may be blocked while others continue to accept traffic. However, eventually all will cycle through the blocked state or the network will reach a new stable operating point. This proposed scheme will not be affected by the introduction of IP security. This is the sole advantage compared to an IntServ solution.

FUTURE WORK

To evaluate our QoS scheme, we have initiated a simulation activity to determine the effectiveness of the admission control scheme and packet loss both during limited network congestions and during massive re-routing causing extensive congestions. The initial simulations confirmed the unsuitability of RED during congestion for TDM calls, while our proposed algorithms were able to limit congestion in a controlled fashion. However, under severe congestion there is still some risk that a high priority packet is lost since buffers are not purged of low priority packets. We would also like to compare our solution with a pure IntServ/RSVP based solution looking at blocking probability and robustness.

CONCLUSIONS

Based on our analysis we recommend that a TDMoIP solution for military tactical networks is based on the direct routing of individual TDM channels. This allows support for military priority as well as offers the most efficient utilization of resources and simplifies network management since the routing of TDMoIP calls are based on the IP routing.

It is not possible to map a particular TDM address to an associated IP address. Existing IETF proposals for VoIP routing, e.g. TRIP (Telephony Routing in IP), are not applicable. Instead, we propose a call routing protocol based on IP multicast that is efficient and requires no changes to the existing TDM signaling. The call setup messages are efficiently flooded to a multicast group with a minimum of state information required. Additionally, a hint based caching scheme can be deployed to minimize delay and overhead even further.

Due to low bandwidth and the need for efficient utilization of resources in military tactical networks, TDMoIP requires support for congestion and call admission control. We recommend a solution based on a combination of ECN and packet dropping. ECN is used to signal congestion while packet drops are monitored in the TDMoIP adapters and if it increases above a defined threshold, new calls are blocked unless they have higher priority than the existing ongoing calls.

ACKNOWLEDGEMENTS

We would like to thanks our colleges at Thales Communications AS in Norway for valuable input to the discussions of some of these topics.

REFERENCES

[1] A. Vainshtein and Y.Stein (Ed.), *Unstructured TDM Circuit Emulation Service over Packet Switched Network (UCESoPSN)*, Internet Draft <draft-vainshtein-pwe3-ucesopsn-00.txt>, Exp: March 2003.

[2] Y. Stein, R. Shaashoua, R. Insler, M. Anavi, *TDM over IP*, Internet Draft <draft-anavi-tdmoip-05.txt>, Exp: September 2003.

[3] T. Johnson, *TDM Circuit Emulation over Packet (CEP-TDM)*, Internet Draft <draft-johnson-pwe3-tdm-00.txt>, Exp: August 2003.

[4] J. Rosenberg, H. Salama, M. Squire, *Telephony Routing over IP (TRIP)*, RFC 3219 (2002).

[5] R. Braden, D. Clark, S. Shenker, *Integrated Services in the Internet Architecture*, RFC 1633 (1994).

[6] R. Braden, L. Zhang, S. Berson. S. Herzog, S. Jamin, *Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification*, RFC 2205 (1997).

[7] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, *An Architecture for Differentiated Services*, RFC 2475 (1998).

[8] K. Nichols, et al., *Definition of the Differentiated Services Field (DS-Field) in the IPv4 and IPv6 Headers*, RFC 2474 (1998).

[9] J. Heinanen, et al., *Assured Forwarding PHB Group*, RFC 2697 (1999).

[10] K. Ramakrishnan, S. Floyd, D. Black, *The Addition of Explicit Congestion Notification (ECN) to IP*, RFC 3168 (2001).

[11] Terzis, J. Krawczyk, J. Wroclawski, L. Zhang, *RSVP Operation Over IP Tunnels*, RFC: 2746 (2000).