high speed local area networks

edited by
o.spaniol and a.danthine

IFIP

north-holland

HIGH SPEED LOCAL AREA NETWORKS

Proceedings of the IFIP TC 6/WG 6.4 International Workshop on
High Speed Local Area Networks (HSLAN)
Aachen, Federal Republic of Germany, 16-17 February, 1987

edited by

Otto SPANIOL
Lehrstuhl Informatik IV
RWTH Aachen
Aachen, FRG

André DANTHINE
Institut d'Electricité Montefiore
Université de Liège
Liège, Belgium

1987

NORTH-HOLLAND
AMSTERDAM · NEW YORK · OXFORD · TOKYO

ISBN: 0 444 70324 1

Published by:
ELSEVIER SCIENCE PUBLISHERS B.V.
P.O.Box 1991
1000 BZ Amsterdam
The Netherlands

Sole distributors for the U.S.A. and Canada:
ELSEVIER SCIENCE PUBLISHING COMPANY, INC.
52 Vanderbilt Avenue
New York, N.Y. 10017
U.S.A.

Library of Congress Cataloging-in-Publication Data

IFIP TC 6/WG 6.4 International Workshop on High Speed Local Area Networks (1987 : Aachen, Germany)
High speed local area networks : proceedings of the IFIP TC 6/WG 6.4 International Workshop on High Speed Local Area Networks (HSLAN) Aachen, Federal Republic of Germany, 16-17 February, 1987 / edited by Otto Spaniol, André Danthine.
p. cm.
Sponsored by the IFIP WG 6.4.
Bibliography: p.
ISBN 0-444-70324-1 (U.S.)
1. Local area networks (Computer networks)--Congresses.
I. Spaniol, Otto, 1942- II. Danthine, A. III. IFIP Working Group 6.4. IV. Title.
TK5105.7.I37 1987
004.6'8--dc19

87-23269
CIP

PRINTED IN THE NETHERLANDS

PREFACE

High Speed Local Area Networks are defined to be communications systems which offer a comparatively high raw data capacity. The lower end of the spectrum is 50 or 100-Megabit-per-second (with increasing tendency) whereas the upper limit is in the Gigabit-per-second range. Several projects or even prototypes deal with systems offering 1.2 Gigabit-per-second [reports about such networks are included in the present book]. For the overwhelming majority of computer network users, however, such announcements are dreams or science fiction since very often the user has to live, even today, with digital lines providing 9.6 kilobit or even less. Thus the HSLAN capacity would result in an improvement of 4 to 5 orders of magnitude.

The installation of HSLAN will concentrate in relatively small geographical areas but in contrast to many classical local area networks the diameter of coverage may be as high as, say, 25 km. This enlarged distance will be required for one of the major applications for HSLAN, namely a highly performant interconnection of local area subnetworks. These LANs are deliberately heterogeneous since they are tailored for specific applications with different traffic scenarios and workloads. Thus HSLAN may and probably will offer a decisive step towards a really open communication which is, moreover, flexible enough to offer the best subnetwork solutions with respect to any given application.

High Speed Local Area Networks have been commercially available for several years in the 50 Megabit/sec range. Due to the enormous progress in the optical transmission domain it is to be expected that in the very near future other products will show up on the market which could also be offered at more reasonable prices. Several research groups in the US, Japan and Europe have entered a very interesting competition with respect to different HSLAN concepts. Within the ESPRIT program of the European Community at least three HSLAN approaches are being developed simultaneously.

At the moment there is not yet a breakthrough for a commonly agreed standard even if FDDI may be considered as a first step in this direction. The situation is similar to many other domains where concurrent approaches have to be tested before a concen-

tration to the (hopefully) most suitable systems together with a de-facto-standardization may be envisaged.

The IFIP WG 6.4 workshop held in Aachen (FRG) on February 16–17, 1987, intended to present the state-of-the-art within the HSLAN area as well as to give directions for the solutions of unsolved problems. Some of the most important questions were:

- Which applications for HSLAN?
- When will products appear on the market for a more reasonable price?
- Which performance is required for connecting stations, i.e. will the limited terminal capacity result in serious bottlenecks (like in some established lower speed LAN)?
- Are HSLAN the adequate solution for LAN subnetwork interconnection?
- Will interconnected HSLAN extend to metropolitan area networks (MAN)?
- Will wideband ISDN and its associated PBX make HSLAN concepts obsolete?

The latter questions could not be answered satisfactorily during the workshop since e.g. information about wideband ISDN was very scarce. However, it seems that HSLAN will remain one of the most important topics in the networking area independently of ISDN progress. There will be a specific market segment for each of the concepts. The other major research questions have been treated in the following sessions:

- HSLAN Technology
- Simulation and Performance
- Network Management
- Services and Integration
- HSLAN Prototypes.

It turns out that enormous progress has been made in concept development and in understanding specific problems arising in the HSLAN environment. Performance aspects have become much clearer, too. The consensus was that future needs for research and development in the HSLAN field should move from the introduction of new schemes and access methods towards the important feature of getting experience with pilot installations in the real environment. Furthermore, upcoming standards should be carefully investigated.

The relevance of the information which was presented at the workshop and which is contained in the present proceedings volume is extremely high. This is due to the fact that papers could be submitted just one month before the conference. The organisation of the event in such a short time would have been absolutely impossible without the enormous engagement of Peter Martini and Thomas Welzel

(both from the Technical University of Aachen). Their work has been very much appreciated by all the workshop participants as well as by the editors of the present book.

A. Danthine
O. Spaniol

ORGANIZING COMMITTEE

André Danthine
University of Liège, Belgium

Peter Martini
Technical University of Aachen, Federal Republic of Germany

Otto Spaniol
Technical University of Aachen, Federal Republic of Germany

Thomas Welzel
Technical University of Aachen, Federal Republic of Germany

CONTENTS

High Speed Local Area Networks
O. Spaniol and A. Danthine (Editors)
Elsevier Science Publishers B.V. (North-Holland)
IFIP, 1987

SYNCHRONISATION SCHEMES FOR A HIGH SPEED OPTICAL FIBRE TDM RING NETWORK

P.A. Davies, A. Jamasebi-Jahromi

Electronic Engineering Laboratories, The University of Kent, Canterbury, Kent, CT2 7NT.

INTRODUCTION

In this paper we shall discuss some of the options available for achieving clock synchronisation in a High Speed Optical Fibre Ring Network which is currently being designed at the University of Kent.

We will begin by describing briefly, the structure of the network that we intend to build and then we will go on to consider in detail, the various schemes for achieving synchronisation and our reasons for selecting a particular method.

THE UKC FIBRE RING

As part of a long term interest in fibre optic local area networks (LANs) we are currently involved in the design phase of a 1.2 Gbps fibre optic ring network which will use TDMA as a multiple access technique. It is not our intention in this paper to argue the case for any given protocol but we will give briefly our reasons for choosing this combination of technologies. Firstly we have chosen an operating channel rate compatible with developments in long haul telecommunications in order to be able to utilise, as far as possible, off-the-shelf components. We have chosen a ring topology because we wish to use point-to-point links. We have selected TDM as being a well developed technique relatively easy to implement at Gbps rates with off-the-shelf components and capable of supporting a wide range of services including voice, data and video. It is important to note that a major aim in the work is to use commercially available components as far as possible in order to produce a practical solution to the problem of integrating services on a high speed channel.

Our prototype system will use 8 channel mux and demux chips produced by gigabit logic. In the later stages we may increase the number of channels. Initially, one of the 8 channels will be used for synchronisation but subject to satisfactory performance we intend to add data to the synchronisation channel at a later stage.

The ring circumference will be of the order of 1 Km or more. At a transmission rate of 1.2 Gbps approximately 750 frames each of 8x144 Mbps channels can be stored on a 1 Km ring. This means that we can achieve an integral number of frames on the network by introducing clock skew of less than 0.2%. This removes the requirement for high speed buffering at the nodes but does require that our clock synchronisation circuitry will accommodate this clock

skew. The proposed system is shown in figure 1.

In the following sections we will examine the various clock recovery schemes available. We will examine optimum and sub-optimum schemes and will indicate how the constraints of working at high bit rate may influence our decision to use a particular scheme.

OPTIMUM CLOCK SYNCHRONISATION TECHNIQUES

The optimum clock recovery procedure is generally acknowledged to involve the "maximum a posteriori" (MAP) estimator.[1] Essentially this technique is based on a maximum likelihood detection. The received signal is observed over K symbols during which the timing misalignment, ε, between the local clock and the incoming data stream is then calculated. The scheme is rarely used in practice and is mainly used for comparison to assess the performance of other synchronisation techniques.

Various sub-optimum synchronisation schemes motivated by the MAP approach have been proposed but all tend to involve the use of integrate and dump circuits. Also the MAP synchroniser is an open loop system and cannot adapt to slowly changing ε. In practice, a close loop form of synchroniser is required because the clock rate of the incoming signal may vary slowly with time. An example of the kind of closed loop synchroniser motivated by the map technique is shown in figure 2. This is known as an "early/late gate synchroniser" and is based on comparing the phase of the output of two integrate and dump circuits. Clearly the implementation of this type of circuit is likely to be complex and not well suited to Gbps operation.

NON OPTIMUM CLOCK SYNCHRONISATION TECHNIQUES

Essentially the optimum and sub-optimum techniques use sophisticated phase detection methods to compare the phase and frequency of the local clock to the clock component of the incoming signal. Our aim is to generate the phase/frequency error signal for Gbps data with off-the-shelf components and hence the phase/frequency detector must be in the simplest possible form. Suitable digital phase/frequency detectors which are available to us include exclusive OR gates and RS flip-flops.

Of course, the easiest way to extract clock from incoming data is to use a filter, commonly a tuned circuit or phase locked loop, to extract the clock frequency directly form the incoming signal. However, as is well known, NRZ type data does not contain any frequency component at the clock frequency. There are various solutions to this problem and they include:

a/ The addition of a line code
b/ Conversion of NRZ data to RTZ data
c/ Use of a non linearity at the receiver
d/ Clock sub-harmonic synchronisation
e/ Clock signal derived from frame synch pulses

The first of these options, (a) the addition of a line code, is not suitable at Gpbs rates for two reasons. Firstly, this will result in an increase in bandwidth producing greater difficulties at both transmitter and receiver. Secondly, the equipment required to generate the line code is likely to be complex and costly even if the code is generated and added to

the sub-channels at 144 Mbps rather than at 1.2. Gbps.[2]

Conversion of NRZ data to RTZ data, as in (b) above, can either be performed at the transmitter or at the receiver. If performed at the transmitter the resulting increase in bandwidth pushes the data rate beyond the capacity of the chips we wish to use. NRZ to RTZ conversion at the receiver will be considered below.

The use of a non linearity at the receiver, (c), is a common technique. Some typical methods for implementing the non linearity include

1/ Square law device
2/ Full wave rectifier
3/ Half bit delay (NRZ to RTZ conversion)
4/ Limiter
5/ Differentiate and square

Block diagrams of some of these methods are shown in figure 3.

Of course the circuitry used to produced the non linearity chosen will need to be capable of operating at the clock frequency and while this may present some design difficulties careful design should eliminate most of the problems and clock extraction based on, for example, half bit delay is a possible candidate.[2] It should be noted though, that all of the "non linearity" techniques rely on processing the transitions in the data. Hence, if sufficient transitions are not present in the original data extra transitions must be introduced by, for example, pulse stuffing or by scrambling the data, neither of which solutions are ideal.

An alternative synchronisation technique which is occasionally used is to lock the phase of the local clock to a sub-harmonic of the clock which is present in the incoming data. Figure 4 shows how this might be achieved. This is a particularly attractive technique for a Gbps system since only the VCO and divide by N network need operate at the clock frequency.

Finally, we should consider the possibility of extracting the clock from the frame synchronisation pulses which are to be transmitted on one of our eight channesl. If a synchronisation sequence of alternating 1's and 0's is transmitted on the synchronisation channel then effectively a clock sub-harmonic at 1/16 of the clock rate is available independent of any data pattern occurring on the data channels.

CLOCK SYNCHRONISATION SCHEMES FOR A 1.2 Gbps RING LAN

Bearing in mind our requirement that we use off-the-shelf components wherever possible, and the resulting restrictions on data rates that such a requirement brings, we feel that our most likely candidates for synchronisation methods are (d) Clock sub-harmonic (Ck/2) synchronisation and (e) Synchronisation to TDM frame pulses (Ck/8 synch). We will now consider these two possibilities in detail.

Clock synchronisation by locking to the TDM frame pulses has some strong arguments in its favour. Firstly, as explained above, a simple frame synchronisation sequence of alternating ones and zeros can guarantee that a transition will occur every 16 data pulses and hence the synchronisation scheme will operate independently of the transmitted data pattern. Of course, such a system will not utilise any data transitions which occur on the data channels between the frame synch pulses and so the system will not be quite

so sensitive to variations in clock phase and frequency as would a more conventional clock extraction scheme. However, it is possible to compensate for this decrease in sensitivity to some extent by increasing the gain of the phase lock loop. Techniques similar to this have been employed with success at channel rates up to 5 Gbps.[3] However, they do involve an extra level of control and a sophistication in hardware built in house that we would prefer to avoid if possible.

Turning now to clock sub-harmonic (Ck/2) synchronisation we arrive at a particularly simple solution to the problems of high speed clock extraction, which will allow us to perform the functions of clock synchronisation and frame synchronisation independently resulting, we would expect, in a more robust system requiring simple control.

The sub-harmonic synchroniser shown in figure 4 includes a bandpass filter. This filer is usually included to prevent the phase lock loop locking to any unwanted sub-harmonic. Of course, the filter could be implemented rather elegantly at Gbps rates in microstrip directly on the printed circuit board carrying the clock extraction circuit. However, in the introduction to the paper we explained that a clock skew of ±0.2% would be required. Since the circuit of figure 4 will tend to operate at the resonant frequency of the filter rather than at the incoming clock rate and as the filter is a limiting factor on a flexible data rate we may remove this filter.

This scheme is, of course, dependent on transitions in the data for correct operation and a long stream of data 1's or data 0's on the channel would severely affect the synchroniser's performance. The data can be randomised to some extent by using a scrambler but this does not give a complete solution to the problem. The scrambling and descrambling should be done at the sub-channel (144 Mbps) rate and not at the transmission rate (1.2 Gbps) where the hardware problems would be severe. Of course, if scrambling is performed on the sub-channels then we must realise that a finite possibility exists that the scrambled channels may combine when multiplexed to form a long stream of 1's or 0's. Ideally, we would like to place a bound on the maximum run of 1's or 0's that can occur on the channel. Fortunately, we can produce this bound if we arrange the synchronisation sequence to be composed of alternating 1's and 0's. This will guarantee a data transition at least every 16 bits.

A block diagram of our chosen clock synchronisation and frame synchronisation system based on the above arguments is shown in figures 4 and 5. This scheme has the following advantages:

1. A minimal number of high speed elements is required. The required elements are straightforward and readily available.

2. Runs of 1's and 0's are bounded by the frame synchronisation sequence.

3. The frame synch sequence is straightforward and frame synchronisation is easy to implement.

4. No control algorithms are required to achieve clock and frame synchronisation .

CONCLUSION

We have looked in some detail at synchronisation techniques for a high speed (Gbps) TDM ring network. We have considered optimum and sub-optimum

structures. We have identified several schemes which would be suitable and have chosen a sub-harmonic (Ck/2) clock extraction scheme combined with a specific frame synchronisation sequence to satisfy our criterion for using readily available off-the-shelf components.

This synchronisation circuit has been simulated at low bit rates and the high speed version (1.2 Gbps) is currently under construction.

ACKNOWLEDGEMENTS

We gratefully acknowledge the support of the Science and Engineering Council of the UK, under whose sponsorship this work is currently being performed.

REFERENCES

1 Lindsey W.C. & Simon M.K.: Telecommunication Systems Engineering, Prentice Hall, 1973.

2 El-Fatatry A.: High Speed Optical Data Transmission over Monomode Fibre at 1300 nm, Ph.D Thesis, University of Kent, 1986.

3 Bergman L.A. & Eng. S.T.: A Synchronous Fibre Optic Ring Local Area Network for Multigigabit/s Mixed Traffic Communications, SAC-3, 6, November 1985, pp. 842-848.

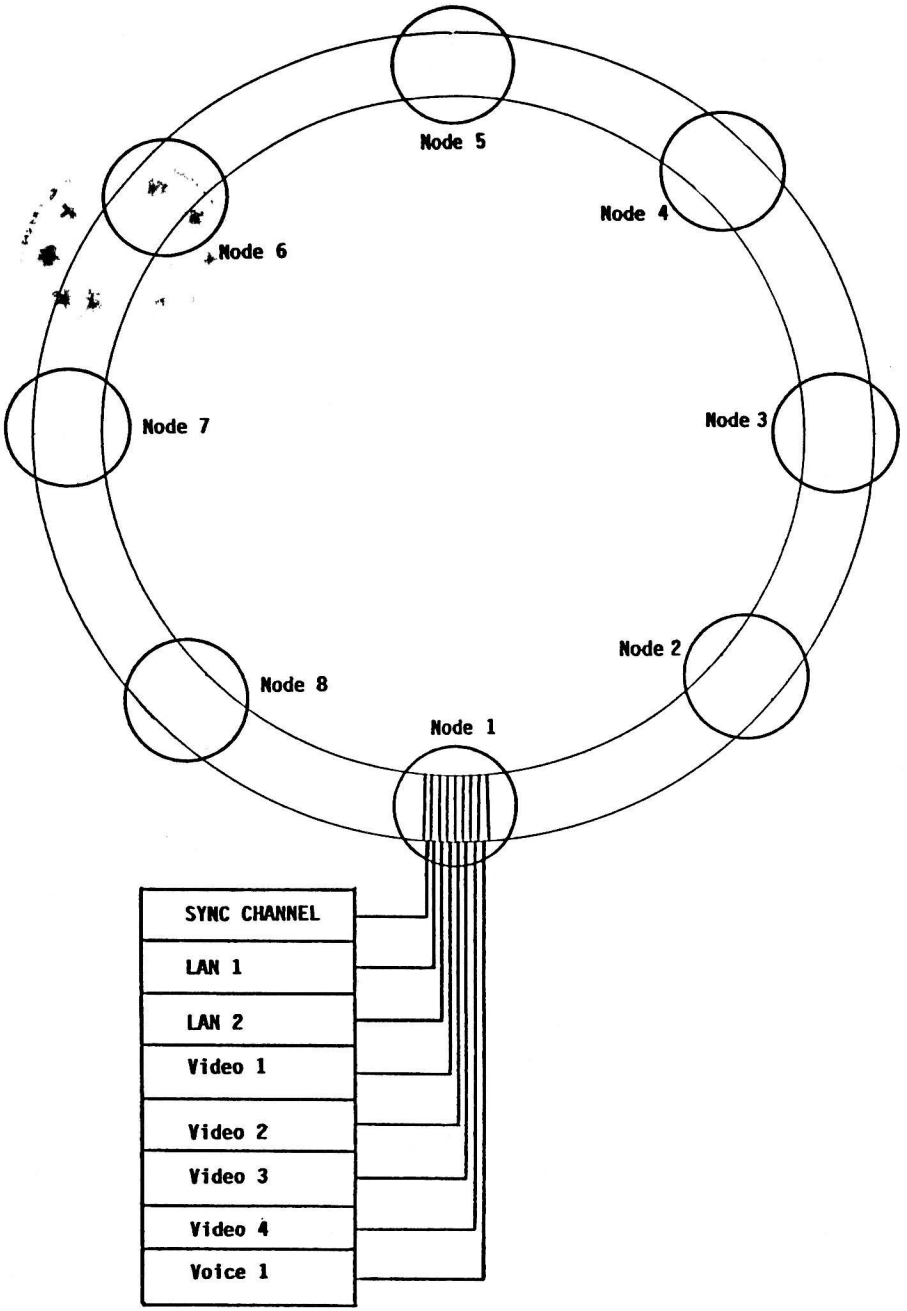

Figure 1 : A high speed optical fibre TDM ring network

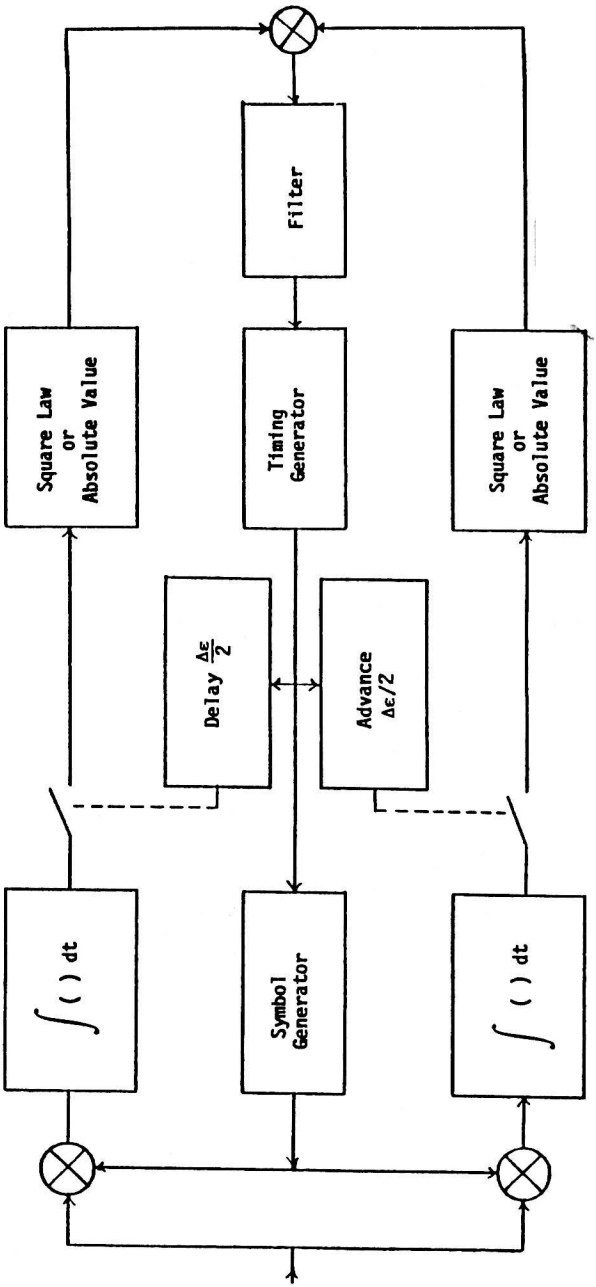

Figure 2 : An early/late gate synchroniser