

Dependable communication synthesis for distributed embedded systems

Nagarajan Kandasamy^{a,*}, John P. Hayes^b, Brian T. Murray^c

^a*Institute for Software Integrated Systems, Vanderbilt University, P.O. Box 1829, Station B, Nashville, TN 37235, USA*

^b*Advanced Computer Architecture Laboratory, University of Michigan, 1301 Beal Avenue, Ann Arbor, MI 48109, USA*

^c*Brighton Technical Center, The Delphi Corporation, 12501 Grand River Avenue, Brighton, MI 48116, USA*

Available online 6 October 2004

Abstract

Embedded control applications such as drive-by-wire in cars require dependable interaction between various sensors, processors, and actuators. This paper addresses the design of low-cost communication networks guaranteeing to meet both the performance and fault-tolerance requirements of such distributed applications. We develop a fault-tolerant allocation and scheduling method, which maps messages onto a low-cost multiple-bus system to ensure predictable inter-processor communication. The proposed method targets time-division multiple access (TDMA) communication protocols, and is applicable to protocols such as FlexRay and TTP which have recently emerged as possible networking standards for embedded systems such as automobile controllers. Finally, we present a case study involving some advanced automotive control applications to show that our approach uses the available network bandwidth efficiently to guarantee message deadlines.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Communication synthesis; Embedded systems; Distributed systems; TDMA protocols

1. Introduction

Embedded computer systems are being increasingly used in cost-sensitive consumer products such as automobiles to replace safety-critical mechanical and hydraulic systems. Drive-by-wire is one example where traditional hydraulic steering and braking are replaced by a networked micro-processor-controlled electro-mechanical system [1]. Sensors measure the steering-wheel angle and brake-pedal position, and processors calculate the desired road-wheel and braking parameters, which are then applied via electro-mechanical actuators at the wheels. Other computerized vehicle-control applications including adaptive cruise control, collision avoidance, and autonomous driving are also being developed [2]. These applications will be realized as *real-time distributed systems* requiring dependable interaction between sensors, processors, and actuators. This paper addresses the design of low-cost communication

networks to meet both the performance and fault-tolerance requirements of such applications.

Related work in communication synthesis for distributed embedded systems belongs to two broad categories—those that assume a fixed network topology and schedule messages to meet their deadlines [3–5], and those that synthesize a topology satisfying message deadlines [6,8].

The authors of [3] assume a fixed network topology and using the controller area network (CAN) protocol, schedule messages by assigning appropriate priorities to help meet their deadlines. Off-line algorithms to schedule both tasks and messages in combined fashion while minimizing overall schedule length are developed in [4,5]. In [7], processors are assigned priorities for network access and the corresponding messages are transmitted using a fixed-priority scheduling scheme.

A network topology satisfying message deadlines can also be constructed from application requirements. Modeling embedded applications as task graphs, [6] estimates the communication delay for inter-processor messages and schedules them on the minimum number of buses using the CAN protocol, while [8] generates point-to-point communication links for messages.

* Corresponding author. Tel.: +1 615 513 1865; fax: +1 615 343 7440.

E-mail addresses: nagarajan.kandasamy@vanderbilt.edu (N. Kandasamy), jhayes@eecs.umich.edu (J.P. Hayes), brian.t.murray@delphi.com (B.T. Murray).

Unlike [3–5,7] which assume a given topology, the approach proposed in this paper synthesizes a fault-tolerant network topology from application requirements. While synthesis methods such as [6] assume an underlying CAN communication protocol and arbitrate bus access using message (processor) priorities, we target time-division multiple access (TDMA) communication protocols where processors are allotted transmission slots according to a static, periodic, and global communication schedule [9]. Also, TDMA protocols such as TTP [10] and FlexRay [11] have recently emerged as possible networking standards for an important class of embedded systems-automobiles.

Rather than generate arbitrary networks, we restrict the topology space to multiple-bus systems. Fig. 1 shows an example where each processor P_i connects to a subset of the communication buses. A co-processor handles message communication independently without interfering with task execution on P_i . A multiple-bus topology allows fault-tolerant message allocation. Also, since communication protocols for the embedded systems of interest are typically implemented over low-cost physical media, individual buses have limited bandwidth; multiple buses may be needed to accommodate the message load.

Given a set of distributed applications modeled as task graphs $\{G_i\}$, our approach constructs a low-cost communication network satisfying both the performance and fault-tolerance requirements of each G_i . Messages are allocated and scheduled on the minimum number of buses $\{B_j\}$ where each B_j has a specified bandwidth. We now summarize the major features of our approach:

- It assumes a multi-rate system where each graph G_i may have a different execution period, $\text{period}(G_i)$.
- It targets a generic TDMA communication protocol.
- It supports dependable message communication by establishing redundant transmission paths between processors, thereby tolerating a bounded number of permanent bus failures.
- It uses network bandwidth efficiently by reusing transmission slots allotted to a processor between the multiple messages sent by it.

Finally, using some representative automotive control applications, we show that the proposed method guarantees predictable message transmission while reducing bandwidth utilization.




Fig. 1. An example multi-bus system where each processor connects to a subset of the communication buses.

The rest of this paper is organized as follows. Section 2 presents an overview of the proposed approach, while Section 3 discusses some preliminaries. The message allocation method is developed in Section 4 and Section 5 presents the case study. We briefly discuss some related issues and conclude the paper in Section 6.

2. Design overview

As the primary objective, we construct a network topology meeting the fault-tolerance and performance goals of the embedded applications. The secondary objective is to minimize hardware cost in terms of communication buses. An iterative method is developed where a feasible network topology satisfying performance goals is first obtained. Its cost is then reduced via a series of steps, which minimize the number of buses by appropriately grouping (clustering) messages while preserving the feasibility of the original solution. Since clustering is NP-complete [12], we use heuristics to obtain a feasible solution.

Fig. 2 shows the main steps of the proposed heuristic approach. For a given allocation of tasks to processors, FT-DESIGN accepts the set of task graphs and processors $\{P_i\}$ as inputs, and returns as output, a low-cost network topology comprising identical buses $\{B_j\}$. Redundant routes are provided for messages with specific fault-tolerance requirements; for a k -fault-tolerant (k -FT) message m_i , k replicas or copies are allocated to separate buses. The network is synthesized assuming a generic TDMA protocol, and can accommodate specific cases such as TTP [10] and FlexRay [11] after some modifications.

We assume that each task graph G_i must meet its deadline by the end of its period, $\text{period}(G_i)$. First, the graph deadline is distributed over its tasks to generate a scheduling range $[r_i, d_i]$ for each task T_i where r_i and d_i denote its release time and deadline, respectively. The initial network topology is obtained by simply allocating each inter-processor message m_i to a separate bus. Without bus contention, m_i 's transmission delay is given by the message size and bus bandwidth, and the overall solution is feasible if all tasks complete before their respective deadlines. Section 3 discusses these initial steps in greater detail.

The number of communication buses in the initial solution is then minimized via an iterative message clustering procedure which groups multiple messages on bus B_j . A message m_i is grouped with an existing cluster C_j if the resulting allocation satisfies the following requirements:

- No two replicas of a k -FT message are allocated to C_j .
- All messages belonging to C_j continue to meet their deadlines.
- The duration (length) of the communication schedule corresponding to C_j does not exceed a designer-specified threshold; if a dedicated co-processor handles message

```

Procedure FT-DESIGN( $\{G_i\}, \{P_i\}$ ) /*  $\{G_i\} :=$  Task graphs,  $\{P_i\} :=$  Processors */
for (each  $G_i$ )
    Distribute  $G_i$ 's deadline to obtain the scheduling range  $[r_i, d_i]$  for each task  $T_i$ ;
for (each  $k$ -FT message  $m_i$ ) begin /* Obtain the initial network topology */
    Determine  $m_i$ 's transmission delay  $tdelay(m_i)$ ;
    Allocate each copy of  $m_i$  to a separate bus  $B_j$ ;
end;
for (each task  $T_i$ ) begin /* Determine task schedulability */
     $w_i :=$  Worst-case response time of  $T_i$  on its allocated processor  $P_i$ ;
    if ( $w_i + tdelay(m_i) > d_i - r_i$ ) return  $\emptyset$ ; /* Solution is infeasible */
end;
 $s_{clust} :=$  SYNTH ( $\{m_i\}$ ); /* Synthesize low-cost topology */
Allocate each cluster  $C_j$  in  $s_{clust}$  to a separate bus  $B_j$ ;
return  $\{B_j\}$ ; /* Return the set of communication buses */
    
```

Fig. 2. The overall approach to fault-tolerant communication network synthesis.

communication as in Fig. 1, the schedule must be compact enough to fit within the available memory.

The proposed clustering approach also uses bus bandwidth efficiently by sharing or re-using transmission slots between multiple messages sent by a processor whenever possible. Each message cluster is allocated to a separate bus in the final topology. Section 4 describes the clustering procedure in greater detail.

3. Preliminaries

This section shows how to obtain the initial solution where tasks are assigned deadlines and scheduled on processors, and messages allocated to separate communication buses.

3.1. Deadline assignment

Initially, only the entry and exit tasks having no predecessors and successors, respectively, have their release times and deadlines fixed. To schedule an intermediate task T_i in the task graph, however, its scheduling range $[r_i, d_i]$ must first be obtained. This is termed the *deadline assignment problem* where the deadline D_i of the task graph G_i must be distributed over each intermediate task

such that all tasks are feasibly scheduled on their respective processors. Deadline distribution is NP-complete and various heuristics have been proposed to solve it. We use the approach of Natale and Stankovic [14] which maximizes the slack added to each task in graph G_i while still satisfying its deadline D_i . Their heuristic is simple, and for general task graphs, its performance compares favorably with other heuristics [13].

We now describe the deadline distribution algorithm. Entry and exit tasks in the graph are first assigned release times and deadlines. A path, $path_i$ through G_i comprises one or more tasks $\{T_i\}$; the slack available for distribution to these tasks is $slack_i = D_i - \sum c_i$ where D_i is the deadline of $path_i$ and c_i the execution time of a task T_i along this path. The distribution heuristic in [14] maximizes the minimum slack added to each T_i along $path_i$ by dividing $slack_i$ equally among tasks. During each iteration through G_i , $path_i$ minimizing $slack_i/n$, where n denotes the number of tasks along $path_i$, is chosen and the corresponding slack added to each task along that path. The deadlines (release times) of the predecessors (successors) of tasks belonging to $path_i$ are updated. Tasks along $path_i$ are then removed from the original graph, and the above process is repeated until all tasks are assigned release times and deadlines.

We use the graph in Fig. 3(a) to illustrate the above procedure. First, the release time of entry task T_1

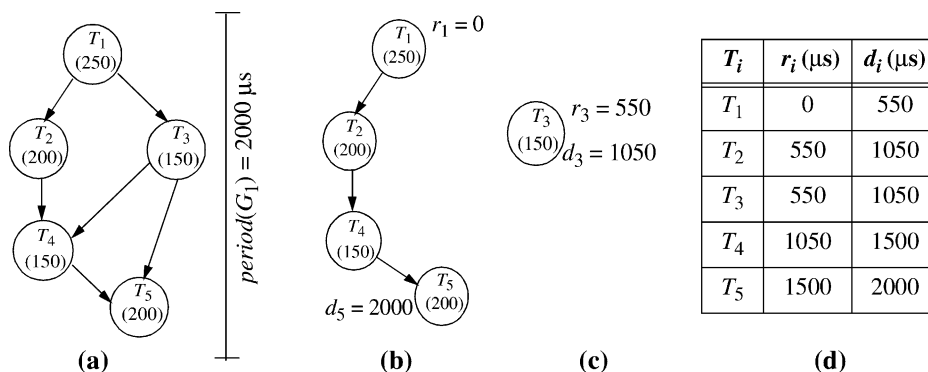


Fig. 3. (a) Example task graph, (b) and (c) paths selected for deadline distribution, and (d) the resulting scheduling ranges for each task.

and the deadline of exit task T_5 are set to $r_1=0 \mu\text{s}$ and $d_5=2000 \mu\text{s}$, respectively. Next, we select the path $T_1T_2T_4T_5$ shown in Fig. 3(b); the total execution time of tasks along this path is $800 \mu\text{s}$, and as per the heuristic, a slack of $(2000 - 800)/4 = 300 \mu\text{s}$ is distributed to each task. Once their release times and deadlines are fixed, these tasks are removed from the graph. Fig. 3(c) shows the remaining path comprising task T_3 which has its release time and deadline fixed by T_1 and T_4 , respectively. Fig. 3(d) shows the resulting scheduling range for each task.

3.2. Task scheduling

Once the scheduling ranges of tasks in the graph are fixed, each T_i may now be considered independent with release time r_i and deadline d_i , and scheduled as such. To tackle multi-rate systems, we use *fixed-priority scheduling* where tasks are first assigned priorities according to their periods [15], and at any time instant, the processor executes the highest-priority ready task. Again, the schedule is feasible if all tasks finish before their deadlines; feasibility analysis of schedules using simple closed-form processor-utilization-based tests has been extensively studied under fixed-priority scheduling [15]. However, in addition to feasibility, we also require task T_i 's response time w_i , given by the time interval between T_i 's release and finish times; the response time is used in the next stage of our algorithm to determine the message delays to be satisfied by the network.

For multi-rate task graphs, the schedules on individual processors are simulated for a duration equal to the least common multiple (LCM) of the graph periods [16]. Since this duration evaluates all possible interactions between tasks belonging to the different graph iterations, the worst-case

response time for each task T_i is obtained. Fig. 4(a) shows a simple multi-rate system comprising two task graphs with periods 2000 and 3000 μs ; Fig. 4(b) and (c) show the task allocation and scheduling ranges, respectively. Fig. 4(d) shows the corresponding schedule for 6000 μs —the LCM of the graph periods. Task response times within this time interval are shown in Fig. 4(e). Multiple iterations of a task are evaluated to obtain its worst-case response time. For example, in Fig. 4(e), the first iteration of tasks T_1, T_2 , and T_4 (in bold) has the maximum response time among the iterations within the given time duration. The task scheduling on processors is successful if, for each task T_i , $w_i \leq d_i - r_i$. However, for the overall solution to be feasible, all messages must also meet their deadlines.

3.3. Initial network topology

A k -FT message m_i sent by task T_i has deadline delay (m_i) = $d_i - r_i - w_i$ where w_i denotes T_i 's worst-case response time. Initially, the network topology allocates a separate communication bus for each message copy. Therefore, in this topology, m_i experiences no network contention and its transmission delay is $\text{size}(m_i)/B_j^{\text{speed}}$ where $\text{size}(m_i)$ and B_j^{speed} denote the message size in bits and bus bandwidth in KB/s, respectively. The solution is feasible if, for each m_i , delay(m_i) is greater than the corresponding transmission delay.

4. Fault-tolerant message clustering

We now develop a clustering approach to reduce the cost of the initial network topology obtained in Section 3 where

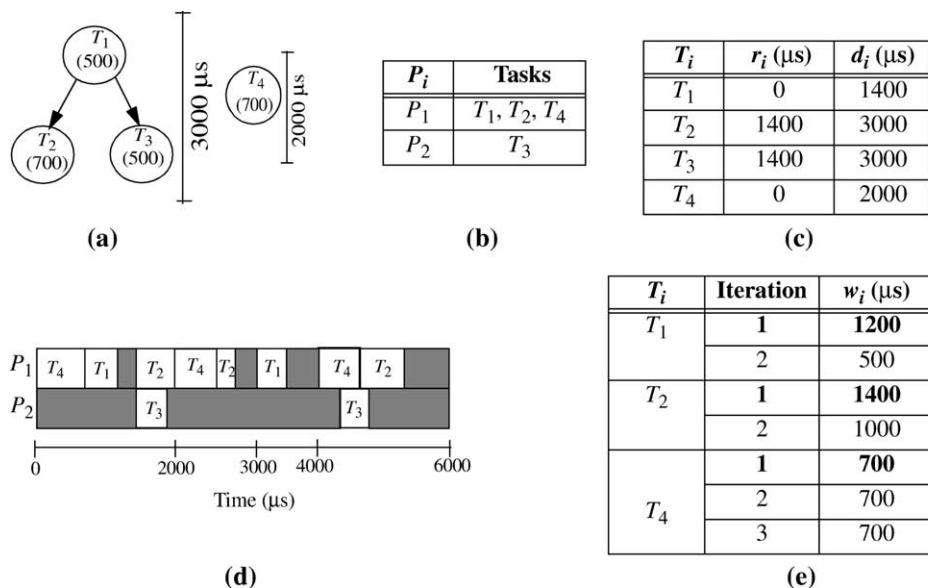


Fig. 4. (a) An example multi-rate system, (b) task-to-processor allocation, (c) task scheduling ranges, (d) task schedule for the duration of the least common multiple of the task periods, and (e) the response times of different task iterations over the simulated time interval.




Fig. 5. A TDMA-based allocation of transmission slots to processors on communication bus B_j .

multiple messages are grouped on a single bus while preserving the feasibility of the original solution. The fault-tolerance requirement of each k -FT message is also satisfied during this procedure.

First, we briefly review message transmission in a typical TDMA communication protocol. As an example, we choose the FlexRay protocol currently under development by a consortium of automotive companies to provide predictable communication for distributed control applications [11]. Fig. 5 shows the TDMA scheme where messages are transmitted according to a static, periodic, and global communication schedule called a *round* comprising identical-sized slots. Each processor P_j is allotted one or more sending slots during a round where both slot size and the number of slots per round are fixed by the system designer. Though successive rounds are constructed identically, the messages sent by processors may vary during a given round.

We now state the fault-tolerant message clustering problem as follows. Given a communication deadline $\text{delay}(m_i)$ for each k -FT message m_i sent by processor P_j , construct TDMA rounds on the minimum number of communication buses such that during any time interval corresponding to $\text{delay}(m_i)$, P_j is allotted a sufficient number of transmission slots to transmit m_i . Allocation of messages to multiple buses is related to *bin-packing* where fixed-size objects (messages) are packed into a bin (round) of finite size while minimizing the number of bins. The general bin-packing problem is NP-complete and heuristics are typically used to obtain a solution [17].

We treat each m_i as a periodic message with period, $\text{period}(m_i)$ equal to its deadline $\text{delay}(m_i)$ and generate message clusters $\{C_j\}$, such that the corresponding TDMA round, $\text{round}(C_j)$ satisfies the constraints previously introduced in Section 2:

- No two replicas of a k -FT message m_i are allocated to C_j .
- The duration of $\text{round}(C_j)$ does not exceed a designer-specified threshold.
- The slots within $\text{round}(C_j)$ guarantee m_i 's deadline, i.e. the time interval between successive sending slots for m_i equals its period.

Each message cluster C_j is allocated to a separate communication bus in the final network topology. Our method also makes efficient use of bus bandwidth by minimizing the number of transmission slots needed to satisfy message deadlines within a TDMA round by reusing slots between messages sent by a processor whenever possible.

We assume an upper bound on TDMA-round duration provided by the designer in terms of the maximum number

of transmission slots n_{\max} and slot duration Δ_{slot} . Typically, the choice of n_{\max} depends on the memory limitations of the communication co-processor such as the number of transmit and receive buffers. Each transmission slot within a round has duration $\Delta_{\text{slot}} = \min_i \{\text{size}(m_i)\} / B_j^{\text{speed}}$ μs . The message period $\text{delay}(m_i)$, originally expressed in microseconds, is now discretized as $\lfloor \text{delay}(m_i) / \Delta_{\text{slot}} \rfloor$ and expressed in terms of transmission-slot intervals. To simplify the notation, we will use $\text{delay}(m_i)$ to denote this discrete quantity from here on.

To guarantee message m_i 's deadline, the corresponding slot allocation must satisfy both its periodicity requirement and a distance constraint between successive m_i transmissions as the following example illustrates. Fig. 6(a) shows an allocation scenario for message m_1 having $\text{delay}(m_1) = 2$ slots within a TDMA round of duration four slots where m_1 requires one slot for transmission. Though m_1 's periodicity requirement may be satisfied by simply allocating sufficient slots within each of its periods, it results in missed deadlines. The interval between successive m_1 transmissions may be as close to one and as far as three slots away. As Fig. 6(a) shows, in the worst case, m_1 may be allocated a transmission slot just before the end of its current period and one immediately at the start of its next period. Clearly, this results in a deadline violation. Similar problems may also occur when multiple messages are clustered. Fig. 6(b) shows TDMA rounds corresponding to messages m_1 and m_2 with periods $\text{period}(m_1) = 2$ and $\text{period}(m_2) = 5$ slots, respectively. Transmission slots are allocated in *first-fit* (FF) fashion where messages are ordered in terms of increasing period and the first available slots allocated to each m_i within the round. The slot allocation in Fig. 6(b) results in a deadline violation where the minimum and maximum distances between successive slots for m_2 are 4 and 6 slots, respectively. Therefore, to guarantee message m_i 's deadline, the corresponding allocation must satisfy a maximum distance between successive m_i transmission slots equal to $\text{period}(m_i)$. Note that in the above example, message deadlines may be satisfied by modifying their periods appropriately. Fig. 6(c) shows the slot allocation for both messages after m_2 's period is modified to four slots. It is easily checked that the distance constraint of two and four slots for successive transmissions of m_1 and m_2 , respectively, is satisfied.

The above discussion suggests that the original message periods may need modification prior to allocating slots within the TDMA round. We adopt a strategy where the periods of all messages within a cluster are constrained to be harmonic multiples of some base period p_{base} , i.e. $\text{period}(m_i) = 2^k p_{\text{base}}$. A similar concept is used while scheduling tasks in real-time systems requiring a specific temporal separation between successive task executions [18,19]. We constrain each m_i 's period to be the maximum integer $\text{period}(m_i) \leq n_{\max}$ satisfying $2^k p_{\text{base}} \leq \text{delay}(m_i) < 2^{k+1} p_{\text{base}}$; if $p_{\min} = \min_i \{\text{period}(m_i)\}$ is the smallest period among the messages, then $p_{\min}/2 < p_{\text{base}} \leq p_{\min}$.

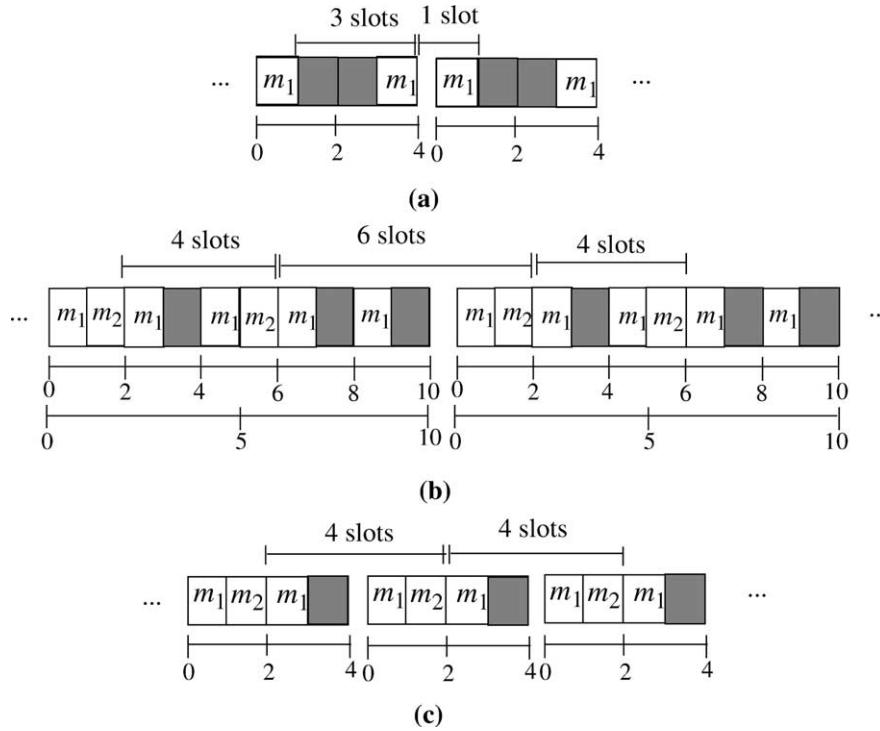


Fig. 6. (a) Message allocation resulting in a missed deadline, (b) a clustering of multiple messages resulting in missed deadlines, (c) a clustering guaranteeing deadlines obtained after modifying message periods appropriately.

Fig. 7 shows the synthesis algorithm to construct the network topology. For each p_{base} value between $[p_{\text{min}}/2, p_{\text{min}}]$, message periods are modified appropriately, and clustered to generate the corresponding topology. Finally, the best solution, in terms of the number of clusters, is chosen.

The CLUSTER procedure in Fig. 8 takes a set of messages s_{msg} as input, their periods modified and sorted in terms of increasing $\text{period}(m_i)$, and returns the set of message clusters s_{clust} as output. During each clustering step, we choose a k -FT message m_i having the minimum period within s_{msg} and allocate it to k separate clusters.

For each m_i , we obtain all feasible message-to-cluster allocations by grouping m_i with each C_j in s_{clust} and generating $\text{round}(C_j \cup m_i)$. If needed, new clusters are created within s_{clust} to accommodate all copies of m_i . If more than k feasible allocations are obtained, then the k best solutions are chosen based on efficient bandwidth use. The exact evaluation criterion is discussed later in this section.

The ALLOC procedure generates a feasible round $(C_j \cup m_i)$. It accepts an existing message cluster C_j and a message m_i and generates a feasible TDMA round (if possible) for the new allocation $C_j \cup m_i$. As discussed above, message m_i 's period, $\text{period}(m_i)$ is first transformed to relate

```

Procedure SYNTH ( $\{m_i\}$ )           /*  $\{m_i\}$  := Message set */
   $p_{\text{min}} = \min_i \{\text{period}(m_i)\}$ ;   /*  $p_{\text{min}}$  denotes the minimum period in the message set */
   $\text{cost}_{\text{min}} :=$  Number of messages in  $\{m_i\}$ ; /* Topology cost where each  $m_i$  is allotted a dedicated bus */
   $s_{\text{msg}} := \emptyset$ ;                 /* Initialize the message set */
  for (each  $p_{\text{base}}$  in  $[\frac{p_{\text{min}}}{2}, p_{\text{min}}]$ ) begin
     $s_{\text{msg}} := \{m_i \mid m_i\text{'s period is the largest integer such that } 2^k \cdot p_{\text{base}} \leq \text{delay}(m_i) < 2^{k+1} \cdot p_{\text{base}}\}$ ;
    Sort messages in  $s_{\text{msg}}$  by increasing period;
     $s_{\text{clust}} := \text{CLUSTER}(s_{\text{msg}})$ ;   /*  $s_{\text{clust}}$  := set of clusters */
     $\text{cost}_{\text{cur}} :=$  Number of clusters in  $s_{\text{clust}}$ ;
    if ( $\text{cost}_{\text{cur}} < \text{cost}_{\text{min}}$ ) begin
       $\text{cost}_{\text{min}} := \text{cost}_{\text{cur}}$ ;
      Store  $s_{\text{clust}}$  as current best solution;
    end;
     $s_{\text{msg}} := \emptyset$ ;
  end;
  return  $s_{\text{clust}}$ ;                 /* Return the best allocation */

```

Fig. 7. Algorithm to synthesize the network topology.

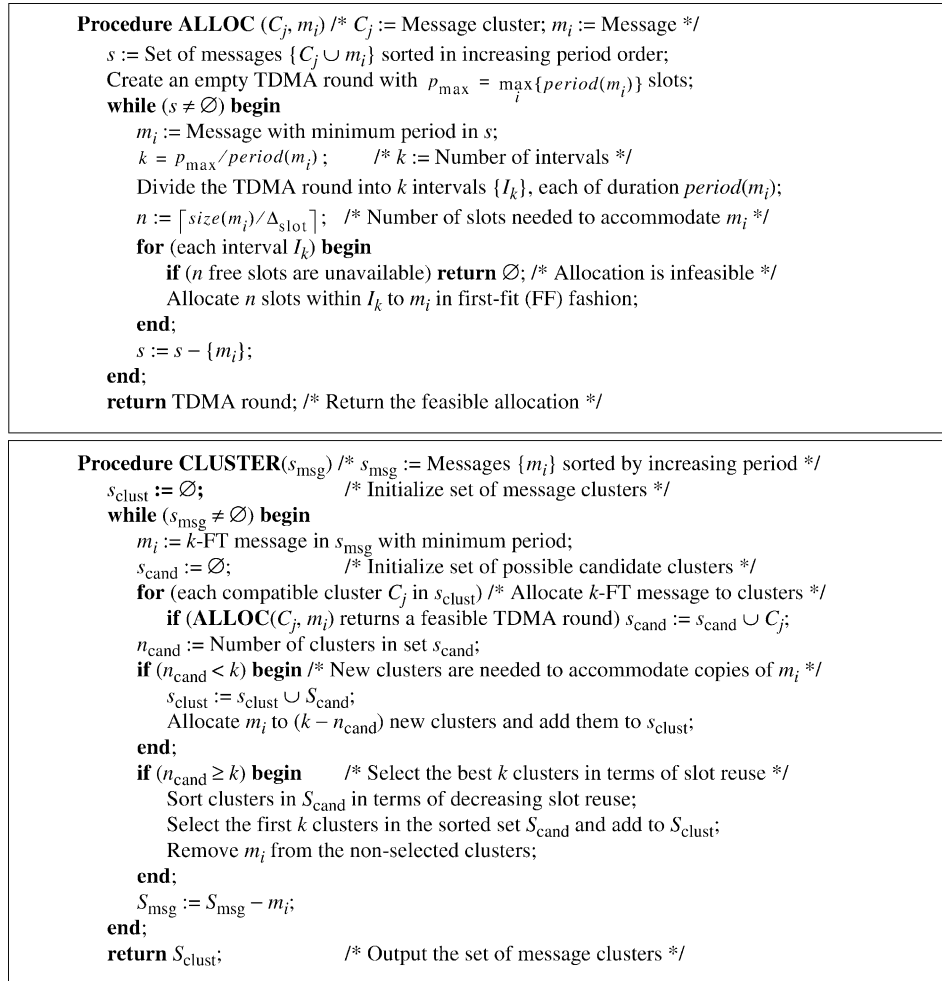


Fig. 8. The clustering algorithm generating the reduced-cost network topology.

harmonically to those in C_j and the messages are sorted in increasing period order. The duration of the new round, $round(C_j \cup m_i)$ is $p_{\max} = \max_{C_j} \{period(m_i)\}$. To allocate transmission slots for the new message m_i , ALLOC divides $round(C_j)$ into k disjoint time intervals $\{I_k\}$ where $k = p_{\max} / period(m_i)$ and I_k has duration $period(m_i)$. Transmission slots are then allotted within each interval using the FF packing strategy. The distance constraint between transmission slots for m_i is guaranteed since the allotted slots occur in the same positions within each interval I_k .

The computational complexity of the synthesis procedure is $O(n^3)$ where n is the number of messages; the outer while loop of CLUSTER iterates through all n messages, and during each iteration, ALLOC explores all message to cluster allocations, a process of complexity $O(n^2)$.

Theorem 1. *Given a cluster $C_j = \{m_1, \dots, m_n\}$ with harmonically related message periods, the TDMA round generated by ALLOC guarantees message transmission deadlines.*

Proof. We show by induction on the number of messages that for each m_i in C_j , and during any time interval corresponding to $period(m_i)$, ALLOC allocates a sufficient

number of transmission slots to send m_i . Assume that the messages $\{m_1, \dots, m_n\}$ are ordered in terms of increasing period. Also, let $I_k^{m_i}$ be the k th interval within the TDMA round containing transmission slots for message m_i .

Consider the base case when allocating the first message m_1 within the TDMA round of duration p_{\max} . The round is divided into intervals $\{I_k^{m_1}\}$, each of duration $period(m_1)$. Initially, all slots within the round are free and the FF packing scheme allots the first available slots within each interval to m_1 . The maximum distance between two successive m_1 transmission slots equals $period(m_1)$, and thus, ALLOC guarantees m_1 's deadline.

Now, assume that transmission slots have been previously allotted for messages $\{m_1, \dots, m_i\}$ to satisfy their respective distance constraints, and let m_{i+1} be the next message being considered. As before, the round is divided into intervals $\{I_k^{m_{i+1}}\}$. Since messages are ordered in terms of increasing period, $I_k^{m_{i+1}} = pI_k^{m_i}$, $p \geq 1$. We consider the following cases:

$p = 1$. Both intervals have the same duration. Since the allocated slots for m_i satisfy its minimum distance constraint, all allotted and free slots occur in the same

positions within $I_k^{m_i}$. The FF scheme chooses the first available slots within each $I_k^{m_i}$ to accommodate m_{i+1} . Therefore, the newly allotted slots for m_{i+1} also occur in the same positions within each interval.

$p > 1$. Since message periods are harmonic multiples of each other, each $I_k^{m_{i+1}}$ spans the same number of m_i intervals. Furthermore, since each m_i interval is identical in terms of the slot occurrences, it follows that each $I_k^{m_{i+1}}$, comprising a sequence of these intervals, is also identical. Therefore, the FF scheme chooses slots occurring in the same positions within $I_k^{m_{i+1}}$ for message m_{i+1} .

□

Theorem 2. Given a cluster $C_j = \{m_1, \dots, m_n\}$ where the message periods are harmonically related, ALLOC guarantees a feasible message allocation if one exists.

Proof. We prove this theorem by showing that if ALLOC fails to find a feasible allocation for C_j , then no other packing scheme will. Given a feasible TDMA round for messages $\{m_1, \dots, m_i\}$, ALLOC fails to find an allocation for the new message m_{i+1} only if the number of transmission slots within an interval $I_k^{m_{i+1}}$ is insufficient. Clearly, rearranging messages within this interval does not increase the number of available slots. Now assume m_{i+1} is successfully allocated within $I_k^{m_{i+1}}$ after removing a previously allocated message m_i . This implies that m_{i+1} now occupies some slots previously used by m_i . However, since $I_k^{m_i} \leq I_k^{m_{i+1}}$, there exists at least one interval where m_i cannot be accommodated (All slots in the interval are allocated). Therefore, no feasible allocation for the cluster C_j exists.

□

4.1. Transmission-slot reuse

Recall that during clustering, each message m_i is treated as periodic with period, $\text{period}(m_i)$. However, if the task T_i transmitting m_i does not execute at that rate, then the bus bandwidth is over-utilized. We can improve bandwidth utilization by reusing transmission slots among the multiple messages sent by processor P_j .

The worst-case arrival rate $\text{arrival}(m_i)$ for each message m_i in a multi-rate system is obtained during schedulability analysis by simulating the corresponding task schedule. It is important to note that $\text{arrival}(m_i)$, expressed in terms of slot intervals, depends on the execution rate of the sender task T_i . Let $\{m_i\}$ be the set of messages sent by a processor within a message cluster C_j . Now, assume message m_{new} , also transmitted by the same processor, to be allotted slots within $\text{round}(C_j)$. If each message m_i is allotted n_i transmission slots within the time interval $\text{period}(m_{\text{new}})$ in $\text{round}(C_j)$, then the number of slots available for reuse by m_{new} is

$$n_{\text{reuse}} = \sum_{m_i} n_i - \sum_{m_i} \left\lceil \frac{\text{period}(m_{\text{new}})}{\text{arrival}(m_i)} \right\rceil n_i$$

where $\text{arrival}(m_i)$ denotes the worst-case arrival rate of message m_i . Therefore, m_{new} is allotted

$$\left\lfloor \frac{\text{size}(m_{\text{new}})}{B_j^{\text{speed}} \Delta_{\text{slot}}} \right\rfloor - n_{\text{reuse}}$$

transmission slots within $\text{period}(m_{\text{new}})$.

Fig. 9 shows an example of slot reuse among four messages sent by a processor assuming a network bandwidth of 250 KB/s and slot width of 50 μs. Fig. 9(a) lists message attributes in increasing order of their periods

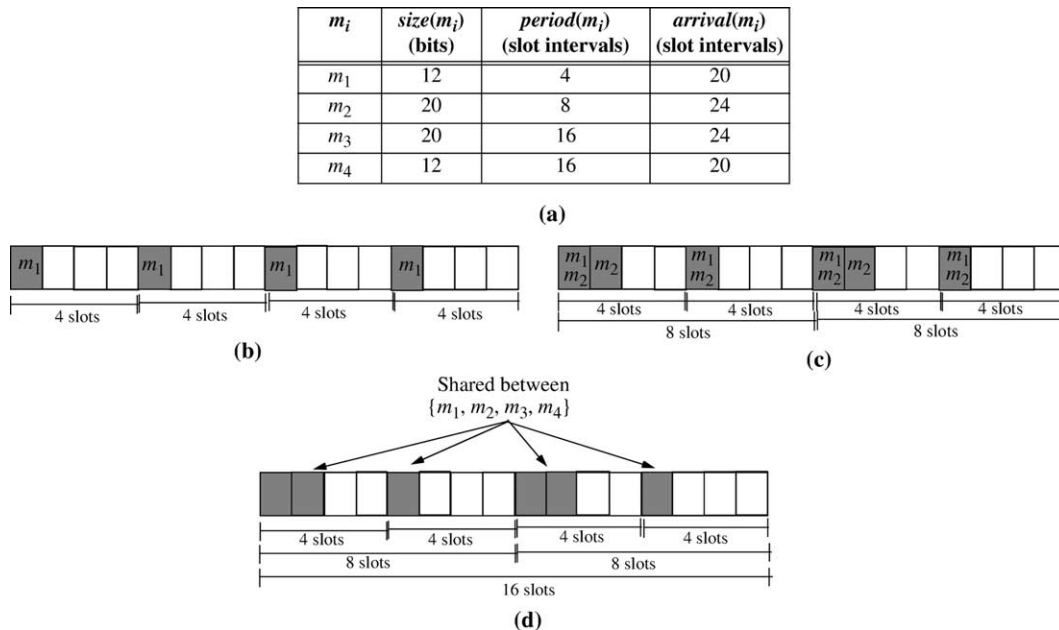


Fig. 9. An example of transmission slot reuse during allocation: (a) messages listed in increasing order of their periods, and (b)–(d) illustration of slot reuse as messages are added to the TDMA round.

and Fig. 9(b)–(d) show how slots are reused between messages while constructing the TDMA round. In Fig. 9(b), m_1 is allotted one transmission slot during any time interval spanning $\text{period}(m_1)$. Before allotting slots for m_2 , the number of reusable slots during $\text{period}(m_2)$ is determined as $2 - \lceil 8/20 \rceil = 1$; m_1 may use one transmission slot during any $\text{period}(m_2)$. Since m_2 requires two slots for transmission, we allot one additional slot per $\text{period}(m_2)$ in Fig. 9(c). When allocating m_3 , we determine the number of reusable slots within $\text{period}(m_3)$ as $6 - (\lceil 16/20 \rceil + \lceil 16/24 \rceil) = 3$. Since, m_3 may reuse three slots during any $\text{period}(m_3)$, no new slots are allotted. A similar argument holds when allocating message m_4 . It is straightforward to extend the ALLOC procedure to include slot reuse.

Theorem 3. *The TDMA round generated by ALLOC for the cluster C_j with transmission-slot reuse guarantees message deadlines.*

Proof. We show by induction on the number of messages that for each message m_i sent by processor P_j , ALLOC allocates a sufficient number of sending slots during any $\text{period}(m_i)$ even when its slots are reused by other messages.

As before, we assume that messages are ordered in terms of increasing period.

Consider the base case where the first message m_1 is allotted to an empty TDMA round. No slots can be reused and m_1 is simply allotted enough transmission slots during any $\text{period}(m_1)$ span. Now, assume messages $\{m_1, \dots, m_i\}$ sent by P_j have been allotted slots satisfying their deadlines and a new message m_{new} , also sent by P_j , needs to be grouped with this set. In the worst case, during any time interval spanning $\text{period}(m_{\text{new}})$, the previously scheduled messages may use up to

$$\sum_{m_i} \left\lceil \frac{\text{period}(m_{\text{new}})}{\text{arrival}(m_i)} \right\rceil n_i$$

slots for transmission. Since m_{new} is sent only during the unused slots (see above discussion), messages $\{m_1, \dots, m_i\}$ continue to meet deadlines even when some of their transmission slots are reused.

□

Given a set of clusters and a new message to be allocated to one, CLUSTER explores all possible cluster-message

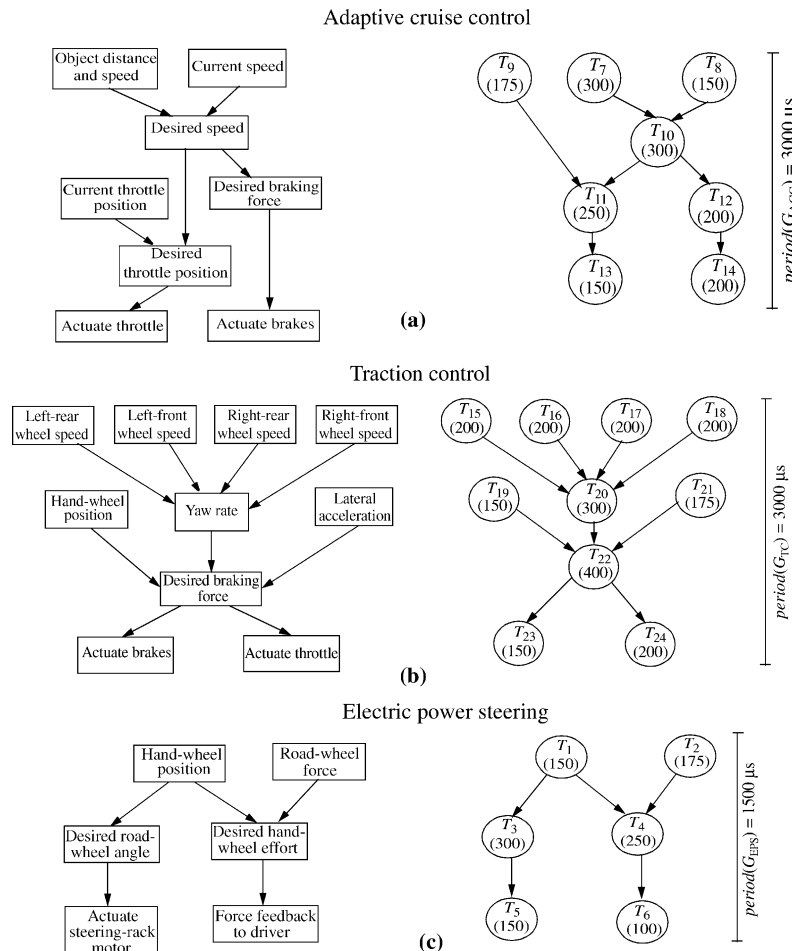


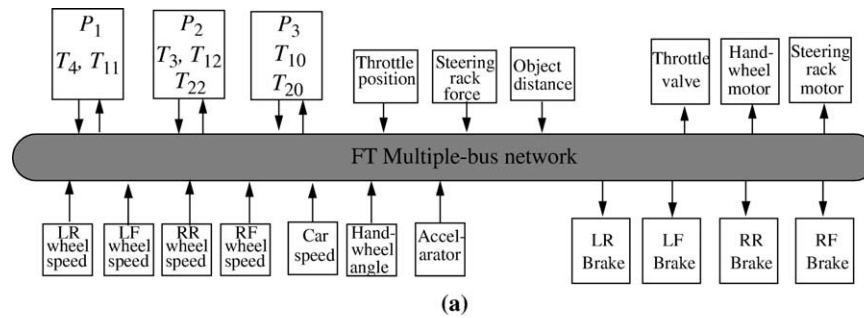
Fig. 10. (a) Adaptive cruise control, (b) traction control, and (c) electric power steering applications, and the corresponding flow-graph representations.

allocation scenarios. Slot reuse is used as the deciding factor in selecting the best allocation since the cluster allocation resulting in maximum reuse minimizes the bandwidth utilization. Finally, when TDMA slots are shared between messages sent by a processor, as in Fig. 9(d), the communication co-processor must correctly schedule their transmission, i.e. given a slot, decide which message to transmit in it. Though this paper does not address message-scheduling logic within the co-processor, an earliest-deadline first approach seems appropriate.

5. Case study

We now illustrate the proposed synthesis method using some advanced automotive control applications as examples. These include adaptive cruise control (ACC),

electric power steering (EPS), and traction control (TC), and are detailed in Fig. 10(a)–(c). The ACC application automatically maintains a safe following distance between two cars, while EPS uses an electric motor to provide necessary steering assistance to the driver. The TC application actively stabilizes the vehicle to maintain its intended path even under slippery road conditions. These applications demand timely interaction between distributed sensors, processors, and actuators, i.e. have specific end-to-end deadlines, and therefore require a dependable communication network. Fig. 11(a) shows the physical architecture of the system where sensors and actuators are directly connected to the network and the designer-specified task-to-processor allocation, while Fig. 11(b) summarizes the various message attributes affecting network topology generation. We assume 1-FT messages throughout. Columns two and three list the sending and receiving tasks for



Message m_i	(Sender, receiver)	$size(m_i)$ (bits)	$delay(m_i)$ (μs)	$delay(m_i)$ (slot intervals)
m_1	(T ₁ , T ₃)			
	(T ₁ , T ₄)	12	300	6
m_2	(T ₂ , T ₄)	12	275	5
m_3	(T ₃ , T ₅)	8	300	6
m_4	(T ₄ , T ₆)	12	350	7
m_5	(T ₇ , T ₁₀)	12	500	10
m_6	(T ₈ , T ₁₀)	12	650	6
m_7	(T ₉ , T ₁₁)	10	1425	28
m_8	(T ₁₀ , T ₁₁)			
	(T ₁₀ , T ₁₂)	12	500	10
m_9	(T ₁₁ , T ₁₃)	10	500	10
m_{10}	(T ₁₂ , T ₁₄)	10	500	10
m_{11}	(T ₁₅ , T ₂₀)	12	475	9
m_{12}	(T ₁₆ , T ₂₀)	12	475	9
m_{13}	(T ₁₇ , T ₂₀)	12	475	9
m_{14}	(T ₁₈ , T ₂₀)	12	475	9
m_{15}	(T ₁₉ , T ₂₂)	10	1100	22
m_{16}	(T ₂₀ , T ₂₂)	10	275	5
m_{17}	(T ₂₁ , T ₂₂)	8	1025	20
m_{18}	(T ₂₂ , T ₂₃)			
	(T ₂₂ , T ₂₄)	12	650	6

(b)

Fig. 11. (a) The physical architecture including task-to-processor allocation and (b) the message attributes required for network topology construction.

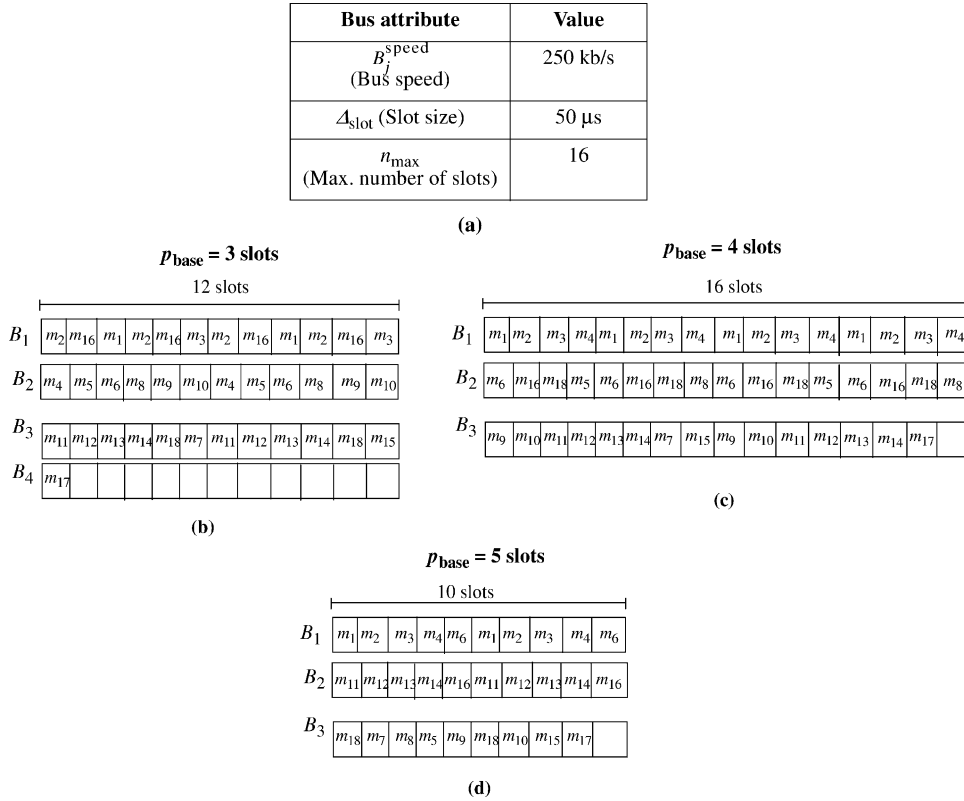


Fig. 12. (a) Example TDMA round specifications and (b) communication schedules generated without slot reuse where message periods are modified to relate harmonically to (b) $p_{\text{base}}=3$, (c) $p_{\text{base}}=4$, and (d) $p_{\text{base}}=5$ slots, respectively.

each message and the message size $size(m_i)$ in bits, respectively, while columns four and five list the communication delay, $delay(m_i)$ for messages in microseconds, and transmission-slot intervals. These delay values are obtained by first assigning deadlines to tasks and then performing a schedulability analysis on their respective processors (see Section 3).

As summarized in Fig. 12(a), we assume a version of the FlexRay communication protocol having a bandwidth of 250 KB/s and a minimum transmission-slot width of 50 μ s. Since m_2 and m_{16} have the minimum period of five slots among all messages, p_{base} may assume values of three, four, or five slots. Fig. 12(a)–(c) show the communication schedules generated by SYNTH (without slot reuse) after modifying the message periods to relate harmonically to each of the above p_{base} values. Those corresponding to p_{base} values of 4 and 5 slots compare best in terms of topology cost.

We now show how to reduce bandwidth utilization by sharing transmission slots between messages. As candidates for slot reuse, consider messages m_3 and m_{10} sent by tasks T_3 and T_{12} , respectively, where both tasks are allocated to processor P_2 . In Fig. 13(a), where message periods are modified using $p_{\text{base}}=3$, m_3 and m_{10} cannot share slots since both have a periodicity of six slots. In Fig. 13(b), however, when their periods are modified as $period(m_3)=4$ and $period(m_{10})=8$ using $p_{\text{base}}=4$ slots, reuse is possible. Note that the EPS application comprising T_3 transmitting

m_3 has a 1500 μ s period-corresponding to the inter-interval time between successive m_3 transmissions. Therefore, in Fig. 12(b), m_3 requires only one of four allocated slots on bus B_1 (Task T_3 , however, may request m_3 's transmission anytime during the round), and m_{10} with a period of eight slots can reuse the one free slot available during any period(m_{10}). A similar argument holds for messages m_4 and m_9 sent by processor P_1 . Fig. 12(c) shows the schedule corresponding to $p_{\text{base}}=5$ slots. Again, slots are reused between messages $\{m_3, m_{10}\}$ and $\{m_4, m_9\}$.

Finally, though the topologies in Fig. 13(b) and (c) have the same cost (three buses each), Fig. 13(b) has a somewhat lower slot utilization of 89.5% compared to 90% for Fig. 13(c). Since the empty slots in Fig. 13(b) may be used to transmit additional (non-critical) messages when compared to Fig. 13(c), we select the topology in Fig. 13(b) as the final solution.

6. Conclusions

This paper has addressed the synthesis of low-cost TDMA communication networks for distributed embedded systems. We have developed a fault-tolerant clustering method which allocates and schedules k -FT messages on the minimum number of buses to provide dependable transmission. The proposed method was illustrated using a case

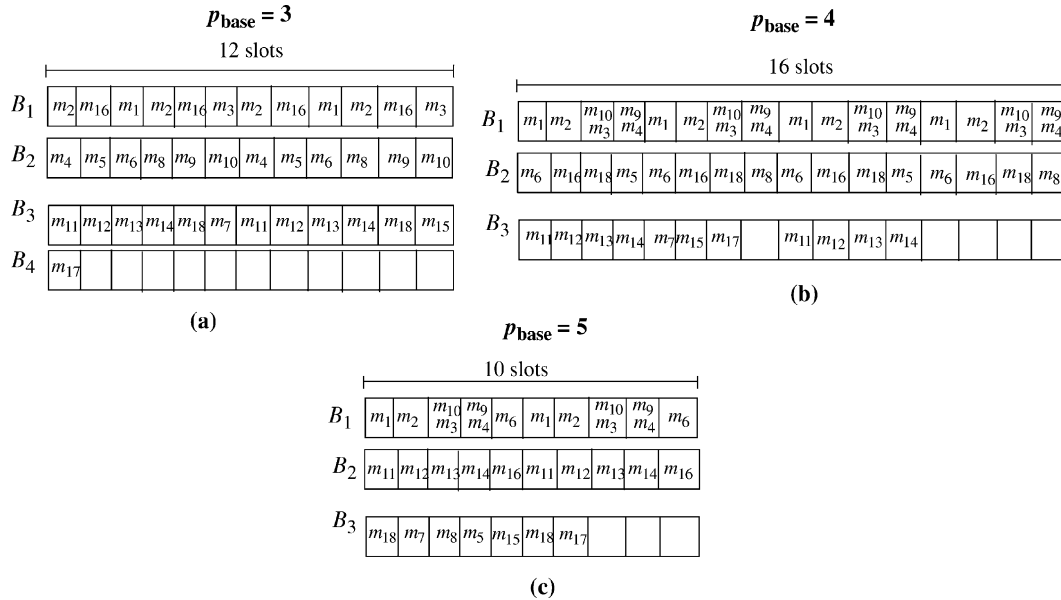


Fig. 13. Communication schedules generated while reusing transmission slots for different values of p_{base} : (a) $p_{base}=3$, (b) $p_{base}=4$, and (c) $p_{base}=5$ slots.

study involving some advanced automotive control applications and it was shown that sharing transmission slots among multiple messages reduces bandwidth consumption while preserving predictable communication. Therefore, the method has the potential to reduce topology cost when applied to larger embedded systems.

This paper does not address the design and implementation of the message scheduler on the co-processors (responsible for transmitting and receiving messages in their respective slots). We also do not address the fault-tolerant allocation of tasks to processors. The message clustering scheme can be easily incorporated into an overall synthesis scheme dealing with both task and message allocation. These issues will be investigated as part of future work.

Acknowledgements

This research was supported by a contract from The Delphi Corporation.

References

- [1] Bretz EA. By-wire cars turn the corner. *IEEE Spectr* 2001;38(4): 68–73.
- [2] Leen G, Heffernan D. Expanding automotive electronic systems. *IEEE Comput* 2002;35(1):88–93.
- [3] Ortega RB, Borriello G. Communication synthesis for distributed embedded systems. In: Proceedings of the international conference on computer-aided design (ICCAD); 1998. p. 437–44.
- [4] Abdelzaher TF, Shin KG. Combined task and message scheduling in distributed real-time systems. *IEEE Trans Parallel Distrib Syst* 1999; 10(11):1179–91.
- [5] Doboli A, Eles P, Peng Z, Pop P. Scheduling with bus access optimization for distributed embedded systems. *IEEE Trans VLSI Syst* 2000;8(5):472–91.
- [6] Yen TY, Wolf W. Communication synthesis for distributed embedded systems. In: Proceedings of the international conference on computer-aided design (ICCAD); 1995. p. 288–94.
- [7] Rhodes DL, Wolf W. Co-synthesis of heterogeneous multiprocessor systems using arbitrated communication. In: Proceedings of the international conference on computer-aided design (ICCAD); 1999. p. 339–42.
- [8] Prakash S, Parker AC. Synthesis of application-specific multiprocessor architectures. In: Proceedings of the ACM/IEEE design automation conference; 1991. p. 8–13.
- [9] Kopetz H. Real-time systems: design principles for distributed embedded applications. Boston, MA: Kluwer Academic Publishers; 1997.
- [10] Kopetz H. TTP—a time-triggered protocol for fault-tolerant real-time systems. In: Proceedings of the IEEE fault-tolerant computing symposium; 1993. p. 524–33.
- [11] Berwanger J, Ebner C, Fluhrer S, Fuchs E, Hedenetz B, Kuffner W, Kruger A, Lohrmann P, Millinger D, Peller M, Ruh J, Schedl A, Sprachmann M. FlexRay—the communication system for advanced automotive control systems. In: Proceedings of the SAE World Congress; 2001. Paper: 2001-01-0676.
- [12] Wolf W. An architectural co-synthesis algorithm for distributed embedded computing systems. *IEEE Trans VLSI Syst* 1997;5(2):218–29.
- [13] Kao B, Garcia-Molina H. Deadline assignment in a distributed soft real-time system. *IEEE Trans Parallel Distrib Syst* 1997;8(12):1268–74.
- [14] Natale MD, Stankovic JA. Dynamic end-to-end guarantees in distributed real-time systems. In: Proceedings of the real-time systems symposium; 1994. p. 216–27.
- [15] Liu CL, Layland J. Scheduling algorithms for multiprogramming in a hard real-time environment. *J ACM* 1973;24(1):46–61.
- [16] Hu X, D'Ambrosio JG, Murray BT, Tang DL. Co-design of architectures for automotive powertrain modules. *IEEE Micro* 1994; 14(4):17–25.
- [17] Johnson DS. Fast algorithms for bin packing. *J Comput Syst Sci* 1974; 3(3):272–314.
- [18] Lin KJ, Herkert A. Jitter control in time-triggered systems. Proceedings of the Hawaii international conference on system sciences; 1996. p. 451–59.
- [19] Han CC, Lin KJ, Hou CJ. Distance-constrained scheduling and its applications to real-time systems. *IEEE Trans Comput* 1996;45(7): 814–26.