

A Versatile Architecture for the Distributed Sensor Integration Problem

S. S. Iyengar, *Senior Member, IEEE*, D. N. Jayasimha, *Member, IEEE*, and D. Nadig

Abstract—The computational issues related to information integration in multisensor systems and distributed sensor networks has become an active area of research. From a computational viewpoint, the efficient extraction of information from noisy and faulty signals emanating from many sensors requires the solution of problems related a) to the architecture and fault tolerance of the distributed sensor network, b) to the proper synchronization of sensor signals, and c) to the integration of information to keep the communication and the centralized processing requirements small. In this paper, we propose a versatile architecture for a distributed sensor network which consists of a multilevel network with the nodes (processing element/sensor pairs) at each level interconnected as a deBruijn network. We show that this multilevel network has reasonable fault tolerance, admits simple and decentralized routing, and offers easy extensibility. We model information from sensors as real valued intervals and derive an interesting property related to information integration in the presence of faults. Using this property, the search for a fault is narrowed down to two potentially faulty sensors or communication links. In a distributed environment, information has to be integrated from “temporally close” signals in the presence of imperfect clocks in a distributed environment. We apply the results of past research in this area to state various relationships between the clocks of the processing elements in the network for proper information integration.

Index Terms—Abstract estimate, clock synchronization, distributed sensor networks, deBruijn networks, fault tolerance, information integration.

I. INTRODUCTION

IN RECENT YEARS, there has been increasing interest in the development of distributed sensor networks (DSN's) for information gathering. This is partly because of the availability of new technology which makes the DSN's economically feasible to implement and the increasing complexity of today's information gathering tasks to which they are applied. These tasks are usually time-critical and rely on the reliable delivery of accurate information for their completion. To meet these requirements, a DSN must be able to dynamically respond to fault conditions, reconfiguring its activities as necessary to compensate for disturbances. Thus, the search for efficient, fault-tolerant architectures for DSN's has become an important

area in research. A DSN consists of a set of sensors, a set of processing elements (PE's), and a communication network interconnecting the various PE's. One or more sensors is associated with each PE. We refer to the PE and its associated sensor(s) as a *node*.

The integration of multiple, disparate sensors into a useful sensor network involves the solution of several different problems. For an excellent discussion of the problems and the current state of the art in multisensor integration, the reader is referred to the survey paper by Luo and Kay [9]. From a computational viewpoint, however, the efficient extraction of information from noisy and possibly faulty signals emanating from many sensors requires the solution of problems relating 1) to the architecture and the fault tolerance of the distributed sensor network, 2) to the proper synchronization of sensor signals, and 3) to the integration of information to keep the communication and the processing requirements small.

Wesson *et al.* [2] were the first to attempt designing efficient networks for distributed sensing. They proposed the hierarchical and committee interconnection topologies. A sensor network based on a fixed number of complete binary trees fully interconnected at their roots (we will refer to this network as a flat tree network) was considered in [11], [12] and the following issues were studied:

- 1) the integration of information in real time when clocks at the nodes are not perfect,
- 2) the transmission of information without incurring heavy communication costs, and
- 3) the fault tolerance of the network to certain types of faults.

A. Scope of the Paper

In this paper, which is a continuation of research reported in [11], [12], we propose a new versatile architecture based on the deBruijn network, (first proposed by Pradhan [3]), which has several advantages over the flat tree network. Specifically, the proposed network has better fault-tolerant properties and supports more nodes than the latter with the same diameter. We show how information integration could be achieved in this network and derive an interesting property related to such integration in the presence of faults.

This paper is organized as follows. Section I-B has a brief overview of sensor integration. The notations and definitions used in the paper are presented in Section I-C. After motivating the need for a new sensor network in Section II-A, we propose a multilevel network with each level having the deBruijn

Manuscript received January 13, 1992. This work was supported in part by The Office of Naval Research under Grant ONR-N00014-91-J-1306, in part by the LEQFS—Board of Regents under Grant LEQFS-RD-A-04, and in part by the National Science Foundation under Grant CCR 8908189.

S. S. Iyengar and D. Nadig are with the Department of Computer Science, Louisiana State University, Baton Rouge, LA 70803.

D. N. Jayasimha is with the Department of Computer and Information Science, The Ohio State University, Columbus, OH 43210. In 1993–1994, he will be at NASA-Lewis Research Lab.

IEEE Log Number 9212756.

interconnection in Section II-B. Algorithms for routing in this network are described in Section II-C. We describe sensor integration in the presence of faults in Section II-D. The fault tolerant properties of the network are the subject of Section III. In a DSN, it is necessary that the clocks on the nodes be synchronized. A variant of a previously known method for synchronizing clocks is described for the network in Section III. In Section IV, we compare topological, routing and fault-tolerant properties of the proposed DSN with other sensor networks. We conclude the paper by highlighting the features of the proposed network and indicate the future directions this area of research could possibly take.

B. An Overview of Sensor Integration

The PE's of a DSN combine the sensor output readings to derive an accurate value of the physical process that the sensors monitor. This process of combining the sensor outputs is called *information integration* or *data fusion*.

The method used to integrate the information passed by the sensors depends on whether the sensors provide competitive information or complementary information. In the former case, each sensor ideally provides identical information. This redundancy of the sensor readings helps in enhancing the reliability and fault tolerance of the network. Also, noise in the signals can be detected and removed. This is because the noise in different sensor signals tend to be uncorrelated while the signals of interest are correlated. It is therefore necessary for the information from the sensors to be combined in a meaningful and effective manner, so that the result is fairly accurate. Complementary information integration is done when only partial information is available from each sensor; such information is then integrated to obtain the result.

Following Marzullo [10], we distinguish between a *concrete sensor* and an *abstract sensor*. A concrete sensor is a device that can be used to sample a physical state variable. An abstract sensor is a piecewise continuous function from a physical state variable to a dense interval of real numbers. The reasons for using an abstract sensor rather than a concrete sensor are detailed in [10], [11]. Determining the function which maps a concrete sensor to an abstract sensor depends on many factors such as the choice of a particular sensor type (e.g., motion detecting sensor, range finding sensor, vision sensor), the compensation that has to be applied to the raw sensor value which is itself dependent on the local values of certain parameters (e.g., design parameters of the sensor), the nature of the application, etc. For instance, if a sensor reads a value to be S and its maximum error is known to be E , then an abstract sensor, albeit simple, could be the interval $(S - E, S + E)$. A PE at a node converts a concrete sensor to an abstract sensor. The abstract sensors are *combined* or *integrated* to form an *abstract estimate*. The particular method of combining depends on the integration algorithm used. To keep the terminology simple, we refer to the abstract sensor as the abstract estimate also. An abstract estimate could, in turn, be combined with one or more abstract estimates to form a new abstract estimate.

Marzullo [10] considers the case of a processor receiving input from several sensors whose outputs are intervals. He




Fig. 1. Integrating six intervals with three incorrect.

gives a fault tolerant integration algorithm which takes as input the intervals representing the sensors and gives as output an interval which always contains the actual physical value. A *correct sensor* is one whose interval contains the actual physical value. Hence, any two correct sensors must intersect since they both contain the physical value being measured.

Marzullo considers the case when at most f ($f < n$) sensors are faulty in a n -sensor network. The physical value would then be contained in *any* of the $(n - f)$ intersecting intervals. Since it is not possible to decide which intersection contains the physical value, the smallest connecting interval containing all the $(n - f)$ intersections is taken to be the output of the processor. It can be seen that this final estimate contains the actual physical value. The final estimate, however, becomes arbitrarily wide as the number of faulty sensors becomes large. In such cases, an integration method described in [12] reduces the width of the final abstract estimate. For simplicity, we will use Marzullo's model for information integration in the proposed network. Fig. 1 illustrates Marzullo's method for integrating six intervals of which three are incorrect (i.e., $n = 6$, $f = 3$).

In this paper, we concentrate on competitive information integration. The architecture described here could be used effectively for complementary information integration in the presence of noise and possibly faulty sensors.

C. Notations and Definitions

We model the DSN by an undirected graph $G = (V, E)$, where each node represents one or more sensors and an associated PE of the network, and each edge represents a communication link of the network. The *length* of a path between two nodes is the number of edges encountered while going from one node to another. The *distance* between two nodes is the shortest length between the nodes. The *diameter* of the network is the largest distance between any two nodes in the network. The *degree* of a node is the number of edges associated with that node. The *degree of the network* is the largest degree of any node in the network.

Let y^f represents a binary number with bit y repeated f times; \bar{y} represents the complement of y , and x represents the *don't care* bit. For example, the binary number 00011xx is represented by $0^31^2x^2$. A node i in a network with $N = 2^k$ nodes has the binary address $i_{k-1}i_{k-2} \dots i_1i_0$ where $i_{k-1}(i_0)$ is the most (least) significant bit. The following definitions describe two address transforming functions append (*app*) and strip (*str*).

Let M be a k -bit number. Then

$$\text{app}(M, y) = My$$

$$\text{str}(i_{k-1}i_{k-2}\cdots i_1i_0) = i_{k-1}i_{k-2}\cdots i_1.$$

For example, $\text{app}(000, 1) = 0001$ and $\text{str}(0010) = 001$. Note that $\text{str}(\text{app}(M, y)) = M$.

Our interest lies in multi-level networks (MLNs) in which each node of the network can be associated with a *level* number. An l -level network has l levels numbered from 0 to $l-1$. The set of nodes at level m to which a node i at level m is connected form the *neighbors* of i . The set of nodes to which i is connected at level $m-1$ form the *parents* of i . The set of nodes to which i is connected at level $m+1$ form the *children* of i . In the MLN that we consider for the proposed DSN, there is a single node called the *root* at level 0, and each node at a higher level number has at most one parent and at most r children. We refer to such a network as a r -ary MLN. The node i at level $m > 0$ has the address $i_{m-1}i_{m-2}\cdots i_1i_0$, where each digit $i_j \in \{0, 1, \dots, r-1\}$ ($0 \leq j < m$). This node i is connected to at most r children nodes whose addresses are $\text{app}(i, 0), \text{app}(i, 1), \dots, \text{app}(i, r-1)$, and to its parent node whose address is $\text{str}(i)$. For every node i at level m , the relation $\Phi_m(i)$ yields the set of nodes to which i is connected at level m . In the network proposed, all but the 0^{th} level of the network have the same interconnection scheme at each level.

Hence two nodes i and j in this network are connected iff

- 1) $j = \text{app}(i, b)$, or
- 2) $j = \text{str}(i)$, or
- 3) $j = \Phi(i)$

where $b \in \{0, 1, \dots, r-1\}$.

A real interval $R = (R_l, R_u)$ is represented by a pair of real numbers; R_l is called the *lower bound* and R_u is called the *upper bound* of the interval R . We shall refer to real intervals simply as intervals.

The width of the interval, $|R|$, equals $(R_u - R_l)$. The *set theoretic intersection* of two intervals, X and Y is defined as

$$X \cap Y = \{c \mid c \in X \text{ and } c \in Y\}.$$

Correspondingly, two intervals are said to *intersect* (or *overlap*) if their set theoretic intersection is non empty. Hence, if the set theoretic intersection of X and Y is non empty then their interval intersection is the interval $(\max(X_l, Y_l), \min(X_u, Y_u))$. A special case of interval intersection is interval inclusion. X *includes* Y if $X_l < Y_l$ and $X_u > Y_u$. The *span* of two intervals X and Y is defined as

$$X \cup Y = (X_l, Y_u).$$

Note that the span operation between two nonoverlapping intervals may result in an interval that includes points not lying in either of the intervals.

Intervals X and Y are said to be *non distinct* if either X includes Y or Y includes X ; otherwise, X and Y are said to be *distinct*.




Fig. 2. A flat tree network with 12 nodes.

II. ARCHITECTURE OF THE DISTRIBUTED SENSOR NETWORK

This section describes the architectural features of the proposed network. We provide the motivation for this architecture by reviewing the past work of other researchers and pointing out the shortcomings of their approaches. In the next subsection, we list desirable features of a DSN and later show how the proposed network provides many of these features.

A. Motivation for a New Architecture

Wesson *et al.* [2] have described two architectures for a DSN. The first is the hierarchical or tree organization and the second is the committee organization whose nodes can send messages to one, some, or all nodes in the network. The hierarchical network has several advantages like constant node degree and easy extensibility. It is not a good choice for a DSN, however, because a faulty node can disconnect an entire subtree. The committee organization does not have this disadvantage but is expensive and is not easily extensible.

It is clear from the above observations that both the committee organization and the tree organization have disadvantages; a combination that uses the merits of both the types of architectures is hence desirable. The flat tree network, referred to earlier, incorporates some of the merits of both these organizations. The nodes in this network are organized as many complete binary trees, the roots of which are completely connected. Fig. 2 shows a flat tree network with 12 nodes. It has some disadvantages, however. For example, integration errors of the lower nodes *accumulate* as the information goes up the hierarchy. One way to overcome this problem is to interconnect nodes in the lower levels of this network.

This motivates our proposal for a class of networks which have a committee organization at each level and an overall hierarchical organization. The versatility of networks with this organization arises from the fact that several topologies could be considered to interconnect nodes at each level. The networks that we propose consists of the flat tree with nodes at each level interconnected as a deBruijn graph. We will show that this class of networks has several advantages such as

- 1) they allow the construction of large networks with bounded degree,
- 2) their diameter of these networks grows only logarithmically with the the number of nodes,
- 3) they admit simple routing schemes, and
- 4) they possess fault tolerant capabilities.




Fig. 3. DG (2, 3).

B. The Proposed Architecture

The proposed DSN is a modified l -level MLN with the top level completely connected and with each of the other levels interconnected as a deBruijn network. The versatility of this organization arises from the fact that several interconnection topologies could be considered to interconnect nodes. Before describing the proposed architecture for DSN, we briefly review the evolution of the deBruijn network and mention its important features.

The use of deBruijn networks as interconnection topologies for fault-tolerant parallel and distributed architectures was first proposed by Pradhan [3], who also proposed fault tolerant VLSI architectures based on this network [5], [13]. Recently, deBruijn networks have gained significant practical importance with the on-going implementation of a 8096 PE deBruijn architecture by JPL for the Galileo project, scheduled for completion by 1995 [13].

An important feature of the deBruijn network is that it can be configured as many useful computational networks in spite of faults. In addition, deBruijn networks have

- 1) a small diameter
- 2) admit simple routing, and
- 3) possess good fault tolerant capabilities.

For a detailed discussion on the aforementioned features of deBruijn networks, see the paper by Samantham and Pradhan [8].

Using graph theoretic notation, the undirected deBruijn network $DG(d, k)$ has $N = d^k$ nodes with diameter k and degree $2d$. We are interested in binary deBruijn networks $DG(2, k)$ which have $N = 2^k$. A node i of the network with the binary address $a_{k-1}a_{k-2} \cdots a_1a_0$ has neighbors:

$$a_{k-2}a_{k-3} \cdots a_1a_0a_{k-1} \quad (i1) \quad (i1)$$

$$a_{k-2}a_{k-3} \cdots a_1a_0\bar{a}_{k-1} \quad (i2) \quad (i2)$$

$$a_0a_{k-1}a_{k-2} \cdots a_2a_1 \quad (i3) \quad (i3)$$

$$\bar{a}_0a_{k-1}a_{k-2} \cdots a_2a_1. \quad (i4) \quad (i4)$$

The address of neighbors $i1$ and $i3$ is obtained by the left shift-end-around operation and the right shift-end-around operation on i respectively—they are called the LR and the RR neighbors of i . The address of nodes $i2$ and $i4$ is obtained by complementing the rightmost bit of $i1$ and the leftmost bit of $i3$ respectively—they are correspondingly called the LRC and the RRC neighbors of i . Fig. 3 shows an eight-node binary deBruijn network with the nodes named with the convention just described.




Fig. 4. MBD with two layers.

The proposed DSN is organized as follows:

- 1) The nodes in the topmost level are called commander nodes. There are four commander nodes that are completely connected.
- 2) The nodes in each of the underlying levels are interconnected as a binary deBruijn network $(l - 1)$.
- 3) Each node X , at level m in the network is connected to two children nodes $\text{app}(X, 1)$ and $\text{app}(X, 0)$ at level $m + 1$ ($m < l - 1$) and is connected to its parent node $\text{str}(X)$ at level $m - 1$.

Henceforth we shall refer to the proposed network as the multi-level binary deBruijn network (MBD). Since the topmost level of the MBD contains 2^2 nodes, it is convenient to assign it level 2. Hence, an l -level MBD has l levels numbered from 2 through $l + 1$. Fig. 4 shows a 2-level MBD—the inter-level connections are shown by dashed lines and the intra-level connections by solid lines. Each node of the MBD has a PE, a clock which maintains real time, an associated sensor which samples the physical variable(s) of interest, and an associated buffer. The PE translates the sensor reading into an abstract estimate, time stamps the estimate with the current time, and places the abstract estimate in the associated buffer. There is also a buffer associated with each link. The PE's connected to the link have access to this buffer. Fig. 5 shows the architectural details of a node of the MBD. (Note: With slight modifications, we could allow for multiple sensors at each node.)

Topological Properties: The following lemmas describe the topological properties of the MBD.

Lemma 1: The number of nodes in MBD with l levels is $4(2^l - 1)$.

Proof: The number of nodes at level m ($2 < m \leq l + 1$),

$$n_m = 2 \times n_{m-1}; \quad n_2 = 4$$

Solving this equation yields the total number of nodes as

$$N = 4(2^L - 1) \quad (1)$$

Lemma 2: The MBD with L levels has degree 7 and diameter $L + 1$.

Proof: The nodes at the top and bottom levels have degree at most 5. Now, consider an internal node in the network. This node is in a deBruijn network and hence has at most 4 neighbors. The same node is also connected to its 2 children nodes and a parent node. Hence a node in the




Fig. 5. Details of node architecture and internode connections.

MBD has degree equal to at most 7. For deriving the diameter of the network, consider the lowermost level in the MBD. This corresponds to $DG(2, l+1)$ with diameter $l+1$. Note that the farthest distance between nodes in the uppermost and lowermost level is only L . Hence, the farthest nodes in the MBD lie in the lowermost level, i.e., the diameter equals $l+1$. From (1), the diameter of the MBD is $O(\log N)$.

Addressing Scheme: Consider a MBD with l levels. The address of a node in this network consists of two parts—

- 1) the level number of the MBD in which it is present. This requires $\lceil \log(l) \rceil$ bits for its representation.
- 2) index of the node in that level. This requires at most $(l+1)$ bits to index a node in any level, because the lowermost level (i.e., level $(l+1)$) contains $2^{(l+1)}$ nodes.

The address of a node in a MBD with l levels, hence needs $\lceil \log(l) \rceil + (l+1)$ bits.

Extensible Issues: To extend a MBD with l levels, we can add the additional nodes at the lowermost level. Thus, extending the network requires a fixed number of interconnections between the new nodes and the nodes at level $(l+1)$ only. Note that the information integration process will not get affected at any other level of the MBD. Additional bits may be needed to address the nodes in the new level.

C. Routing

We show that messages can be routed efficiently in a decentralized manner in the MBD. We first consider routing within a level and then consider routing across levels. To evaluate the routing complexity, we assume that a message takes unit time to traverse a link.

Intra-Level Routing: Routing in the top level takes unit time step since the nodes are completely connected. Routing in a deBruijn network is a well studied problem—we will consider the routing algorithm presented in [4]. In this algorithm, tag bits are appended to the message at the source before routing. These tag bits are used by intermediate nodes to compute the address of the next node in the path. This method assumes that all the nodes in the path are fault-free.




Fig. 6. Type 1 routing tag.

Hence the algorithm will fail if any of the intermediate nodes or links are faulty.

In this section we describe two distributed routing algorithms PATH.1 and PATH.2 in which the address of the next node is computed at the previous node in the path. PATH.1 takes $\Theta(\log N)$ steps in a deBruijn network with N nodes, and PATH.2 takes $O(\log N)$ steps.

Let a binary deBruijn network have $N = 2^k$ nodes and let $S = s_{k-1}s_{k-2} \dots s_1s_0$ be the source node that sends a message to the destination node $D = d_{k-1}d_{k-2} \dots d_1d_0$.

The message consists of the data and the message header. The message header contains a routing tag whose content depends on the type of routing being performed. Two types of routing tags are used—one for normal routing (Type 1) and the other for fault tolerant routing (Type 2).

The Type 1 routing tag contains the source and destination node addresses, a counter (z), and an interlevel routing bit (IRB). The IRB bit is *set* if the source and the destination nodes are in different levels and is *reset* if the nodes are in the same level. The number of *message hops* from the source node to the current node is recorded in the counter, and is used to generate the address of the next node in the path. Fig. 6 shows a Type 1 routing tag.

PATH.1 Algorithm

From the construction of the deBruijn network we know that the source node has the following neighbors— $d_0s_{k-1}s_{k-2} \dots s_1$ and $s_{k-2} \dots s_1s_0d_{k-1}$. Using this property we can now generate two routes by appending successive bits of the destination node to the source address.

Route 1

- $(z = 0) \quad s_{k-1}s_{k-2} \dots s_1s_0 \quad (\text{source})$
- $(z = 1) \quad d_0s_{k-1}s_{k-2} \dots s_1$
- $(z = 2) \quad d_1d_0s_{k-1} \dots s_2$
- \vdots
- $(z = k) \quad d_{k-1}d_{k-2} \dots d_1d_0 \quad (\text{destination})$

Route 2

- $(z = 0) \quad s_{k-1}s_{k-2} \dots s_1s_0 \quad (\text{source})$
- $(z = 1) \quad s_{k-2}s_{k-3} \dots s_0d_{k-1}$
- $(z = 2) \quad s_{k-3} \dots s_0d_{k-1}d_{k-2}$
- \vdots
- $(z = k) \quad d_{k-1}d_{k-2} \dots d_1d_0 \quad (\text{destination})$

Clearly Route 1 and Route 2 take exactly $k = \log N$ steps.

Let $i_{k-1}i_{k-2} \dots i_1i_0$ be the address of the node under consideration. The following steps (executed by each node) describe the PATH.1 algorithm:

- 1) If the label of the node is the same as the destination address in the routing tag, then accept the message.
- 2) Otherwise, check the value of the routing tag counter z . The address of the next node in the path is

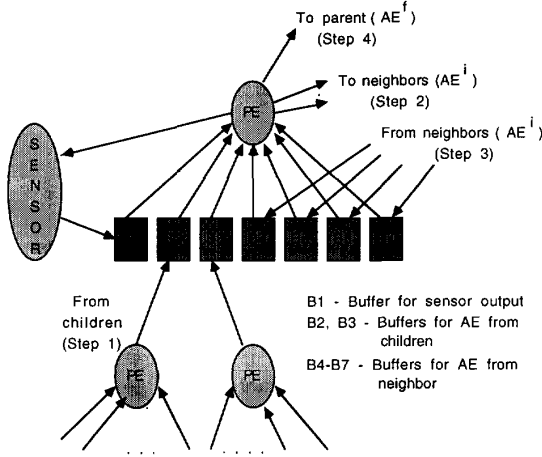


Fig. 11. Information integration process.

integrate them with abstract estimate from local sensor to get AE^i .

Step 2: Send AE^i to neighbors.

Step 3: Receive abstract estimate from neighbors and "compare" with own abstract estimate to compute AE^f . Identify any faulty node in the process.

Step 4: Send AE^f to parent node.

This process of information integration ensures that only the "correct" estimates move up to the commander nodes in the network. Note that the width of the estimates moving upwards is bounded by the width of one of the correct estimates of that level of the MBD. An incorrect estimate would be received by a parent only when the child or the link connecting the two nodes is faulty.

Fig. 11 shows the complete information integration process at a node in the network.

E. Fault Tolerant Issues

In a large network it is unrealistic to expect all the nodes or links along a path to be fault-free at all times. When some nodes or links fail, an alternative path that avoids the faulty node or link must be derived. One of the major advantages of our network over the network proposed in [11] is that abstract estimates can be routed around faults using the interconnections between nodes at the same level.

A node is faulty if it sends an incorrect abstract estimate to its parent or to any of its neighbors. Link faults can be detected if a node does not receive the abstract estimate of its neighbor during the comparison step. When a node (a node failure is assumed to be equivalent to the failure of all links associated with it) or link failure is detected, any of the following actions can be taken.

- 1) The fault can be ignored during the integration process. After integration is complete and abstract estimates have been sent to the upper level of the MBD, the node which

detected this fault can run a diagnostic algorithm on the faulty node or link after isolating it.

- 2) If the faulty node or link is in the path of an abstract estimate transmitted towards its destination, this abstract estimate can be re-routed around the fault to the destination. After the integration process is complete the node which detected the fault can run a diagnostic algorithm on the faulty component.

The MBD network provides fault tolerance by taking both of the remedial actions mentioned previously.

Suppose a node X , with children Y and Z , is faulty. In the flat tree [11] network the subtree rooted at X is unusable. In the MBD network, however, the abstract estimates of Y and Z are also read by the neighbors of Y and Z . Thus the abstract estimates of Y and Z get factored into the final abstract estimates produced by the neighbors of Y and Z . Hence the subtree rooted at X does not become unusable—only the faulty node is unusable. Moreover, X is identified as a faulty node during the comparison step because its abstract estimate (which it sends to its neighbors) may not contain the physical value.

If action 2) is taken by the neighbor of the faulty node, then it must reroute the abstract estimate received, around the faulty node to the destination. This means that the destination node must wait for more time to receive the abstract estimate, because additional hops may be required for rerouting the message. This requires that the value of γ (maximum difference in time that a node can tolerate between intervals that can be integrated—please see the next section on clock synchronization) be increased to maintain the "near synchronous" behavior of the sensor network. Note that by increasing the value of γ , the network would tolerate single node/link fault but the process of sensor integration would be slowed down. Samantham and Pradhan [8] mention that four additional hops are enough to avoid a single node fault in a binary deBruijn network. Since the nodes in every level (except the top level) in the MBD are arranged in a binary deBruijn network, the value of γ will have to be increased by four time units.

In the remaining part of this section, we show one way of avoiding a single node fault using exactly four hops, when routing in any level of the MBD except the topmost level. Let $s_{k-1}s_{k-2} \dots s_1s_0$ be the source node and $d_{k-1}d_{k-2} \dots d_1d_0$ be the destination node. Application of the PATH.1 algorithm yields the following path:

$$(z = 0) \quad s_{k-1}s_{k-2} \dots s_1s_0 \text{ (source)}$$

$$(z = 1) \quad s_{k-2}s_{k-3} \dots s_0d_{k-1}$$

$$(z = 2) \quad s_{k-3} \dots s_0d_{k-1}d_{k-2}$$

$$\vdots$$

$$(z = m-1) \quad s_{k-m}s_{k-m-1} \dots s_0d_{k-1} \dots d_{k-m+1} \quad (11)$$

$$(z = m) \quad s_{k-m-1}s_{k-m-2} \dots s_0d_{k-1} \dots d_{k-m+1}d_{k-m} \quad (12)$$

$$(z = m+1) \quad s_{k-m-2} \dots s_0d_{k-1} \dots d_{k-m}d_{k-m-1} \quad (13)$$

$$\vdots$$

$$(z = k) \quad d_{k-1}d_{k-2} \dots d_1d_0 \text{ (destination)}$$

Assume that either node $i2$ or the link between $i1$ and $i2$ has failed. We now show an alternative route (reroute) between $i1$ (rerouting source) and $i3$ (rerouting destination) that takes only four additional hops.




Fig. 9. Integration step.

The proof is by induction. The base case for $n = 3$ is straightforward to prove by enumeration (see Fig. 9). Consider p intervals and assume that the lemma holds for less than p intervals. Consider the first $(p - 1)$ of the p sorted intervals. Since we know that there can be at most one fault, either all the $(p - 1)$ intervals intersect or exactly $(p - 2)$ intervals intersect.

Case 1: Exactly $(p - 2)$ intervals intersect (there is one fault among the $(p - 1)$ intervals): by the induction hypothesis, there are at most two $(p - 2)$ interval intersections— $A (A_l, A_u)$ and $B (B_l, B_u)$; let $A_l < B_l$. Further, A and B are non-intersecting (i.e., $A_u < B_l$)—otherwise they would have formed a $(p - 1)$ intersecting interval. Since there can be one fault at most, the p th interval has to be correct and has to intersect with B giving rise to one $(p - 1)$ intersecting interval.

Case 2: All $(p - 1)$ intervals intersect: Let the intersecting interval be $C (C_l, C_u)$. By the induction hypothesis there are at most two $(p - 2)$ interval intersections $D (D_l, D_u)$ and $E (E_l, E_u)$ and these overlap—i.e., $D_l \leq E_l$, $E_l < D_u$. Further, since C is the intersection of D and E , $C_u = \min(D_u, E_u)$ and $C_l = E_l = (p - 1)l$, the lower bound of the $(p - 1)$ st interval. Now the p th interval can intersect with some or all the intervals C , D , and E . Several cases arise:

Case 2a): The p th interval (p_l, p_u) intersects C —the p intersecting interval is then $(p_l, \min(p_u, C_u))$, i.e., $(p_l, \min(D_u, E_u, p_u))$. The p th interval could intersect either D (if $D_u > E_u$) or E (if $E_u > D_u$) but not both to form another $(p - 1)$ intersection.

Case 2b): The p th interval (p_l, p_u) does not intersect C . Hence, $p_l > \min(D_u, E_u)$. The p th interval intersects either D (if $D_u > E_u$) or E (if $E_u > D_u$) but not both both to form a $(p - 1)$ intersection.

The $(p - 1)$ intersecting intervals arising from Case 2 are the interval C and at most one more from Cases 2(a) and 2(b). The lemma then follows from the above cases.

A direct consequence of Lemma 3 is the following theorem. Using the theorem, the search for a faulty node is narrowed down to at most nodes for each fault.

Theorem 1: Given a set of n intervals containing at most one faulty interval,

- 1) there is no faulty interval if there is no $n - 1$ -interval intersection,
- 2) the interval not intersecting with an $n - 1$ -interval intersection is faulty if there is exactly one $n - 1$ -interval intersection, and
- 3) there are two potentially faulty intervals if there are two $n - 1$ -interval intersections one of which is incorrect.

In case 3), the two potentially faulty nodes can be traced by taking the set difference of the interval names that belong to each $(n - 1)$ interval intersection.




Fig. 10. Comparison step.

We now describe the information integration process. For convenience, we will refer to the information integration of abstract estimates between distinct levels as “integration” and refer to the information integration within a level as “comparison.”

Abstract estimates move upward from the leaf nodes to the commander nodes. Every non-leaf node of the network combines the abstract estimates of its two children and the local sensor (sensor associated with this PE) to arrive at a new abstract estimate (AE^i). This step is called the “integration” step.

In the integration step, we assume that at most one of the three (local sensor and 2 children) received abstract sensor estimates is incorrect. The new abstract estimate is found from the three cases (refer Fig. 9) that could arise (Theorem 1). If there are two 2 interval intersections, then the smallest interval containing these intervals forms the new abstract estimate. It can also be shown [10] that this new estimate is at most as wide as one of the input abstract estimates.

Next, each node sends its AE^i to all its neighbors. When a node receives AE^i s from its neighbors, it combines them to arrive at a new estimate AE^f . This step is called the “comparison” step and the algorithm used to combine the estimates is similar to the one described for the integration step. In this step, however, a node combines 3, 4, or 5 estimates depending on the number of its neighbors (2, 3, or 4 respectively).

Since the MBD can tolerate at most one fault (node or link) per level, one of the estimates received from a neighbor could be incorrect. Hence, when a node receives i ($i = 3, 4$ or 5) intervals in the comparison step, it chooses the smallest interval containing all (which is at most two as shown in lemma 3) the $i - 1$ -interval intersections as the output. The width of this abstract estimate is again at most as wide as one of the input correct intervals.

Fig. 10 shows the comparison process in a node of the network.

If there are two $i - 1$ -interval intersections in the comparison step, then we know that there exists an incorrect interval. Identifying the faulty node which sent this incorrect interval requires the diagnostic testing of at most two nodes as we showed in Theorem 1. Once a node has been identified as faulty, appropriate action can be taken so as to either “repair” the faulty node or replace it and notify the parent and children of the faulty node. In this paper, we are not concerned with the problems of identifying the cause of faulty behavior and attempting to rectify that cause.

The following steps summarize the process of information integration:

Step 1: Receive abstract estimates from children and




Fig. 7. PATH.1 route between 011 and 101.

$d_z i_{k-1} i_{k-2} \dots i_1$ (Route 1) OR $i_{k-2} i_{k-3} \dots i_1 i_0 d_{k-z-1}$ (Route 2).

- 3) Increment the counter z and route the message to the next node.

Fig. 7 shows a path from node 011 to node 101 in a DG(2, 3) network using Route 2 of the PATH.1 algorithm.

PATH.2 Algorithm

PATH.2 algorithm routes the message along the shortest path between the source and destination nodes. To find the shortest path, we treat the node addresses as binary strings and use a string matching algorithm described in [1].

Find the largest x such that $s_{x-1} s_{x-2} \dots s_1 s_0 = d_{k-1} d_{k-2} \dots d_{k-x}$, and the largest y such that $s_{k-1} s_{k-2} \dots s_{k-y} = d_{y-1} d_{y-2} \dots d_1 d_0$. We can compute x and y in $O(k)$, i.e., $O(\log N)$ time. The following three cases arise depending on the relationship between x and y .

Case 1: ($x > y$)—the shortest path is given by the following sequence of nodes:

($z = 0$) $s_{k-1} s_{k-2} \dots s_1 s_0$ (source)
 ($z = 1$) $s_{k-2} s_{k-3} \dots s_0 d_{k-x-1}$
 ($z = 2$) $s_{k-3} \dots s_0 d_{k-x-1} d_{k-x-2}$
 :
 ($z = k - x$) $s_{x-1} s_{x-2} \dots d_1 d_0$ (destination)

The destination $s_{x-1} s_{x-2} \dots d_1 d_0 = d_{k-1} d_{k-2} \dots d_1 d_0$ is reached after $(k - x)$ steps.

Case 2: ($x < y$)—the shortest path is given by the following sequence of nodes:

($z = 0$) $s_{k-1} s_{k-2} \dots s_1 s_0$ (source)
 ($z = 1$) $d_y s_{k-1} \dots s_2 s_1$
 ($z = 2$) $d_{y+1} d_y s_{k-1} \dots s_2$
 :
 ($z = k - y$) $d_{k-1} d_{k-2} \dots s_{k-y+2} s_{k-y+1}$ (destination)

The destination is reached after $(k - y)$ steps.

Case 3: ($x = y$)—choose either of the above routings to obtain the shortest path.

For this algorithm, the routing tag counter is initiated to $k - x$ or $k - y$.

Fig. 8 shows the shortest path between nodes 011 and 101 in a DG(2, 3) network using the PATH.2 algorithm. In this case $x = 1$ and $y = 2$; hence, the shortest path is of length 1. Using the PATH.1 algorithm yields a path length of 3.

The following steps describe the PATH.2 algorithm (as executed by node i).

In practice, both of the algorithms would be implemented using shift/and complement operations at each step. this would




Fig. 8. PATH.2 route between 011 and 101.

obviate the need for the routing tag counter thus reducing the message header requirements; further, an expensive increment operation is replaced by a shift/complement operation. Note that the value of x and y need not be computed by all nodes in the path. Instead, the value of x or y can be transmitted in the message header.

Inter-Level Routing: Let the source (S) and destination (D) nodes be at levels L and $L - X$ respectively. At the source the inter-level routing (IRB) bit is set to "1" to indicate that the source and destination nodes are in different levels. Further, when the IRB bit is set, the routing tag counter is not incremented in order to maintain a proper value of the counter for intra-level routing following the inter-level routing.

The source node S first routes the message to its parent $str(S)$. This procedure is repeated recursively till the message is received by a node at the same level as the destination node D . The IRB bit is reset to "0" now, and the source address is replaced by the address of the node that received the message. The message can be then routed to the destination using PATH.1 or PATH.2 algorithm.

When the destination is at a higher level than the source, routing can be similarly done by using $app()$ to generate the address of the next node in the path till the message reaches the same level as that of the destination node. The message can then be routed using PATH.1 or PATH.2 algorithm. Note that messages in the MBD are usually routed from higher to lower levels only since the final integration is done by the commands.

D. Information Integration

In this section, we describe the process of information integration in the MBD. The idea behind the integration is to a) keep the communication requirements small—this is done by communicating the abstract estimate as a *single* interval and b) maintain accuracy by ensuring that the physical values of interest is *always* contained in the abstract estimate.

Since the deBruijn network has a connectivity of 2, the MBD can tolerate at most one node fault or link fault per level (except at the topmost level which is fully connected). We first prove some results related to fault tolerance when abstract estimates (or intervals) are to be integrated in the presence of faults in the network.

Lemma 3: Consider n ($n \geq 3$) intervals of which at most one can be faulty. Then there can be at most two $(n - 1)$ distinct interval intersections among these n intervals.

Proof: Without loss of generality, assume that the n intervals (i_l, i_u) ($1 \leq i \leq n$) are sorted in increasing order by their lower bounds.

Type	IRB	Source Address	Destination Address	Routing Tag Counter	Source Address (i_1)	Destination Address (i_3)	RC	RRB
------	-----	----------------	---------------------	---------------------	--------------------------	-------------------------------	----	-----

Fig. 12. Type 2 routing tag.




Fig. 13. Fault tolerant routing between 001 and 110 when 011 is faulty.

When i_1 receives a message (consisting of the abstract estimate and the Type 1 tag), it appends four fields to the Type 1 tag which enable rerouting of the message—(1) source address (i_1), (2) destination address (i_3), (3) reroute counter (RC), and (4) rerouting bit (RRB). We shall refer to the tag, formed by appending reroute fields to the Type 1 tag, as Type 2 tag. Figure 12 shows a Type 2 tag. To initiate rerouting, i_1 increments z and sets $RRB = "1"$. When $RRB = "1"$, a node does not increment z ; instead it uses RC to compute the address of the next node in the reroute. When the message reaches i_3 , i_3 removes the reroute fields from the tag, and increments z . Routing from i_3 then proceeds normally using PATH.1 or PATH.2 algorithms.

The alternative route between i_1 and i_3 is shown here:

$$\begin{aligned}
 (z = m-1) & s_{k-m} s_{k-m-1} \cdots s_0 d_{k-1} \cdots d_{k-m+1} \\
 (z = m) & s_{k-m-1} s_{k-m-2} \cdots s_0 d_{k-1} \cdots d_{k-m+1} \bar{d}_{k-m} \\
 (z = m) & s_{k-m-2} s_{k-m-3} \cdots s_0 d_{k-1} \cdots \\
 & d_{k-m+1} \bar{d}_{k-m} d_{k-m-1} \\
 (z = m) & d_{k-m-1} s_{k-m-2} \cdots s_0 d_{k-1} \cdots d_{k-m+1} \bar{d}_{k-m} \\
 (z = m) & d_{k-m} d_{k-m-1} \cdots s_0 d_{k-1} \cdots d_{k-m+1} \\
 (z = m) & d_{k-m-1} s_{k-m-2} \cdots s_0 d_{k-1} \cdots d_{k-m+1} d_{k-m} \\
 (z = m+1) & s_{k-m-2} s_{k-m-3} \cdots s_0 d_{k-1} \cdots \\
 & d_{k-m+1} d_{k-m} d_{k-m-1}.
 \end{aligned}$$

The above route takes 6 steps—only 4 more than the normal route between i_1 and i_3 . Figure 13 shows fault tolerant routing in a DG(2, 3) (level 3) between nodes 001 and 110 when the node 011 is faulty. This alternative route can be chosen when a faulty node is encountered in the path to the destination node. Hence, the routing algorithms given earlier can be easily adapted to take the alternative path in case of faults. This rerouting algorithm is also more adaptive to faults than the one presented in [8] since our algorithm does not require that the presence of a fault be known to the source node as the other algorithm does.

Finally, since the network can sustain one node or link fault at every level, the MBD network with l levels and $N = 2(2^l - 1)$ nodes can sustain l , i.e., approximately $\log N$, node or link faults.

III. CLOCK SYNCHRONIZATION ISSUES

So far we have assumed that any two abstract estimates can be integrated. In reality, since the sensor outputs typically change as a function of time, only estimates that are temporally

“close to each other” must be integrated if meaningful results are desired. This is achieved by time-stamping each estimate. The condition under which two estimates may be integrated is given at the end of the next subsection. In a distributed environment such as ours, there is no central synchronized clock which regulates the activities of each node. Instead, each node is under the control of its own clock. Since the sensor responds to real-time events, it is convenient for the clock to provide the real, i.e., physical time. Further, since the estimates from different sensors have to be integrated, it is necessary for the time provided by the clocks of the sensors to be “close to each other.” The clock at each node may not be accurate because of a variety of reasons such as clock shift, change in temperature, etc. Each clock must therefore synchronize with a more accurate clock. We assume the existence of a central time server which when requested for the time at t , provides the time $C(t)$. The PE’s in our DSN spatially lie within tens of feet from each other, hence the existence of a single time server for the clocks on all the PE’s can be assumed. The central time server itself periodically synchronizes with a universal time server, which is always accurate and lies outside our environment.

The following definitions are used.

- ϵ Maximum allowable deviation in time of a clock on a PE.
- α Maximum allowable deviation in time of a clock on the central time server.
- κ Maximum allowable drift rate in time of a clock on a PE.
- σ Maximum allowable drift rate of the clock in the central time server.
- δ Channel transmission delay. •
- ξ Delay in receiving the message sent by the central time server to any PE.
- γ Maximum difference in time that a node can tolerate between intervals that can be integrated.

We use the clock model described in [11] to synchronize the clocks in the MBD. We summarize the basic results of the model in the next two subsections.

A. Clock Behavior and Synchronization

Let $C_p(t)$ be the time provided by the clock on PE p at time t (t is the time provided by a universal time server). We assume that the clocks run continuously rather than in discrete “ticks.” Hence, $(dC_p(t))/dt$ denotes the rate at which the clock is running at time t . We also assume that this rate is nonnegative; hence, the time on the clocks monotonically increase.

We now state the conditions on the clocks for proper synchronization.

Clock Condition 1: The deviation in time of each clock is bounded, i.e., for PE p , there exists $\epsilon_p \ll 1$ and $\alpha \ll 1$ such that

$$|t - C_p(t)| \leq \epsilon_p \quad (2a)$$

$$|t - C(t)| \leq \alpha. \quad (2b)$$

Clock Condition 2: Between synchronizations, the drift rate of the clock is bounded, i.e., for PE p , there exists $\kappa_p \ll 1$ and $\sigma \ll 1$ such that

$$\left| \frac{dC_p(t)}{dt} - 1 \right| \leq \kappa_p \quad (3a)$$

$$\left| \frac{dC(t)}{dt} - 1 \right| \leq \sigma.$$

Clock Condition 3: The clock on each of the PE's and the central time server advance monotonically.

For simplicity, we assume that ϵ_p and κ_p is the same for all PE's, i.e., $\epsilon_p = \epsilon$ and $\kappa_p = \kappa$. From Clock Condition 1 we have,

Synchronization Bound: If p and q are two PE's then

$$|C_p(t) - C_q(t)| \leq 2\epsilon. \quad (4)$$

Let δ_{\min} and δ_{\max} be the minimum and maximum values of delay for a message sent by a PE to its neighbor. Let γ be the maximum difference in time that a node can tolerate between intervals that can be integrated. Note that the value of γ will depend on the longest path between leaf nodes and commander nodes, which is equal to L in a MBD with L levels.

The following lemma and theorem state precisely the conditions for combining abstract sensor estimates which are temporally "close to each other." The discussion in this section follows closely the discussion presented in [11]. We therefore state all the results without proof. The interested reader is referred to [11].

Lemma 4: Let a message be received by PE p at $C_p(t)$. Then this message was sent in the interval $(C_p(t) - 2\epsilon - \delta_{\max}, C_p(t) + 2\epsilon - \delta_{\min})$.

The time stamp of an abstract estimate may not belong to the interval given above if the channel is faulty. The following theorem gives the condition under which estimates are "temporally close" and may be integrated.

Theorem 2: Let the three proper abstract sensor estimates I_1 , I_2 and I_3 be received by PE p at times

$$C_p(t1) < C_p(t2) < C_p(t3)$$

respectively. Then I_i ($i = 2, 3$) can be integrated, iff

$$(C_p(ti) - C_p(t1) + 4\epsilon + \delta_{\max} - \delta_{\min}) \leq \gamma.$$

Since the clocks on the central time server and each of the PE's drift, they have to be periodically reset. We now state a bound on the time period between synchronizations. Let T_s be the time period of synchronizations of the central time server, and let T_c be the time period between synchronizations of the PE p .

The central time server synchronizes itself every T_s seconds with a perfect universal time server which exists outside the environment of the DSN. The central time server also synchronizes the clock on a PE every T_c seconds.

Let ξ_{\min} and ξ_{\max} be the minimum and maximum of the delay in receiving the message sent by the central time server

to any PE. Also, let T_s^i and T_c^i be the periods corresponding to T_s and T_c as observed by the central time server. The bounds on these observed periods can be shown to be [11]:

Theorem 3: The time period as observed by the central time server between synchronizations of the central time server is bounded by

$$T_s^i \leq \alpha \left(\frac{1}{\sigma} - 1 \right).$$

Theorem 4: The time period as observed by the central time server between synchronizations of the clock on a PE is bounded by

$$T_c^i \leq \frac{\epsilon - 2\alpha - (\xi_{\max} - \xi_{\min})}{\kappa} - \alpha.$$

IV. COMPARISON WITH COMMERCIAL SENSOR NETWORKS

In previous sections we described the topological and fault-tolerant properties of the MBD. We also saw that routing in the MBD is very simple. Nearly all sensor networks used in process control industries now are based on the bus or ring systems. With the need for large scale DSN's (combined with the need for high sampling rates), the common data path in the bus and the high diameter/low connectivity of the ring make them both unsuitable to support the communication required among the nodes of the DSN. Further, since data fusion is by nature hierarchical, a hierarchical interconnection network would be most suitable for such a function.

Table I compares the topological, routing and fault-tolerant properties of the MBD with the bus and ring networks. It can be seen that the MBD is a good alternative to the bus and ring networks.

V. CONCLUSION

The effective use of multiple sensor systems requires the solution of various problems relating to sensor models, the architecture of the sensor network, the integration of information at each node of the network, the cost of information transmission, and the fault tolerance of the network. The integration of information in real time requires the clocks at each of the nodes be synchronized. Synchronization of clocks is a nontrivial task in such distributed sensor networks. In an earlier paper [11], some issues related to the architecture of DSN's, information integration, and clock synchronization had been addressed. This paper extends the study by considering a more sophisticated architecture for DSN's which has a number of advantages including the ability to tolerate single node or link faults at each level.

Since our focus has been primarily on computational issues, we have chosen to represent sensor output information by Marzullo's simple and elegant model which is based on real valued intervals. We have also used a relatively simple information integration algorithm. We are aware that sensor modeling is itself a very detailed area of study [7] and that very sophisticated methods exist for information integration. We have also assumed that the output of each sensor is a physical value. The above discussion and results easily extend

TABLE I

	Bus	Ring	MBD*
Communication	$O(N)$	$O(N)$	$O(\log N)$
Diameter	$O(1)$	$O(N)$	$O(\log N)$
Degree	1	2	7
Routing	Simple	Simple	Simple
Fault-tolerance	1	2	$O(\log N)$ †
Cost (number of links)	1	$N - 1$	$< 3.5N$

MBD* multilevel deBruijn network.

† Assuming of one fault per level.

to the case when the output of a sensor is a vector rather than a single value.

This study could be extended in several directions. A straightforward extension is to assign weights to the abstract estimates produced as a function of its level in the hierarchy. We also plan to investigate more sophisticated fault tolerant strategies for the deBruijn network than the scheme presented here. A future goal of our project is to investigate the computation and communication requirements of more sophisticated integration algorithms executing on large scale DSN's.

ACKNOWLEDGMENT

The authors would like to thank Dr. D. K. Pradhan for his discussions of the paper and the referees for their comments. The authors also like to thank Dr. Madan for his comments on this paper.

REFERENCES

- [1] D.E. Knuth, J.H. Morris, and V.R. Pratt, "Fast pattern matching in strings," *SIAM J. Computing*, vol. 6, pp. 323-350, 1977.
- [2] R. Wesson, *et al.*, "Network structures for distributed situation assessment," *IEEE Trans. Syst., Man, Cybern.*, pp. 5-23, Jan. 1981.
- [3] D. K. Pradhan, "Interconnection topologies for fault-tolerant parallel and distributed architectures," in *Proc. 10th Int. Conf. Parallel Processing*, Aug. 1981, pp. 238-242.
- [4] D. K. Pradhan and S. M. Reddy, "A fault-tolerant communication architecture for distributed systems," *IEEE Trans. Comput.*, vol. c-31, no. 9, 1982.
- [5] D. K. Pradhan, "Dynamically restructurable fault-tolerant processor network architecture," *IEEE Trans. Comput.*, vol. c-34, May 1985.
- [6] A.-H. Esfahanian and S. L. Hakimi, "Fault-tolerant routing in deBruijn communication networks," *IEEE Trans. Comput.*, vol. c-34, 1985.
- [7] H. F. Durrant-Whyte, "Sensor models and multisensor integration," *Int. J. Robot. Res.*, vol. 7, no. 6, 1988.
- [8] M. R. Samantham and D. K. Pradhan, "The deBruijn multiprocessor network: A versatile parallel processing and sorting network for VLSI," *IEEE Trans. Comput.*, vol. 38, 1989.
- [9] R. C. Luo and M. G. Kay, "Multisensor integration and fusion in intelligent systems," *IEEE Trans. Syst., Man, Cybern.*, vol. 19, pp. 901-927, Sept/Oct., 1989.
- [10] K. Marzullo, "Tolerating failures of continuous-valued sensors," *ACM Trans. Comput. Syst.*, vol. 4, pp. 284-304, Nov. 1990.
- [11] D. N. Jayasimha, S. S. Iyengar, and R. L. Kashyap, "Information integration and synchronization in distributed sensor networks," *IEEE Trans. Syst., Man, Cybern.*, vol. 21, pp. 1032-1043, Sept. 1991.
- [12] L. Prasad, S. S. Iyengar, R. L. Kashyap, and R. N. Madan, "Functional characterization of sensor integration in Distributed Sensor Networks," *IEEE Trans., SMC*, vol. 21, Sept/Oct. 1991.
- [13] D. K. Pradhan, "Fault-tolerant VLSI architectures based on deBruijn graphs (Galileo in the Mid Nineties)," *DIMACS Series in Discrete Mathematics and Theoretical Comput. Sci.*, vol. 5, 1991.



S. Sitharama Iyengar (M'88-SM'90) is the chairman of the Computer Science Department and Professor of Computer Science at Louisiana State University. He has directed LSU's Robotics Research Laboratory since its inception in 1986. He has been actively involved with research in high-performance algorithms and data structures since receiving the Ph.D. degree in 1974, and has directed more than 18 Ph.D. dissertations at LSU. He has served as principal investigator on research projects supported by the Office of Naval Research, the National Aeronautics and Space Administration, the National Science Foundation/Laser Program, the California Institute of Technology's Jet Propulsion Laboratory, the Department of Navy-NORDA, the Department of Energy (through Oak Ridge National Laboratory, Tennessee), the LEQFS-Board of Regents, and Apple Computers.

He has edited a two-volume tutorial on autonomous Mobile Robots and has edited two other books and more than 150 publications—including 90 archival journal papers in areas of high-performance parallel and distributed algorithms and data structure for image processing and pattern recognition, autonomous navigation, and distributed sensor networks. Iyengar was a visiting professor (fellow) at JPL, the Oak Ridge National Laboratory, and the Indian Institute of Science. He is also an Association for Computing Machinery national lecturer, a series editor for *Neuro Computing of Complex Systems*, and area editor for the *Journal of Computer Science and Information*. He has served as guest editor for the *IEEE TRANSACTIONS ON SOFTWARE ENGINEERING* (1988); *Computer* magazine (1989); the *IEEE TRANSACTIONS ON SYSTEM, MAN, AND CYBERNETICS*; the *IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING*; and the *Journal of Computers and Electrical Engineering*.

Dr. Iyengar was awarded the Phi Delta Kappa Research Award of Distinction at LSU in 1989, won the Best Teacher Award in 1978, and received the Williams Evans Fellowship from the University of Otago, New Zealand, in 1991.



D. N. Jayasimha (M'91) received the bachelor's degree in electronics engineering from the University Visvesvaraya College of Engineering, Bangalore, India, the master's degree in computer science from the Indian Institute of Science, Bangalore, India, and the Ph.D. in computer science from the University of Illinois, Urbana.

He is presently working at NASA Lewis Research Center as a Visiting Scientist. Since 1988 he has been an Assistant Professor at the Department of Computer and Information Science, The Ohio State University, Columbus, OH. His research interests are in communication and synchronization aspects of parallel computing, parallel architectures, and distributed sensor networks.

He is a member of ACM and Sigma Xi.



Deepak S. Nadig was born in Bangalore, India, in 1968. He received the B.Tech degree in Electrical Engineering from Indian Institute of Technology, Bombay, India in 1989.

From 1989 to 1991, he worked for Tata Consultancy Services, Madras, India, where he was a member of the Building Blocks project for IBM Research Labs, Germany. Since 1991 he has been a graduate student in the Department of Computer Science and a research assistant in the Robotics Research Laboratory at Louisiana State University.

His research interests include distributed systems, fault-tolerant computing, and computer networks.