Correspondence

Creating Disjoint Paths in Gamma Interconnection Networks

Nian-Feng Tzeng, Po-Jen Chuang, and Chwan-Hwa Wu

Abstract— **The Gamma interconnection network (GIN) is composed of 3×3 basic building blocks, with interconnecting patterns between stages following the plus-minus-2^i functions. In this paper, we consider modifications to the GIN by altering the interconnecting patterns between stages so as to achieve high terminal reliability between any source-destination pair, resulting in the reliable GIN (REGIN). A type of REGIN's ensures totally disjoint paths in existence from any source to any destination, thereby capable of tolerating an arbitrary single fault. If several building blocks (i.e., 3×3 switches) are fabricated in one chip with very large scale integrated (VLSI) technology, the layout area and the pin count are less for the REGIN than for its GIN counterpart as a result of the change in the interconnecting patterns, giving rise to potential cost reduction. The terminal reliability of the REGIN is derived and compared with that of a compatible GIN. In addition, the performance of the REGIN is evaluated using simulation.**

Index Terms—**Fault tolerance, Gamma interconnection networks, network performance, redundant paths, routing algorithms, terminal reliability.**

I. Introduction

The multistage interconnection network (MIN) can be employed in a multiprocessor to connect processors and memory modules together, using multiple stages of small crossbar switches of a fixed size. The Gamma interconnection network (GIN) [1], like the augmented data manipulator [2] and the inverse augmented data manipulator [3], is one type of MIN composed of 3×3 switches, with interconnecting patterns between stages following the plus-minus-2^i functions. While the GIN offers multiple paths from any source to most of the destinations, there is only a unique path between a source S and a destination D when S equals D, and all the multiple paths existing between certain (S, D) pairs in the GIN share a single common route for many stages before starting to fork, making terminal reliability between such a pair unsatisfactorily low. In Fig. 1, for example, the only two paths between (6, 14) share the same route until reaching stage 3, where the two paths fork, as marked.

A MIN with redundant paths is desirable for high-performance multiprocessor systems, because it can tolerate faults by routing requests over alternative paths, achieving enhanced reliability. As the GIN is not a fault-tolerant network despite its relatively high hardware complexity, one might wonder if the GIN can be made fault-tolerant without increasing its hardware complexity (or even with its cost reduced, if possible). This work aims to find such a profitable design. It should be noted that an interesting variation of the Gamma network, called the Monogamma network, was considered by Raghavendra and Parker [4]. The Monogamma network has the same hardware complexity as the Gamma network, but it can provide multiple paths between every source-destination pair. Unfortunately, the multiple paths between a communication pair in such a network are not necessarily disjoint and a certain single internal switch failure could make the communication from a source to a destination impossible.

In this correspondence, we consider modifications to the GIN by altering the interconnecting patterns between stages so that a modified GIN provides alternative paths between every (S, D) pair, including the case of $S = D$. Such a network is referred to as a reliable **GIN** (REGIN), as it exhibits high terminal reliability between every communication pair. Furthermore, one type of REGIN's has totally disjoint paths from any source to any destination, making it possible to tolerate an arbitrary single fault. In addition, the REGIN enjoys potential cost reduction, if several building blocks (i.e., 3×3 switches) are fabricated in one chip with VLSI technology, because the pin count and layout area are less for the REGIN than for its GIN counterpart.

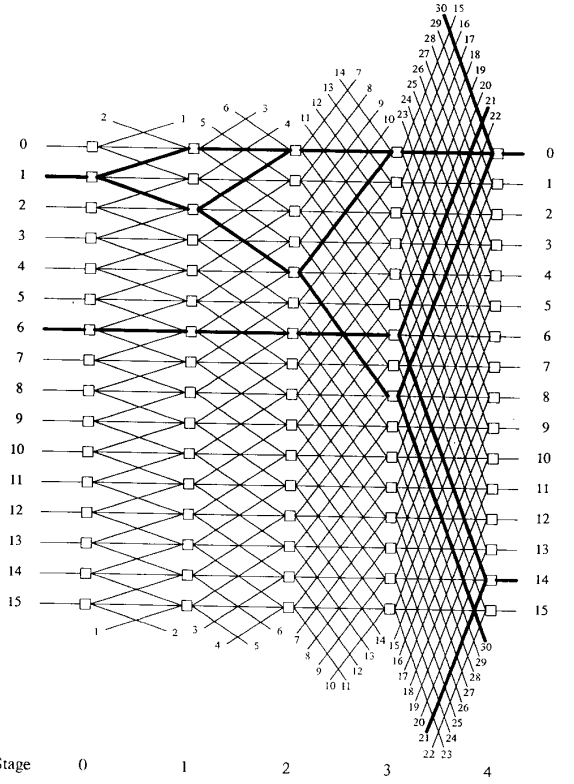

Fig. 1. A Gamma interconnection network with size 16. (Bold lines indicate alternative paths between (1, 0) and between (6, 14).

Manuscript received September 6, 1991; revised July 16, 1992 and February 4, 1993. This work was supported by NSF under Grant MIP-8807761. A preliminary version of this work was presented at the 21st International Symposium on Fault-Tolerant Computing, June 1991.

N.-F. Tzeng is with the Center for Advanced Computer Studies, University of Southwestern Louisiana, Lafayette, LA 70504.

P.-J. Chung was with the Center for Advanced Computer Studies, University of Southwestern Louisiana, Lafayette, LA 70504. He is now with the Department of Electrical Engineering, Tamkang University, Taipei, Taiwan, R.O.C.

C.-H. Wu is with the Electrical Engineering Department, Auburn University, Auburn, AL 36849.

IEEE Log Number 9205227.

TABLE I
THE NUMBER OF ALTERNATIVE PATHS FOR EVERY TAGIN A NETWORK WITH SIZE 16

Tag	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
GIN_4	1	5	4	7	3	8	5	7	2	7	5	8	3	7	4	5
$MGIN_4$	7	6	8	6	6	4	4	2	2	2	4	4	6	6	8	6
$REGIN_4^3$	4	6	6	6	4	6	4	4	2	5	5	6	5	7	6	5
$REGIN_4^2$	7	7	6	4	4	5	5	4	3	4	3	4	4	7	7	7

(The number of total paths amounts to $3^4 = 81$ in every network.)

A. Gamma and Monogamma Interconnection Networks

A size $N = 2^n$ GIN consists of $n + 1$ stages, with each stage involving N switches [1]. Every switch in intermediate stages is a 3×3 crossbar, whereas a switch in the first stage and a switch in the last stage are of sizes 1×3 and 3×1, respectively. The N sources are numbered from 0 to $N - 1$, starting with the top one; so are the N switches in each stage and the N destinations. The stages are numbered from 0 to n rightwards. Connecting patterns between stages are based on the plus-minus-2^i functions, namely, the j^{th} switch in stage i, $0 \leq i < n$, is connected to three switches in stage $i + 1$ according to the functions: $f_i^p(j) = j + 2^i (mod\ N)$, $f_i^m(j) = j - 2^i (mod N)$, and $f_i^s(j) = j$. The three functions define the lower, the upper, and the straight connections originating from switch j in stage i. A GIN with size 16 is depicted in Fig. 1.

A request from source S in a GIN of size $N = 2^n$ (denoted by GIN_n) can be routed to its destination D under the guidance of an n-digit tag, which represents the difference (modulo N) between D and S. Each tag digit can be 1, 0, or -1, and digit d_i is used at stage i, in a way that the lower (or upper) connection is taken if d_i equals 1 (or -1), and the straight connection is taken when d_i is 0. For the sake of convenience, -1 is denoted by $\bar{1}$. The GIN makes use of the binary fully redundant number system to represent each tag $T = D - S(\mathrm{mod}\ N)$. With this number system, a tag $T \neq 0$ has multiple representations, corresponding to multiple paths in the GIN. As an example, $000\bar{1}$, $00\bar{1}1$, $0\bar{1}11$, and $\bar{1}111$ all represent the same value -1, and they correspond to four possible paths, say, from source 1 to destination 0 illustrated in Fig. 1.

For a given $N = 2^n$, there are 3^n different tag representations, since the tag consists of n digits. The tag values lie in the range of $-(N - 1)$ to $(N - 1)$, and the values $(i - N)$ and i are equivalent. In Fig. 1, for example, $T = -1 = (15 - 16)$ has an equivalent tag value of 15, which can be denoted by 1111; the path corresponding to this tag value from source 1 to destination 0 is also shown in the figure. The total number of different paths (i.e., representations) for a tag value T in GIN_n can be calculated by a recurrence relation [1]:

$$\Phi_n(T) = \begin{cases} \Phi_{n-1}(\frac{T}{2}) & \text{if } T \textit{ is even} \\ \Phi_{n-1}(\frac{T-1}{2}) + \Phi_{n-1}(\frac{T+1}{2}) & \text{if } T \textit{ is odd.} \end{cases} \quad (1)$$

The number of alternative paths in GIN_4 is listed in Table I. It can be seen from the table that there is only a unique path for tag = 0, and also the number of alternative paths is small for certain tag values. When T equals $N/2$, for example, there are only two alternative paths, and when T is $N/4$ or $3N/4$, there are only three alternative paths. (Notice that this fact holds for any sized GIN.) Moreover, those alternative paths for such a T share the same common route before starting to fork at a later stage.

The Monogamma interconnection network (MGIN) is constructed from a GIN with the last stage removed and one stage added at the input side; the connecting pattern between the added stage and stage 0 is identical to that between stages 0 and 1 [4]. The numbers of alternative paths in an MGIN with size 16 are also listed in Table I. While there are multiple paths between any source-destination pair in the MGIN, those existing paths are not necessarily disjoint. For tag = 7, as an instance, two alternative paths exist and they both share the same common route after the first three stages, unable to guarantee reliable communication between every source-destination pair with the tag.

All the paths from a source to a destination in the GIN form a ladder structure, based on which the *terminal reliability* of any (S, D) pair (which is defined as the probability of having at least one path for communication between the pair) can be evaluated [1]. Ladder structures representing all the paths between the (1, 0) pair and between the (1, 4) pair in Fig. 1 are shown in Fig. 3(a). Switches are associated with reliabilities, which denote the probabilities that they are still operational (after a given working period). The failure at a connection is considered to be a switch fault. Faults are assumed to arise randomly and independently.

The terminal reliability results of all tag values in GIN_6 are plotted in Fig. 4, where the reliabilities of a 3×3 switch, a 1×3 switch, and a 3×1 switch are assumed to be 0.9, 1.0, and 1.0, respectively (i.e., switches in the first stage and the last stage are fault-free). As expected, terminal reliability is relatively low if the tag value equals 0, $N/2$, $N/4$, $3N/4$, etc. The GIN involves several vulnerable communication pairs. It is thus desirable to ensure the existence of *disjoint* paths between each (S, D) pair, not just multiple paths.

II. RELIABLE GAMMA INTERCONNECTION NETWORKS

A source can reach any destination in the GIN without using any upper (or any lower) connections. This is reflected in the tag generation process studied by McMillen and Siegel [5], where a tag is called the positive (or negative) dominant tag if its tag bits correspond to only the straight and lower (or upper) connections. Using the third type of connections simultaneously makes it possible to create multiple paths in the GIN. Due to the symmetric connecting style characterized by functions $f_i^p(j)$ and $f_i^m(j)$, however, the GIN cannot produce multiple paths if the tag value is 0 and also fails to offer disjoint paths between many source-destination pairs. It appears advantageous by changing the connecting patterns between certain stages such that the upper and the lower connections are no longer symmetric throughout the entire network, if the resultant network guarantees the existence of disjoint paths between every (S, D) pair.

For GIN_n, if the function specified the upper connection from a switch in stage $n - 1$ to a switch in the last stage is changed to $\hat{f}_{n-1}^m(j) = (j - 1) mod\ N$ (the same as $f_0^m(j)$), as shown in Fig. 2(a), with the remaining functions left unchanged, two disjoint paths exist from any S to a $D = S$. All the paths between source 9 and destination 9 in Fig. 2(a) are highlighted. In fact, it will be proved later that disjoint paths are always present from any source to any destination in such a network. An additional possible benefit associated with the network given in Fig. 2(a) is that the wrap-around connections between the last two stages are reduced from 16 (in the GIN illustrated in Fig. 1) to 9. This reduction could translate to a cost savings if several switches are fabricated in one chip using VLSI technology, because fewer pins and less layout area for connections are then needed. For example, if four rows of switches in Fig. 2(a) (inside a dotted rectangle) are housed in a single chip, the pin count for total connections per chip is reduced by 6, and less area is taken by connections between the last two stages.

The GIN can generally be modified in a way that one or several consecutive stages immediately before the last stage have the connecting patterns altered. A reliable GIN of interest is a "modified" GIN such that $\hat{f}_i^m(j)$ changes from $j - 2^i (mod\ N)$ to $j - 1 (mod\ N)$, for all $\alpha \leq i < n$, and the rest of the connection functions remain identical to those in the GIN, denoted by $REGIN_n^\alpha$. $REGIN_4^3$ and $REGIN_4^2$ are depicted in Fig. 2. A REGIN has multiple paths between any source-destination pair, even for S=D, as indicated by bold lines.

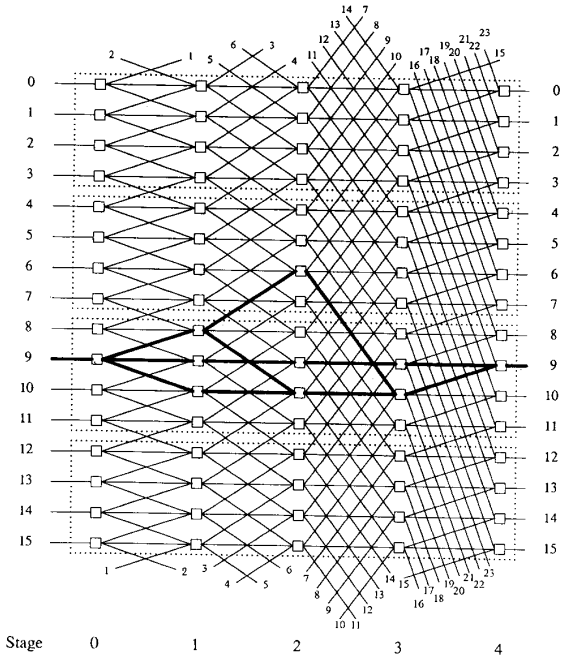

(a) REGIN$_4^3$

(a)

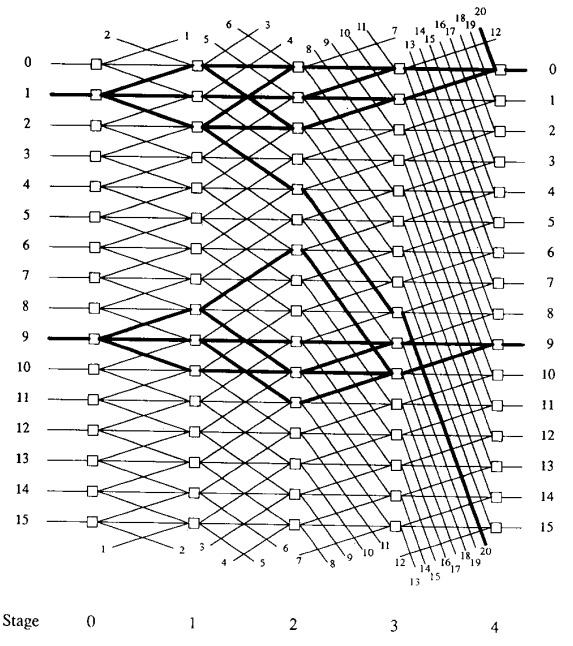

(b) REGIN$_4^2$

(b)

Fig. 2. Two possible reliable GIN's under consideration. (Alternative paths between the (9, 9) pair are highlighted in both networks, and the paths between the (1, 0) pair in REGIN$_4^2$ are also shown.)

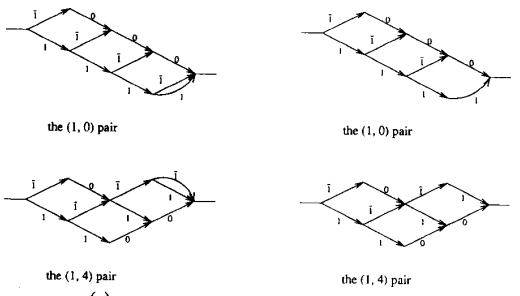

Fig. 3. (a) Ladder diagrams representing all paths between two given communication pairs in GIN$_4$. (b) Degenerative ladder structures representing all paths present between two pairs in DGIN$_4^3$.

Each request in a REGIN$_n^\alpha$, like that in GIN$_n$, carries a routing tag of n digits, with each digit equal to $\bar{1}$, 0, or 1. The tag value, again, is the difference between destination and source. The weight of every tag digit is determined by the connecting patterns. Specifically, the i^{th} tag digit, d_i, has the weight of -2^i if d_i is $\bar{1}$ and $i < \alpha$, and of -1 if d_i is $\bar{1}$ and $i \geq \alpha$; it has the weight of 2^i when d_i equals 1 for any i. In other words, the tag is represented by means of a "modified" binary redundant number system. Digit d_i is used at stage i when routing the request through a REGIN: it takes the lower (or upper) connection for $d_i = 1$ (or -1), and takes the straight connection for $d_i = 0$. Routing complexity is exactly the *same* in a REGIN as in its GIN counterpart.

A. Calculating Alternative Paths

For a given tag, different "modified" binary redundant number systems yield different amounts of possible representations. Since each representation corresponds to one possible path, the REGIN's under consideration offer varying numbers of alternative paths for a tag. It is interesting to find out the number of redundant paths in REGIN$_n^\alpha$, $0 < \alpha < n$. A recurrence equation for the number of redundant paths can be obtained by considering i) every path which can reach destination D from a switch, say switch SW_c, in stage α (SW_c is called a *candidate switch*), and then ii) all the paths which originate from source S and terminate at SW_c. Every path between SW_c and D is found from destination D backward, along a connection at a stage in the reverse direction, one stage after another until stage α is reached.

Since switch D in stage n (the last stage) is connected to three switches in stage $n-1$, namely, switches $D+1 \ (mod\ N)$, D, and $D-2^{n-1} \ (mod\ N)$, a path can reach D from any one of these three switches in stage $n-1$, giving rise to three possible paths from stage $n-1$ to D. Similarly, any one of the three switches in stage $n-1$ is in turn connected to three switches in stage $n-2$, resulting in nine possible paths from stage $n-2$ to D. This process repeats, and the number of possible paths to D triples when moving one stage backwards, as a switch j in stage i, $\alpha < i \leq n$, is connected to exactly three switches in the immediate prior stage, stage $i-1$, which are switches $j+1 \ (mod\ N)$, j, and $j-2^{i-1} \ (mod\ N)$. In REGIN$_n^\alpha$, there are $3^{n-\alpha}$ total paths from stage α to destination D.

Consider a given tag value T. During the process of going backwards from D, the tag value T is changed by a certain amount at each move to reflect the difference $(mod\ N)$ between the current position and S (rather than between D and S). The amount changed depends on the connection taken at that move. Specifically, the amount is 1 (or -2^i) if the connection from switch j in stage $i+1$ to switch $j+1$ (or $j-2^i$) in stage i is taken; it is 0 if the connection from switch j in stage $i+1$ to switch j in stage i is used. Every path

from D backward to a switch in stage α, switch SW_c, determines an accumulated amount δ, dictated by the sequence of connections that form the path; this path is referred to as λ. The tag value $T+\delta$ reflects the difference between S and location SW_c in stage α. Since the connecting patterns from stage 0 to stage α are identical to those of the GIN, the number of possible paths from S to SW_c can be estimated using the formula developed for the GIN. Any such possible path may concatenate path λ to form a complete path from S to D, and all the paths so formed are distinct with one another. For a given path λ, there are $\Phi_\alpha(T+\delta)$ total distinct paths between S and D, where Φ_α is defined in Eq. (1) for the GIN and δ is the tag amount changed when path λ is chosen. Let Λ be the set of those $3^{n-\alpha}$ paths from stage α to D in REGIN_n^α, then, the total number of alternative paths for a given tag T in REGIN_n^α is expressed by

$$\Omega_n^\alpha(T) = \sum_{\lambda\in\Lambda} \Phi_\alpha(T+\delta). \tag{2}$$

It is obvious that (2) becomes (1) when α equals n, and (2) gives the same value for $n=1$ and for $n=0$. Every path in a REGIN, like in the GIN, is associated with a tag representation. Since there are totally 3^n different tag representations for any REGIN_n^α, the summation of alternative paths over all T's, $0 \le T < 2^n$ (note that T and $T-2^n$ are equivalent), should amount to 3^n, regardless of the value of α. This can readily be shown by using (2). As a result, the REGIN_n distributes those 3^n tag representations to 2^n possible T's in a way determined by α. Ω_4^α for every tag under various α values is provided in Table I. The number of paths between every (S, D), pair in any REGIN_4^α is greater than or equal to 2; multiple paths always exist from any S to any D in a REGIN.

Evaluating the Pin Count The pin counts required for the GIN and the REGIN are different. Suppose that $r, r>1$, rows of switches in a network are fabricated in one single chip. There are r switches in stage i of the GIN housed in a chip, and those r switches connect to $\eta = 2 \times \min(2^i, r)$ switches of stage $i+1$ which locate outside the chip. The pin count contributed by the connections originating from stage i and terminating at stage $i+1$ is thus equal to η, where min is the minimum function. Likewise, the r switches in stage $i+1$ need to connect η switches of stage i which locate outside the chip, contributing η to the pin count. The number of total pins needed for connections per chip for GIN_n is $2r + 4 \times \sum_{i=0}^{n-1} \min(2^i, r)$, where the first term accounts for the connections for r sources and r destinations.

The pin count for REGIN_n^α depends not only on n but also on α. For a given n, a smaller α leads to a lower pin count. The amount of total pins for connections per chip for REGIN_n^α can be derived in a similar way, yielding the pin count of $2r + 2 \times \sum_{i=0}^{n-1} \min(2^i, r)^+ 2 \times \sum_{i=0}^{\alpha-1} \min(2^i, r) + 2(n-\alpha)$.

A pin count savings always results, and the savings increases as either n or r grows. The savings is non-negligible even for a medium-sized network. When r equals 4, for example, 84 pins per chip are needed for GIN_6, but only 66 pins per chip are required for REGIN_6^3. This pin savings coupled with shrunk layout area for connections in a REGIN could translate to cost reduction, making the REGIN practically attractive.

III. Terminal Reliability Estimation

While all the alternative paths for any T in the GIN may be represented by a simple ladder diagram [1], the alternative paths in a REGIN no longer constitute a simple structure. Our reliability evaluation is based on exploring two sets of paths such that the paths in one set and the paths in the other set have no internal switch in common, where an *internal switch* refers to a switch located in an internal stage i, $0 < i < n$. Consider a *degenerative* REGIN_n^α obtained from REGIN_n^α by removing the upper connections between stage i and stage $i+1$, for all $\alpha \le i < n$, namely, by omitting all connections that are specified by functions $\hat{f}_i^m(j) = j - 1 (mod\ N)$. Let DGIN_n^α represent the degenerative REGIN_n^α so obtained, then, DGIN_n^α has the same connecting patterns as GIN_n until stage α, and is without any upper connections from stage α on. Let Ξ denote the set of all possible paths from S to D in DGIN_n^α. It is obvious that any path in Ξ is a path existing from S to D in REGIN_n^α.

Next, we consider all the paths, if any, in DGIN_n^α that originate from S and terminate at switch $D+1 (mod\ N)$ of stage $n-1$. These paths can reach destination D, if a connection from switch $D+1 (mod\ N)$ of stage $n-1$ to switch D in the last stage is present; let $\tilde{\Psi}$ be the collection of all these paths and ζ be such a connection. Assuming that Ψ is the set of paths: $\{\lambda | \lambda \equiv \tilde{\lambda}$ concatenates $\zeta\&$ switch D in the last stage, for all $\tilde{\lambda} \in \tilde{\Psi}\}$. Since a connection ζ always exists in $\text{REGIN}_n^\alpha, \alpha < n$, if path λ belongs to Ψ, λ is also present between S and D in REGIN_n^α. Note that Ψ may be an empty set for certain (S, D) pairs in a given REGIN_n^α. If Ψ is not empty, then no path in Ξ shares a common internal switch with any path in Ψ, as described generally in the next theorem.

Theorem 1: Consider the paths from S to D in REGIN_n^α. Any path in Ξ and a path in Ψ have no internal switch in common.

Proof: Let switch j be referred to as an even (odd) switch, if j is even (odd). We observe that in a DGIN (just like in the GIN), an even (odd) switch in stage 0 connects an odd (even) switch in stage 1 when an upper or a lower connection is taken, and connects an even (odd) switch in stage 1 when a straight connection is taken. In addition, from stage 1 on, an even (odd) switch can reach only even (odd) switches in later stages, regardless of what connections are used in the DGIN.

First, let us consider the case when both S and D are even (odd). Because the tag $D-S$ is even, a path in Ξ takes a straight connection, arriving at an even (odd) switch in stage 1. Based on our observation above, the path can reach only even (odd) switches in subsequent stages. On the other hand, a path in Ψ takes either an upper or a lower connection and visits an odd (even) switch in stage 1, since the tag now is $D+1-S$, an odd value. According to the above observation, the path can solely reach odd (even) switches in stages i, $1 < i < n$.

Next, we consider the case when S being even (odd) and D being odd (even). Since the tag $D-S$ is odd, a path in Ξ takes either an upper or a lower connection, arriving at an odd (even) switch in stage 1, and reaches only odd (even) switches in subsequent stages. In contrast, a path in Ψ takes a straight connection and visits an even (odd) switch in stage 1, because the tag now is $D+1-S$, an even value. Apparently, the path can solely reach even (odd) switches in stages i, $1 < i < n$.

Therefore, a path in Ξ and a path in Ψ never have an internal switch in common. □

The subsequent theorem characterizes a special type of REGIN's capable of tolerating arbitrary single faults.

Theorem 2: Totally disjoint paths exist from any source to any destination in REGIN_n^{n-1}

Proof: Consider the three classes of tag values: (1) $0 \le T \le N/2-2$, (2) $T = N/2-1$, and (3) $N/2 \le T \le N-1$. For any class (1) tag, it is easy to show that Ψ is non-empty. According to Theorem 1, this theorem follows immediately. For the class (2) tag, Ξ involves disjoint paths from source S to the two switches in stage $n-1$: $S+T (mod\ N)$ and $S+N-1 (mod\ N)$, which then join switch D in the last stage through the straight and the lower connections, respectively. A class (3) tag T, like a class (1) tag, yields a nonempty Ψ, because a path from S to switch $S+T+1 (mod\ N)$ of stage $n-1$ always exists, suggesting the presence of disjoint paths. This completes the proof. □

We derive the terminal reliability of REGIN_n^{n-1} (i.e., $\alpha = n-1$) only, as such a REGIN is fault-tolerant due to the presence of totally disjoint paths. Terminal reliabilities contributed by the paths in Ξ and by the paths in Ψ are calculated separately. We observe that all the paths in Ξ form a degenerative ladder structure derived from the (regular) ladder diagram by removing one link between the last two stages, so do the paths in Ψ (recall that the ladder diagram represents all the paths between an (S, D) pair in the GIN). Degenerative ladder structures representing the paths in Ξ between the (1, 0) pair and between the (1, 4) pair in DGIN_4^3 are depicted in Fig. 3(b); the degenerative ladder structures representing the paths in Ψ between the same two pairs are simpler and are not given. As before, faults are assumed to happen randomly and independently, and switches in the first and last stages are perfectly reliable. The terminal reliability of a degenerative ladder structure which denotes all the paths in Ξ can be obtained in the same way as what was described in [1] for the ladder diagram, because the removed link is redundant in the (regular) ladder structure. Similarly, the degenerative ladder structure representing all the paths in Ψ can be constructed and its terminal reliability is evaluated (note that for tag $T = N/2 - 1$, Ψ is empty, according to the proof of Theorem 2, and thus no such a degenerative ladder structure exists). Let terminal reliabilities due to the paths in Ξ and due to the paths in Ψ be denoted respectively by R_Ξ and R_Ψ. Terminal reliability due to all the paths in $\Xi \cup \Psi$ is equal to $R_\Xi + R_\Psi - R_\Xi \times R_\Psi$, as switches in the first and last stages are assumed fault-free.

The terminal reliability so calculated versus tag for REGIN_6^5 is plotted in Fig. 4, where the reliability of a 3×3 switch is assumed to be 0.9 as before. Compared with the GIN, REGIN_6^5 enjoys a significant terminal reliability enhancement, particularly at those tags where the GIN suffers from relatively low reliability values. Terminal reliability as a function of switch reliability for REGIN_6^5 is illustrated in Fig. 5, where the results of tag equal to $N/2$, $N/4$, $N/8$ are given because those tags yield low terminal reliability values in the GIN. (Notice that tag = 0 and tag = $N/2$ in the GIN lead to the same lowest terminal reliability.) As can be seen, the REGIN always exhibits a significant terminal reliability improvement over its GIN counterpart under any switch reliability.

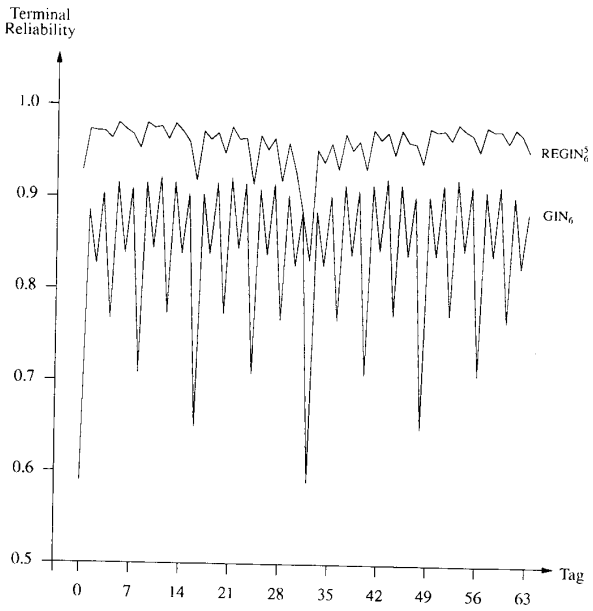

Fig. 4. Terminal reliability comparison between GIN_6 and REGIN_6^5.

IV. Experimental Performance Evaluation

The performance results of various REGIN's under different network parameters are obtained through simulation for comparison. Our study again focuses on REGIN_n^{n-1}'s, which can tolerate any single fault. A network under simulation operates in a packet-switched mode, where every request involves only a single packet and is routed according to the strategy mentioned in Section 2. Each request in a network of size 2^n carries with it a routing tag of n digits, and the tag is selected randomly from all possible tag representations for the tag value of the request. At every network input, a table is used to store all possible tag representations of each tag value, and when a request is generated, the table is looked up for selecting a tag representation; each possible representation of a given tag value has the same probability of being selected.

A simulated network is assumed to be synchronous, with the *network cycle* being the smallest time unit in which a request can be transmitted from one switching element to another. The subsequent assumptions are made in our simulation:

1) Requests arriving at each network input follow a Poisson distribution with an average rate of λ requests per network cycle.
2) The destinations of generated requests are uniformly distributed across all network outputs.
3) A queue with capacity η (which can accommodate η requests) is associated with every switch output port and each network (primary) input.
4) A request heading for a switch output port whose queue is full cannot advance and remains in its original location; it reattempts in the next network cycle.
5) Each switch output port can accept up to three requests during a network cycle. If multiple requests compete for the same switch output port at the same time, a conflict results and the earliest request (in terms of its generation time) is allowed to proceed to the next stage, with the rest contending requests kept in the associated queue.
6) When a request arrives at a network input whose queue is completely occupied, it is not admitted to the network and is dropped. A dropped request is reissued subsequently.
7) A network output can accept one request per network cycle; an accepted request is removed from the network immediately.

Assumption 6) makes our simulation reflect the realistic situation, as a dropped request is often regenerated in practice. The measures of interest in this experimental study are bandwidth and mean delay. Bandwidth is the average number of requests accepted at each network output per cycle. Mean delay indicates the number of network cycles spent by a typical request from its source to its destination; the delay of a request is counted from the time the request is generated (rather than from the time it is admitted to the network) until it leaves the network. The results collected from our simulation are depicted in Figs. 6 and 7, where every point given represents a steady state result. Note that bandwidth, though a collected measure under various request arrival rates (λ), is chosen as one axis for plotting delay-bandwidth characteristics in the figures.

Mean delay versus bandwidth of various REGIN_n^{n-1}'s with fixed queue capacity $\eta = 2$ is shown in Fig. 6 for different network sizes 2^n, $n = 4$, 5, and 6. The simulation results of GIN_n are included in the figure (shown by dotted curves) for comparison. It is observed that for any given n , REGIN_n^{n-1} and its GIN_n counterpart have about the same performance, with the difference falling within tolerable simulation accuracy. This reveals that altering the interconnecting

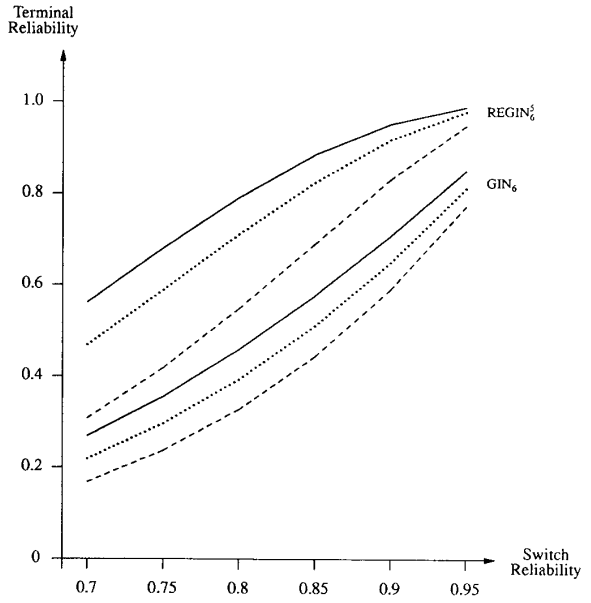

Fig. 5. Terminal reliability vs. switch reliability for GIN_6 and $REGIN_6^5$. (Solid curves show the results for tag = $N/8$, dotted curves show the results for tag = $N/4$, and dashed curves show the results for tag = $N/2$).

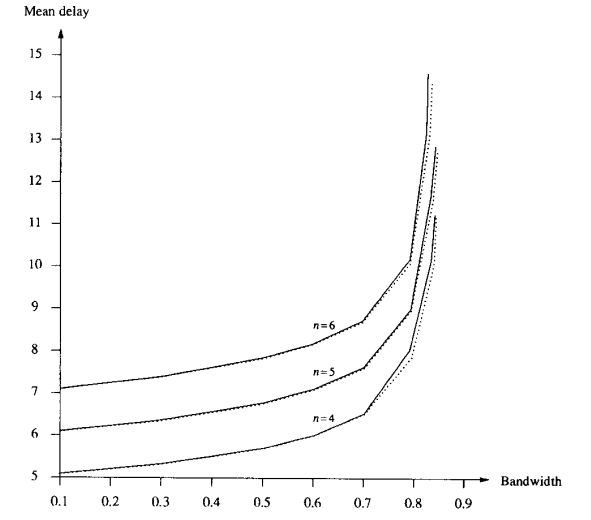

Fig. 6. Mean delay versus bandwidth of various $REGIN_n^{n-1}$, s with queue capacity $\eta = 2$. (The GIN_n results are shown by dotted curves.)

patterns between GIN stages leads to no performance loss, despite a significant gain in terminal reliability and possible cost reduction. It is interesting to notice in the figure that the performance gaps between different sized networks remain unchanged for bandwidth up to roughly 0.7, at which point networks are not heavily loaded yet. The gaps are attributed to the different numbers of stages requests travel in different sized networks. When the network load becomes heavy, the queueing delay in each stage grows substantially and the performance gaps between different sized networks change.

The performance of $REGIN_6^5$ with various queue capacities is illustrated in Fig. 7. It can be seen from the figure that a REGIN

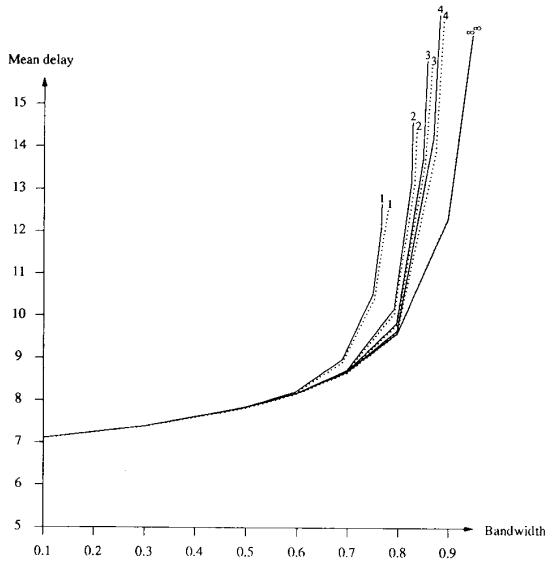

Fig. 7. Mean delay versus bandwidth of $REGIN_6^5$ for various queue capacities. (The GIN_6 results are shown by dotted curves. The number at the end of each curve indicates queue capacity η.)

with queue capacity η equal to 3 can deliver performance close to that of a REGIN with infinite queue capacity. Compared with GIN_6 (shown by dotted curves), $REGIN_6^5$ has almost identical performance for any given queue capacity.

V. Conclusion

In this correspondence, we have investigated "modified" Gamma interconnection networks obtained simply by altering connecting patterns between certain consecutive stages of the Gamma network. A "modified" network exhibits high terminal reliability between every source-destination pair and is dubbed as a REGIN. One type of REGIN's can guarantee the existence of totally disjoint paths between any communication pair. Routing complexity is the same for a REGIN as for its GIN counterpart. All the paths from any source to any destination in a REGIN can be represented by a "modified" binary redundant number system. Results on the distribution of paths for different tags in REGIN's are derived. If several rows of switching elements are fabricated in one chip using VLSI technology, a REGIN could lead to reduced cost, because the pin count per chip decreases and the layout area taken by connections shrinks. It is found through simulation that $REGIN_n^{n-1}$ delivers virtually identical performance as its GIN counterpart.

References

[1] D. S. Parker and C. S. Raghavendra, "The Gamma Network," *IEEE Trans. Computers*, vol. C-33, pp. 367–373, Apr. 1984.

[2] H. J. Siegel and S. D. Smith, "Study of Multistage SIMD Interconnection Networks," *Proc. 5th Annu. Int'l Symp. Computer Architecture,* Apr. 1978, pp. 223–229.

[3] R. J. McMillen and H. J. Siegel, "Performance and Fault-Tolerance in the Inverse Augmented Data Manipulator Network," *Proc. 9th Annu. Int'l Symp. Computer Architecture,* Apr. 1982, pp. 63–72.

[4] C. S. Raghavendra and D. S. Parker, "Reliability Analysis of an Interconnection Network," *Proc. 4th Int'l Conf. Distributed Computing Systems,* May 1984, pp. 461–471.

[5] R. J. McMillen and H. J. Siegel, "Routing Schemes for the Augmented Data Manipulator Network in an MIMD System," *IEEE Trans. Computers,* vol. C-31, pp. 1202–1214, Dec. 1982.