PERFORMANCE EVALUATION OF THE BACKTRACK-TO-THE-ORIGIN-AND-RETRY ROUTING FOR HYPERCYCLE-BASED INTERCONNECTION NETWORKS

Nikitas J. Dimopoulos, Don Radvan, Kin F. Li Electrical and Computer Engineering Department University of Victoria Victoria, B.C. Canada

ABSTRACT[†]

In this work, we present the Hypercycles, a class of multidimensional graphs, which are generalizations of the n-cube. These graphs are obtained by allowing each dimension to incorporate more than two elements and a cyclic interconnection strategy. Hypercycles, offer simple routing, and the ability, given a fixed degree, to chose among a number of alternative size graphs. These graphs can be used in the design of interconnection networks for distributed systems tailored specifically to the topology of a particular application. We are also presenting a back-track-to-the-origin-and-retry routing, whereupon paths that block at intermediate nodes are abandoned, and a new attempt is made. Intermediate nodes are chosen at random at each point from among the ones that form the shortest paths from a source to a destination. Simulation results that establish the performance of a variety of configurations are presented. In addition our initial attempt of constructing a Hypercycle based router is discussed.

1.0 Introduction

Message passing concurrent computers such as the Hypercube[11, 17], Cosmic Cube[15], MAX[12, 13], consist of several processing nodes that interact via messages exchanged over communication channels linking these nodes into one functional entity.

There are many ways of interconnecting the computational nodes, the Hypercube, Cosmic Cube, and the Connection Machine[18] having adopted a regular interconnection pattern corresponding to a binary n-dimensional cube, while MAX adopts a less structured, yet unspecified topology.

Several recent studies attempt extensions and generalizations of the basic tenets of the n-cube. Broder et. al. [4] have proposed product graphs[14] of small "basic" graphs. Their prime concern is to synthesize fault tolerant networks with a given degree of coverage. In these multidimensional graphs, they define a single route from a source to a destination, as the product of routes in each of the constituent dimensions. Routing is exhausted in each dimension before another dimension is considered. Bhunyan and Agrawal [3] have introduced the generalized hypercubes (GHC) which are also graph products of fully connected "basic" graphs. The mixed radix system [2] is used to express the properties of these graphs and their routing. Wittie [19] gives a good overview and comparison of several interconnection networks including the spanning bus and dual bus hypercubes. These are essentially binary n-cubes with broadcast busses connecting the processors in each dimension.

The advantages of having a regularly structured interconnection are many-fold, and they have been proven time and again in their being incorporated in many recent designs [6,11,12,13,15,17,18]. In these structures, easy deadlock-free routing [7] can be accomplished by locally computing each successive intermediate node -for a path that originates at a source node and terminates at a destination node- as a function of the current position and the desired destination. Many regular problems (such as the ones found in image processing,

† This work has been supported by the Natural Sciences and Engineering Research Council Canada, under grant #OGP0001337 physics etc.) have been mapped on such regular structures, and run on the corresponding machines exhibiting significant speedups. In contrast, embedded real-time applications [12, 13], tend to exhibit variable structures that do not necessarily map optimally to an n-cube. In addition, since the size of a binary n-cube is given as 2^n , it means that a particular configuration cannot be expanded but in predefined quantum steps. For example, if a given embedded application requires a system comprised of 9 nodes, the next larger n-cube with 16 nodes must be chosen. This constitutes a significant increase in resource allocation, especially in power-mass limited environments.

Hypercycles [13] can be considered as products of "basic" graphs that allow, as compared to the Generalized Hypercubes (GHC) [3], a richer set of component "basic" graphs ranging in complexity from simple rings to the fully connected ones used in the GHC. Also, contrary to Broder et al [4], we define the component graphs and provide analytical expressions for routing, our aim being twofold:

- (a) To provide computer interconnection networks that match the node requirements of a given embedded system.
- (b) To increase throughput of a given network by providing routing expressions that can be computed analytically (and hence are candidates for VLSI implementation) and which provide a maximum number of alternate paths from a source to a destination. The existence of alternate paths guarantees that a message will not be blocked waiting for its single route to be freed, but it would in turn search for the availability of alternate paths. This strategy also provides for fault protection, since a faulty path can be marked permanently busy, and thus messages can be routed around it.

The Hypercycles, being regular graphs, retain the advantages of easy routing and regularity. Yet, since we are dealing with a class, rather than isolated graphs, we have the flexibility of adopting any particular graph (from the class) that closely matches the requirements of a given application.

This work is divided into these parts. Section 2.0 introduces the Mixed Radix System, Section 3.0 presents some basic graph terminology and notation, while Section 4.0 introduces the Hypercycles and discusses their properties. Sections 5.0 and 6.0 discuss routing and evaluate the performance of several example Hypercycles.

2.0 Mixed Radix Number System

The mixed radix representation [2], is a positional number representation, and it is a generalization of the the standard b-base representation, in that it allows each position to follow its own base independently of the other.

Given a decimal number M factored into r factors as $M = m_1 \times m_2 \times m_3 \times \cdots \times m_r$ then any number $0 \le X \le M-1$ can be represented as

$$(X_{m_1,m_2,...,m_r} = x_1 x_2 ... x_r |_{m_1,m_2,...,m_r}$$
 where $0 \le x_i \le (m_i - 1)$;

$$i = 1, 2, ..., r$$
 and the x_i 's are chosen so that $X = \sum_{i=1}^{n} x_i w$

i

where
$$w_i = \frac{M}{m_1 m_2 \cdots m_i}$$

3.0 Graph Notation

An undirected graph G is defined as: $G=(N,\ E$), where N is the set of nodes

 $N = \{ \alpha_i ; i = 1, 2, ..., N \}, \text{ and } \mathcal{E} \text{ the set of edges}$ $\mathcal{E} = \left\{ e_{ij_i} = \left(\alpha_i, \beta_{j_i} \right) j_i = 1, 2, ..., d_i ; i = 1, 2, ..., N \right\}$ with $\alpha_i, \beta_{j_i} \in \mathbb{N}$ and d_i the degree of node α_i . The degree of a

graph, d(G), is the maximum of the node degrees. A walk in G [5] is a sequence of edges $e_1 e_2 \dots e_l$, such that if $e_i = (\alpha_i, \alpha_{i+1})$ then

 $e_{i+1} = (\alpha_{i+1}, \alpha_{i+2})$ and $e_i \in \mathbb{E}$. The distance, dis (γ, δ) , between nodes γ and δ is defined as the shortest walk between γ and δ if any, otherwise, dis $(\gamma, \delta) = \infty$. The diameter k, is the maximum distance between any pair of nodes. A graph is regular, if all nodes have the same degree.

4.0 Hypercycles.

An *r*-dimensional Hypercycle, is the following regular undirected graph: $\mathbf{G}_{m}^{p} = \left\{ \mathbf{N}_{m}^{p}, \mathbf{E}_{m}^{p} \right\}$ where

 $m = m_1, m_2, m_3, ..., m_r$ a mixed radix, $\rho = \rho_1, \rho_2, ..., \rho_r$; $\rho_i \le m_i / 2$ the connectivity vector, determining the connectivity in each dimension which ranges from a cycle $(\rho_i = 1)$ to fully connected $(\rho_i = \lfloor m_i / 2 \rfloor)$, and $\mathcal{N} \quad \stackrel{\rho}{m} = \{0, 1, 2, ..., M-1\}$. Given $\alpha, \beta \in \mathcal{N} \quad \stackrel{\rho}{m}$ then $(\alpha, \beta) \in \mathcal{E} \quad \stackrel{\rho}{m}$ if and only if there exists

 $\alpha, \beta \in \mathbb{N}$ if and only if there exists

 $l \le j \le r$ such that $\beta_j = (\alpha_j \pm \xi_j) \mod m_j$ with $l \le \xi_j \le \rho_j$ and $\alpha_i = \beta_i$; $i \ne j$

Hypercycles, have degrees [8] $d = \sum_{i=1}^{n} f(m_i, \rho_i)$ where $f(m_i, \rho_i) = \begin{cases} 2\rho_i & \text{if } 2\rho_i < m_i \\ m_i - 1 & \text{if } 2\rho_i = m_i \end{cases}$

and diameter k 1

$$k = \sum_{i=1}^{L} \left\lceil \frac{\lfloor m_i/2 \rfloor}{\rho_i} \right\rceil$$

The n-cube is a Hypercycle, with $M = 2 \times 2 \times \dots \times 2 = 2^n$ and $\rho = 1, 1, 1, \dots, 1$.

4.1 Routing

Hypercycles, have routing properties that are similar to those of the n-cubc. Given nodes $(\alpha)_{m_1m_2...m_i...m_r} = \alpha_1 \alpha_2....\alpha_r$ and $(\alpha^*)_{m_1m_2...m_i...m_r}$ $= \alpha_1 \alpha_2...\xi_...\alpha_r$, a walk, from node α to node α^* , can be constructed as follows: $\alpha_1 \alpha_2...\alpha_i...\alpha_r$, $\alpha_1 \alpha_2...\xi_2...\alpha_r$, ..., $\alpha_1 \alpha_2...\xi_...\alpha_r$, such that§

$$\xi_{j_{i}+1} = \begin{cases} \left(\xi_{j_{i}} + \rho_{i}\right) \mod m_{i} & \text{if} \left[\left(\xi - \xi_{j_{i}}\right) \mod m_{i} = \left|\xi_{j_{i}}, \xi\right|\right] > \rho_{i} \\ \left(\xi_{j_{i}} - \rho_{i}\right) \mod m_{i} & \text{if} \left[\left(\xi_{j_{i}} - \xi\right) \mod m_{i} = \left|\xi_{j_{i}}, \xi\right|\right] > \rho_{i} \\ \xi & \text{if} \left|\xi_{j_{i}}, \xi\right| \le \rho_{i} \end{cases}$$

$$\xi_{0} = \alpha_{t} \qquad \xi_{l_{max}} = \xi$$
Eqn. 4.1.1

We call the length l_{max} of such a walk, the distance along dimension *i*

Given an origin $(\alpha)_{m_1,m_2,\ldots,m_r} = \alpha_1 \alpha_2 \ldots \alpha_r$ and a destination $(\beta)_{m_1,m_2,\ldots,m_r} = \beta_1 \beta_2 \ldots \beta_r$ and if q_i denotes their distance along dimension *i*, the total distance between the origin and the destination nodes, denoted as dis (α,β) and defined as the sum of the individual distances along all the dimensions, is given as

dis
$$(\alpha, \beta) = q = \sum_{i=1}^{n} q_i$$
. For these nodes, there are a total of [‡]

$$I = \begin{pmatrix} q \\ q_1, q_2, \dots, q_r \end{pmatrix} = \frac{q!}{q_1! q_2! \cdots q_r!}$$

distinct walks of length q that connect them. These are constructed by sequentially modifying the source address, each time substituting a source digit by an intermediate walk digit, until the destination is reached. The following walk connects source to destination.

source = $\alpha_1 \alpha_2 \alpha_3 \dots \alpha_r$; $\alpha_1 \xi_1 \alpha_3 \dots \alpha_r$; $\alpha_1 \xi_1 \psi_1 \dots \alpha_r$; $\alpha_1 \xi_2 \psi_1 \dots \alpha_r$; $\alpha_1 \xi_2 \psi_2 \dots \alpha_r$; ...; $\alpha_1 \xi_2 \beta_3 \dots \alpha_r$;...; $\beta_1 \beta_2 \beta_3 \dots \beta_r$ = destination

Figure 1a., gives an example of two distinct walks of equal length that connect a source to a destination, for a Hypercycle.

4.2 Average Distance Calculation Given an r-dimensional hypercycle G_m^{ρ} , and denoting by n_l the number of nodes at distance *l* from a source node $(\alpha)_{m_1 m_2 \dots m_r} = \alpha_1 \alpha_2 \dots \alpha_r$, the average distance between any two nodes in G_m^{ρ} can be calculated as[8]

$$\overline{q} = \frac{\sum_{l=1}^{l} ln_l}{m_l m_2 \cdots m_r}$$

Some typical distances are given in TABLE 1.

5.0 Deadlock Avoidance in Routing.

In section 4.1, we have given a method that establishes at least one path from a source to a destination node. In this part, we are concerned with optimally choosing one of the paths. Routing must be efficient and deadlock free. Deadlock occurs when resources (in this case node to node communication segments) are allocated so that the completion of a partial path requires a segment already allocated to a different partial path which in turn waits for a segment in the first partial path. It is obvious that no messages can propagate over the deadlocked paths, and the only remedy is to break the already established and deadlocked partial paths and try again.

Deadlock may occur easily in cases where the segments that form the paths are chosen at random. Certain routing algorithms (e.g. virtual channels, e-cube routing[7]) prevent deadlocks by ordering the resources (channels) to be allocated. Thus a lower order resource cannot be committed if a needed higher order resource cannot be obtained. The disadvantage of this approach in an interconnection network is that it limits the number of paths connecting a source to a destination to exactly one, even though several alternate free paths

¹ The function $\lfloor x \rfloor$ denotes the

largest integer smaller than or

equal to x, while [y] denotes the

smallest integer larger than or

equal to y.

[•] We define $|a, b| = \min\{(a - b) \mod m_i, (b - a) \mod m_i\}$

[‡]For the definition of a multinomial number, see [1] pp 32.

may exist at a particular moment. We are proposing to adopt a strategy where deadlocks are avoided by requiring a blocked partial path to backtrack to its origin and retry.

6.0 Backtrack-to-the-origin-and-retry routing

For Hypercycle-based interconnection networks, because of the existence of cycles in each dimension, the use of an e-cube type routing that prevents deadlocks, is impractical. We are proposing instead a deadlock avoiding routing strategy. According to our backtrack-to-the-origin-and-retry routing we identify, at each node, all nodes that can be used for the continuation of the path. For all such identified nodes, we also identify the corresponding ports that can be used in order to continue the path. Since several paths may be forming in parallel, some of these ports may already be allocated to some other path. After excluding all the allocated ports, we select one of the remaining free ports at random. The subsequent link in the path is established is then established through the selected port, and the procedure repeats itself until the destination is reached, or no free ports could be found. If no free ports are to be found at an intermediate node in the path, then a break is returned to the origin (through the already established partial path to the blocking node), the partial path is dissolved, and a new attempt for the creation of the required path is initiated. This routing strategy avoids deadlocks through backtracking, and also guarantees that the formed path will be of a minimum length, since each subsequent link is selected according to equation 4.1.1. The backtrack-to-the-origin-and-retry routing is a type of two-phase locking [16], where as resources we consider the various links necessary for the completion of the sourceto-destination circuit.

We have used Extend^{TM†} to construct a simulator capable of simulating any Hypercycle based network. For this simulator, we implemented both the backtrack-to-the-origin-and-retry as well as the e-cube routing strategies. The e-cube routing can only be used for binary cube networks. For each node, we assumed a Poisson message generator which generates packets with uniform distribution of destinations. Each packet carries the destination address which is used for routing. Links are assigned priorities, so that collisions can be resolved. We assumed a packet transmission time (over an established source to destination path) of 100 simulation-clock ticks. We use the simulator to obtain the throughput and delay characteristics of several networks for both e-cube and backtrack-tothe-origin-and-retry in terms to the offered load. Both the offered load and the throughput were normalized in terms to the maximum capacity of each network taken to be proportional to the number of links in the corresponding graph. The average delay was expressed in actual time units necessary to establish a source-to-destination circuit. Simulation results are depicted in figs. 2, 3, 4, 5 and 6.

As it was expected, the performance of the backtrack-to-theorigin-and-retry for both binary cubes and hypercycles of similar sizes, is clearly superior to that of the e-cube as it can be seen in figs. 2 and 3. This is attributed to the fact that the backtrack-to-the-originand-retry can use alternative paths to the destination instead of the single path allotted by the e-cube routing. The additional advantage of the backtrack-to-the-origin and-retry is its inherent fault tolerance. Indeed, if one of links in the network failed, it could be marked as permanently busy, and packets would be routed around it. This obviously is not the case for the e-cube routing.

Figure 2. further shows that under heavy loads, e-cube routing has slightly higher throughput rates for a given load than backtrackto-the-origin. Figure 4 shows the reason for this anomaly. Backtrack routing does not effectively route packets of longer distances because the path is dissolved as soon as a blocked route is encountered, which has a high probability of occurrence under a heavy load. An adaptive algorithm which alternates between backtrack-to-the-origin and e-cube routing under heavy loads may solve this instability.

For graphs of higher degrees, figs. 5 and 6 show the effect of alternate paths on system delay and throughput. The 7 node Hypercycle has only one path choice per source/destination connection (because it is a fully connected graph). This effectively

reduces the system to e-cube style routing only. The 15-node (G $\frac{21}{53}$)

and destination and therefore provide higher system throughput. Generally, system throughput and delay are functions of both average distance and the average number of alternate paths between any two nodes.

The backtrack-to-the-origin-and-retry routing, as discussed above, is currently being implemented in hardware. Figure 7 gives the structure of a computational node in a concurrent computer that incorporates Hypercycles and backtrack-to-the-origin-and-retry routing. Figure 8 gives a block diagram of an r-dimensional Hypercycle router. Figure 9 shows the details of the next port generator.

As it can be seen in Figure 8, we are implementing our routing as a system having four modules. The destination address is used in the Next-Port Generator to generate all possible ports that can be used in forming the path to the required destination address. Subsequently, the Port Validator masks out the ports which are currently used by other paths. Finally, The Port Selector, selects at random one of the validated ports which is then used to continue the circuit towards the required destination. For the random number generator, we use a cellular automaton [1] to generate a 27 bit random number, which we use to obtain its modk where k is the number of valid ports incoming to the Port Selector. It is worth noting that the system is programmable, in the sense that it needs the parameters m, ρ as defined earlier and which define the structure of the network, as well as ξ which the address of the current node.

We have currently finished the design of the Next Port Generator in 74LS logic. The block diagram of the Next Port Generator is given in Figure 9. We have obtained a propagation delay of less than 270 ns by using the 74LS technology, and we are confident that it will be drastically reduced (to less than 100 ns) for a custom VLSI implementation. We are using this implementation technology in order to validate our design. We are currently completing the implementation of the Random Number generator using the same technology. Our next objective is to port our design in a gate array so that speed and compactness could be accomplished.

7.0 Conclusions and Discussion

In this work, we presented the Hypercycle, a class of multidimensional graphs, which are essentially generalizations of the n-cube.

Although these graphs are not the densest possible, they are attractive, because of their simple routing. Similarly to the n-cube, the destination address is used to sequentially route a message through intermediate nodes as outlined in section 4.1. Also, since the node addresses are represented in a mixed radix as a sequence of r-digits, each one of these digits is processed independently and in parallel with the remaining digits. Thus the hardware involved in the routing can be made fast (because of the parallelism) and simple (since each module need only handle arithmetic $modm_i$, as compared to arithmetic $modm_1m_2...m_r$ needed when all the address digits are necessary as is the case with such networks as the chordal rings [10], or the cube connected cycles [51].

The graphs presented in this study, are generalizations of some well known graphs such as the binary n-cube, 2- and 3-dimensional meshes, and rings, which are included as special cases. Examples of some special cases are depicted in Figure 1.

REFERENCES

- 1. Berge C., *Principles of Combinatorics* Academic Press, New York and London, 1971.
- Bhuyan, L. N., and D. P. Agrawal, "Design and Performance of Generalized Interconnection Networks" *IEEE Trans. Comput.* Vol. C-32, No. 12, pp. 1081-1090, Dec. 1983.
- Bhuyan, L. N., and D. P. Agrawal, "Generalized Hypercube and Hyperbus Structures for a Computer Network" *IEEE Trans. Computers.* Vol. C-33, No. 4. pp. 323-333, (Apr. 1984).

and 16-node (a_{44}^{22}) graphs offer two alternate routes between source

[†] Extend is a trademark of Imagine That inc.

- A. Broder, D. Dolev, M. Fischer, B. Simons "Efficient Fault 4. Tolerant Routings in Networks" Proceedings of the 16th Annual ACM Symposium on the Theory of Computing pp. 536-541, (May 1984).
- 5. Carlsson, G. E., J. E. Cruthirds, H. B. Sexton, and C. G. Wright "Interconnection Networks Based on a Generalization of Cube-Connected Cycles" IEEE Trans. Comput., Vol. C-34, No. 8, pp. 769-772, Aug. 1985. E. Chow, H. Madan, J. Peterson "A Real-Time Adaptive
- 6. Message Routing Network for the Hypercube Computer" Proceedings of the Real-Time Systems Symposium, pp. 88-96, San Jose CA., (Dec. 1987) Dally, W.J., and C. L. Seitz "Deadlock-Free Message Routing
- 7. in Multiprocessor Interconnection Networks" IEEE Trans. Comput. Vol. C-36, No. 5, pp. 547-553, May 1987.
- Dimopoulos, N., R.D. Rasmussen, G.S. Bolotin, B.F. Lewis, R. M. Manning "Hypercycles, Interconnection 8. Networks with simple Routing Strategies" Proceedings of the Canadian Conference on Electrical and Computer Engineering, pp. 577-580, Vancouver, B.C., Canada, Nov. 1988. Hortensius, P.D., R. D. McLeod, H.C. Card " Parallel
- 9. Random Number Generation for VLSI Systems using Cellular Automata" IEEE Trans. Computers Vol. 38, pp. 1466-1473, October 1989.
- 10. Imase, M., T. Soncoka, and K. Okada, "Connectivity of Regular Directed Graphs with Small Diameters" IEEE Trans. Comput., Vol. C-34, No. 3, pp. 267-273, Mar. 1985.

- Peterson, J.C., J. O. Tuazon, D. Lieberman, M. Pniel "The 11. MARK III Hypercube -Ensemble Concurrent Computer" Proceedings of the 1985 International Conference on Parallel Processing pp. 71-73, Aug. 20-23 1985.
- Rasmussen, R. D., G. S. Bolotin, N. J. Dimopoulos, B. F. 12. Lewis, and R. M. Manning "Advanced General Purpose Multicomputer for Space Applications" Proceedings of the 1987 International Conference on Parallel Processing. pp. 54-57. (Aug. 1987)
- Rasmussen, R. D., N. J. Dimopoulos, G. S. Bolotin, B. F. Lewis, and R. M. Manning "MAX: Advanced General Purpose Real-Time Multicomputer for Space Applications" 13 Proceedings of the IEEE Real Time Systems Symposium pp. 70-78, San Jose, CA., (Dec. 1987).
- G. Sabidussi "Graph Multiplication" Math. Zeitschr. Vol. 72, pp. 446-457 (1960). 14.
- Scitz, C. L. "The cosmic cube" CACM vol. 28, pp.22-33, 15. Jan. 1985
- Tanenbaum, A. .S., Operating Systems: Design and 16. Implementation Prentice Hall, 1987.
- Tuazon, J. O., J. C. Peterson, M. Pniel, and D. Lieberman 17. "Caltech/JPL MARK II Hypercube Concurrent Processor" Proceedings of the 1985 International Conference on Parallel Processing pp. 666-673, Aug. 20-23 1985.
- Waltz, D. L. "Applications of the Connection Machine" IEEE 18. Computer January 1987, pp.85-97
- L. D. Wittie "Communication Structures for Large Networks 19. of Microcomputers" IEEE Trans. on Computers Vol. C-30, No. 4, pp.264-273. (Apr. 1981).

<u>IABLE I.</u>						
	GRAPH	<u> </u>	DEGREE	DIAMETER	NODES	AVERAGE DISTANCE PER NODE
m	=	62	6	2	12	1.454
r m	=	53	6	2	15	1.571
р т	=	522	6	3	20	1.894
р т	=	333	6	3	27	2.077
р т	=	3322	6	4	36	2.4
r m	=	32222	6	5	48	2.723
ρ m	=	11111	6	6	64	3.047
ρ	=	$\overline{1}\overline{1}\overline{1}\overline{1}\overline{1}\overline{1}\overline{1}$	6	6	68	3 403
ρ	=	112	د ا	6	00	3 434
m P	=		0	Q	04	3.434
m P	=	357	0	0	105	3.015
m P	=	555 111	6	6	125	3.629
m	=	222222	27	7	128	3.528





b. Binary 3-cube G_{222}^{III}



Figure 1. Examples of Hypercycles.



Figure 3. Delay vs. offered load for the 4-cube using backtrack-tothe-origin and e-cube routing. The offered load is normalized to the capacity of the interconnection network. The delay is normalized to the data transmission time.

1







Figure 4. Average distance (of successful packets) vs. offered load for the 4-cube using backtrack-to-the-origin and e-cube routing.



Figure 5. Throughput vs. offered load for several graphs of degree six. The offered load and throughput are normalized to the capacity of the interconnection network.



Figure 6. Delay vs. offered load for several graphs of degree six. The offered load and delay are normalized to the capacity of the interconnection network.







Figure 8. Block diagram of the Hypercycle routing engine. Routing in each dimension is done in parallel.



Figure 9. Details of the Next-Port Generator in the Hypercycle routing engine. This generator is valid for all ρ such that $2\rho< m.$