Routing in Hypercycles. Deadlock Free and Backtracking Strategies†

N. J. Dimopoulos , M. Chowdhury, R. Sivakumar, V. Dimakopoulos

Electrical and Computer Engineering Department, University of Victoria
Victoria, B.C. Canada

Hypercycles, is a class of multidimensional graphs, which are generalizations of the n-cube. They are obtained by allowing each dimension to incorporate more than two elements and a cyclic interconnection strategy. Hypercycles, can be used in the design of interconnection networks for distributed systems tailored specifically to the topology of a particular application. Routing strategies, including deadlock-free and deadlock-avoiding have been developed and their relative performance established through simulations.

1.0 Introduction

Message passing concurrent computers utilize interconnection networks, such as the binary n-cube, toruses, k-ary n-cubes etc. to interconnect the processing nodes. Hypercycles[1] is a class of interconnection networks, which can be represented as products of "basic" graphs ranging in complexity from simple rings to fully connected. They are regular graphs, and exhibit routing similar to the n-cube. Because of the range of different topologies included in the class, networks that closely match the requirements of a given application can be found.

2.0 Hypercycles.

An r-dimensional Hypercycle, is denoted as $\mathcal{G}^{\rho}_{m} = \left\{\mathcal{N}^{\rho}_{m}, \mathcal{E}^{\rho}_{\mathrm{m}}\right\}$. $\mathcal{N}^{\rho}_{m}$ and $\mathcal{E}^{\rho}_{\mathrm{m}}$ the sets of nodes and edges, $m = m_1, m_2, m_3, \ldots, m_r$ a mixed radix, $= r_1, r_2, \ldots, r_r$; $r_i \leq m_i /2$ the connectivity vector, determining the connectivity in each dimension (for a ring $r_i = 1$) for fully connected $r_i = \lfloor m_i /2 \rfloor$). For $\alpha, \beta \in \mathcal{N}^{\rho}_{m}$ then $(\alpha , \beta) \in \mathcal{E}^{\rho}_{\mathrm{m}}$ if and only if there exists $1 \leq j \leq r$ such that $\beta_j = (\alpha_j \pm \xi_j) \,\mathbf{mod}\, m_j$ with $1 \leq \xi_j \leq \rho_j$ and $\alpha_i = \beta_i$; $i \neq j$

3.0 Deadlock-Free routing in Hypercycles..

Deadlocks occur when resources (in the case of circuit switching, node to node communication channels) 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. Routing in hypercubes (e-cube) avoids deadlocks by ordering the channels to be allocated. A lower order channel is not committed if a needed higher order channel cannot be obtained. We have proven [2] that hypercycles expressed as products of graphs which possess a deadlock-free routing, have a deadlock-free routing similar to the e-cube. Such hypercycles include the fully connected ones, and one dimensional hypercycles $\mathcal{G}^{\rho}_{m}$ with $m = 4\rho$

4.0 Random Selection and Backtracking Routing Strategies.

Deadlock-free routing policies, limit the number of paths connecting a source to a destination to exactly one, even though several alternate free paths may exist. In order to utilize the multitude of paths between any source-destination pair, and also not to deadlock the system, a path is allowed to chose at random any of the eligible

† Supported by the Natural Sciences and Engineering Research Council Canada, and IRIS

and available channels in order to advance. If there is a block, then the path backtracks and retries again. In backtracking, there are several options.

Backtrack to the origin. Blocked paths retreat to their origin and try again. Random selections of subsequent links at intermediate nodes diminish the probability that the path which gave rise to the block will be formed again.

Backtracking to a random intermediate node. Blocked paths backtrack to random intermediate nodes, and retry. This preserves some part of the established path and presumably longer paths would have a better chance of succeeding.

Persist at the block before backtracking. Blocked paths persist for a certain period of time at the block. In order not to deadlock, paths are allowed to persist for finite times only. In a variation, longer partial paths persist for longer periods Thus paths which have progressed, are given a better chance of completing.

Mixed Longer paths are allowed to follow a deadlock-free routing for hypercyclcles where this is possible, shorter paths follow one of the backtracking strategies. This strategy guarantees that all the longer paths will eventually reach their destination, while allowing the shorter paths to utilize alternative routes.

Simulation results of the routing policies discussed above for the binary 4-cube. are depicted in Fig. 2. In summary, the backtracking strategies exhibited better performance as compared to the e-cube routing for low traffic. As the network stated saturating, the longer paths were at a disadvantage. *Backtrack to the origin* and *backtrack to an intermediate node*, exhibited much slower saturation as compared to the *e-cube*, which seems to suggest that in a well designed system, they offer a much larger bandwidths compared. It is interesting to note that the mixed and the persist strategies are exhibiting gradual e-cube behavior in the sense that they saturate much earlier but while saturated, they tend not to be biased towards shorter paths.

REFERENCES

1. Dimopoulos, N.,et al "Performance Evaluation of the Backtrack to the Origin and Retry Routing for Hypercycle based Interconnection Networks" *Proceedings, 10th International Conference on Distributed Systems*, Paris, pp. 278-284 (June 1990).
2. Dimopoulos, N.,et al" Routing and Broadcasting in Hypercycles. Deadlock Free and Backtracking Strategies'" *under review, IEEE Trans. Paral. and Distr. Systems.*

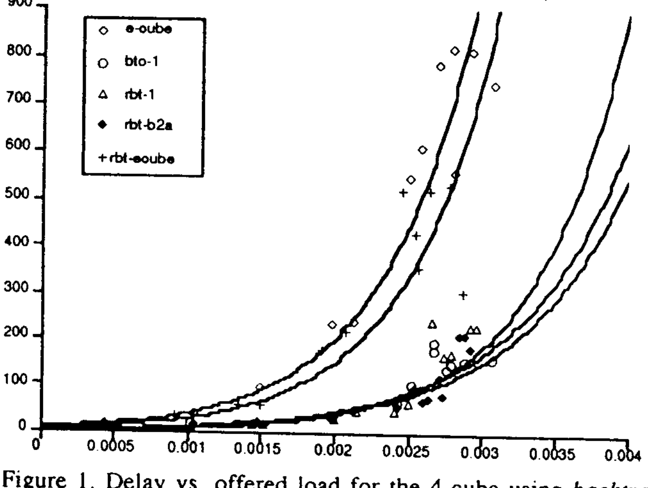

Figure 1. Delay vs. offered load for the 4-cube using *backtrack-to-the-origin* (bto-1) , *random backtracking* (rbt-1), *persistent backtracking* (rbt-b2a), *mixed routing* (rbt-ecube) and *e-cube* (e-cube) routing. The offered load is normalized to the capacity of the interconnectionnetwork. The delay is normalized to the data transmission time.