A New Approach For Optimal MIMD Queueless Routing Of Omega and Inverse-Omega Permutations On Hypercubes

Jean-Pierre Jung[1], Ibrahima Sakho[2]
UFR MIM, Université de Metz
Ile du Saulcy BP 80794 - 57012 Metz Cedex 01 – France
[1]jpjung@univ-metz.fr [2]sakho@univ-metz.fr

Received: January 18, 2011 Accepted: April 6, 2011 DOI: 10.5296/ npa.v3i1.563

Abstract

Omega permutations constitute the subclass of particular permutations which have gained the more attention in the search of optimal routing of permutations in hypercubes. The reason of this attention comes from the fact that they are permutations for general-purpose computing like the simultaneous conflict-free access to the rows or the columns of a matrix. In this paper we address the problem of the optimal routing of omega and inverse-omega permutations on hypercubes under the MIMD queueless communication model. We revisit the problem through a new paradigm: the so-called graphs partitioning in order to take advantage of the recursive structure of the hypercubes topology. We prove that omega and inverse-omega permutations are partitionable. That is any omega (resp. inverse-omega) permutation on n-dimensional hypercube can be decomposed in two independent permutations on two disjoint (n-1)-dimensional hypercubes. We also prove that each one of these permutations is also an omega (resp. inverse-omega) permutation. It follows that any omega (resp. inverse-omega) permutation on n-dimensional hypercube is routable in at most n steps of data exchanges, each step realizing the partition of the hypercube.

Keywords: interconnection network, omega network, hypercube, permutations, perfect shuffle, maximum matching of bipartite graph, graph partitioning, parallel processing, MIMD queueless routing.

1. Introduction

The processors interconnection network (IN) is the heart of the no remote memory access parallel computers in which the inter-processors communications are realized by exchanging messages over the processors interconnection links. The performance of such computers depends greatly on the performance of their processors IN. Among others performance criteria, the scalability of an IN for massive parallelism, its capability of messages deadlock-free routing on shortest paths, its capability of simulating others IN and its management facility are essential. In the research for IN which fulfil these criteria, hypercubes constitute a very attractive alternative. Indeed, the incremental construction of hypercubes confers to them interesting topological properties [1] which allow them to meet most of these essential performance criteria.

Consequently, on one side several commercial parallel machines using them as IN have been built over the years [2]. They have also been used as a means of interconnecting and extending switching matrices in ATM cross connects [3] or proposed as a model for new ATM switches with low complexity and high performance [4]. Nowadays, they continue to be an attractive solution for multicore processor IN [5]. On the other side, several theoretical research works like the one in [6] have been done on different aspects of their use as IN. Among these theoretical researches, one of the most challenging, since at least a quarter of a century, is their rearrangeability under queueless routing constraint, that is their capability to route optimally any permutation such that each node holds only one message throughout the routing.

Omega permutations constitute the subclass of particular permutations which have gained the most attention in the search of optimal routing of permutations in hypercubes. The motivation of this attention comes from the fact that they are permutations for general-purpose computing like the simultaneous conflict-free access to the rows or the columns of a matrix. In this paper we address the problem of the optimal routing of omega permutations. We revisit the problem through a new paradigm the so called partitioning for taking advantage of the recursive structure of the topology of hypercubes.

The remainder of the paper is organised in six sections. Section II gives the problem formulation and some basic definitions related to hypercube, permutations and routing. Section III presents the related works. Section IV introduces the mathematical foundation used to develop the proposed routing algorithm. Section V characterizes partitionable omega permutations then exhibit partitions which assure to resulting permutation to be omega permutation and proposes a routing algorithm. Section VI shows how to deduce a routing algorithm for an inverse omega permutation from its related omega permutation. Section VII concludes the paper and presents some perspectives to improve the length of the routes and so reducing the exchanges steps to the minimum.

2. Problem Formulation

2.1 Definitions

2.1.1 n-Dimensional hypercube

A n-dimensional hypercube, nD-hypercube, is a graph $H^{(n)} = (V, E)$ where :

- V is a set of 2^n nodes $u=0, 1, \ldots, 2^n-1$ denoted by their binary code $(u_{n-1}u_{n-2} \ldots u_0)$ where $u_i \in \{0, 1\}$,
- E is the set of edges {u, v} whose binary codes differ on only one bit.

It is well known that for $n > 0$, a nD-hypercube is obtained by interconnecting two (n-1)D-hypercubes in any one of its dimensions $0 \leq i \leq n-1$. So any hypercube $H^{(n)}$ can be viewed as any of the n couples of $(H^{(n)}_{0,i}, H^{(n)}_{1,i})$ of (n-1)D-hypercubes obtained by restricting the nodes of $H^{(n)}$ to $D_{x,i}$. Fig. 1 illustrates such a view in dimension 3 for a 4D-hypercube.

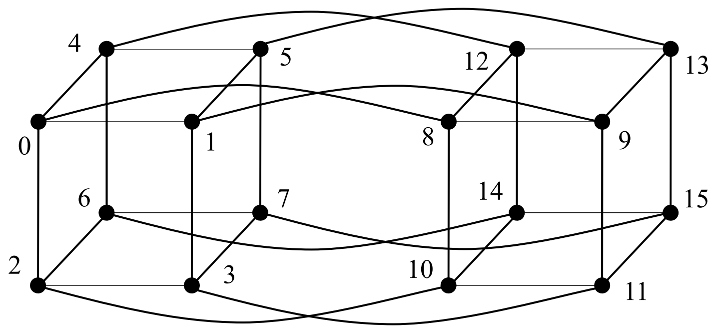

Figure 1. The 4D-hypercube viewed as the interconnection in the dimension 3 of two 3D-hypercubes

2.1.2 Permutation

A permutation on a nD-hypercube $H^{(n)} = (V, E)$ is a bijective map π from V onto itself which associates each node $u=(u_{n-1}u_{n-2} \ldots u_0)_2$ of $H^{(n)}$ with one and only one node $(\pi(u))= (\pi_{n-1}(u)\pi_{n-2}(u)\ldots\pi_0(u))_2$. It is represented by the sequence $\pi = (\pi(u); u=0, 1, \ldots, 2^n-1)$.

2.1.3 Perfect shuffle

The left (resp. right) perfect shuffle is the permutation σ which associates each node $u=(u_{n-1}u_{n-2} \ldots u_0)$ with the node $\sigma(u)=(u_{n-2}u_{n-3}\ldots u_0u_{n-1})$ (resp. $(u_0u_{n-1}u_{n-2}\ldots u_1)$). Table 1 illustrates the left perfect-shuffle on the 3D-hypercube. In the sequel, except if specified, we consider left perfect-shuffles.

Table 1. The perfect schuffle σ=(0, 2, 4, 6, 1, 3, 5, 7) on the 3D-hypercube.

u	u_2	u_1	u_0	$\sigma_2(u)$	$\sigma_1(u)$	$\sigma_0(u)$	$\boldsymbol{\sigma(u)}$
0	0	0	0	0	0	0	**0**
1	0	0	1	0	1	0	**2**
2	0	1	0	1	0	0	**4**
3	0	1	1	1	1	0	**6**
4	1	0	0	0	0	1	**1**
5	1	0	1	0	1	1	**3**
6	1	1	0	1	0	1	**5**
7	1	1	1	1	1	1	**7**

2.1.4 nD-omega network

The nD-omega network is a dynamic multistage interconnection network of n stages which interconnects 2^n inputs u = 0, 1, …, 2^n-1to 2^n outputs v = 0, 1, …, 2^n-1and which is organized as follows:

- each stage is constituted of 2^{n-1} 2x2-switches and interconnects 2^n inputs to 2^n outputs numbered 0, 1, …, 2^n-1,

- each 2x2-switch is a simple exchange element which can be set either straight or crossed to interconnect the two inputs u and u+1 of its stage, either straightforwardly or in crossing them, to the two outputs u and u+1 of its stage,

- the output u of the k^{th} stage is connected to the input $\sigma(u)$ of the $(k+1)^{th}$ stage.

Fig. 2 illustrates such a network for n=3. The top 2x2-switch of the 1^{st} (resp. 3^{rd}) stage is set straight (resp. crossed).

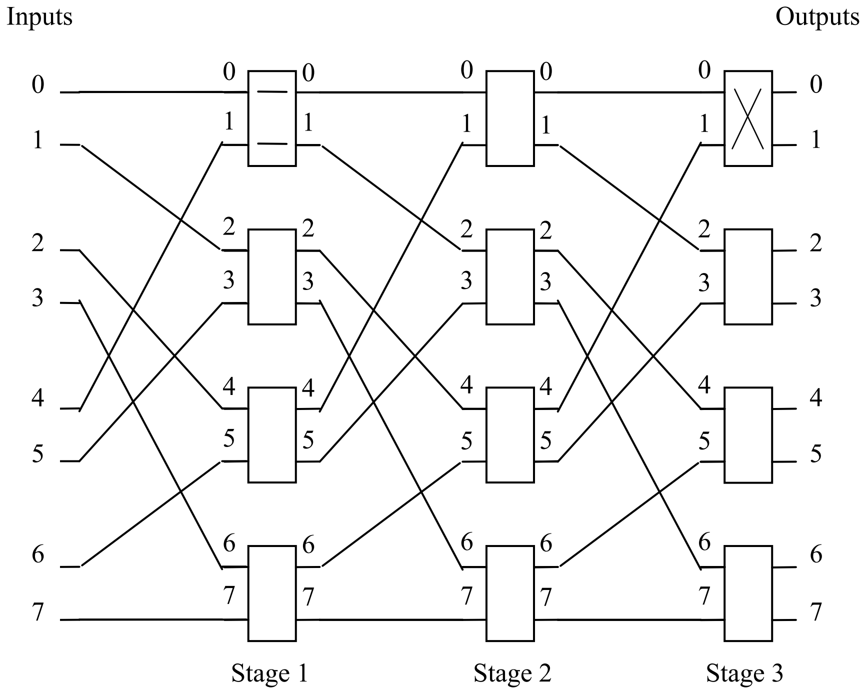

Figure 2. The 3D-omega interconnection network.

It comes from this interconnection logic that each of the 2x2-switches can be viewed both as:

- a node of the balanced binary tree rooted at a 2x2-switch of the 1^{st} stage and whose the leaves are the nD-omega outputs numbered from top to down from 0 to 2^n-1 by an up-down concatenation of the branches edges labels (0 (resp. 1) for an up (resp. a down) edge),

a node of the balanced binary tree rooted at a 2x2-switch of the n^{th} stage and whose the leaves are the nD-omega inputs numbered from top to down $\xi(0)$, $\xi(1)$, …, $\xi(2^n-1)$ where $\xi(u) = (u_0u_1\ldots u_{n-1})_2$ by an down-up concatenation of the branches edges labels (0 (resp. 1) for an up (resp. a down) edge).

2.1.5 nD-inverse omega network

The nD-inverse omega network is the network obtained in inverting the nD-omega network such that its inputs (resp. outputs) become the outputs (resp. inputs).

2.1.6 One pass routing

An one pass routing of a permutation π on the inputs of a nD-omega network consists in defining 2^n paths each of which conveys one message from the input u to the output $\pi(u)$ such that there are no paths which share a same stages interconnection link.

As a result of the tree structure of its 2x2-switches, the routing in nD-omega network verifies the following properties:

- the route which connects the input u to the output v is obtained in connecting, at each stage k, the input link of the corresponding 2x2-switches on the route to its top (resp. down) output link if $v_{n-k} = 0$ (resp. 1).
- for any input u and any output v, there is only one route which connects u to v.
- not all permutations on nD-omega networks are routable in one pass.

2.1.7 Omega permutation

An omega permutation is a permutation on the nD-hypercube which can be routed in one pass on the nD-omega network.

2.1.8 Inverse omega permutation

An inverse omega permutation is a permutation on the nD-hypercube which can be routed in one pass on the inverse nD-omega network.

2.2 MIMD queueless routing

Let π be a permutation on a nD-hypercube network with bidirectional links and a set of 2^n messages of the same size each one located at one node u and destined for the node $\pi(u)$. Routing π under Multiple Instruction Multiple Data (MIMD) queueless communication model [7] consists in conveying all the messages to their respective destination such that:

- there is no restriction that all communications in a given routing step must occur in the same hypercube link direction,
- in a given hypercube link only one message may be sent in a given direction per routing step,
- each hypercube node may store only a single message between routing steps and then needs only a queue of size 1 to hold each message.

Clearly, MIMD queueless routing a permutation on a hypercube consists in a sequence of global and synchronous exchanges of messages between neighbour nodes such that no more than one message is located at each node after each exchange step.

Because messages have the same size, the complexity of such a routing is of the order of the number of required exchange steps. Therefore an optimal routing is the one with the minimal exchange steps. For an arbitrary permutation on a nD-hypercube it is well known, from e-cube routing [8], that this number is at least equal to n.

In this paper we address the problem of optimal MIMD queueless routing of omega and inverse-omega permutations on nD-hypercube.

3. Related works

Optimal routing of permutations on nD-hypercubes is, since a quarter of century, one of the most challenging open problems in the theory of IN. It has so been extensively well studied and, under several communication models, several routing paradigms have been proposed in the literature. In [9] Szimansky considers the offline routing in circuit-switched and packet switched commutation models under all-port MIMD communication model. Under the circuit-switched hypothesis he proves that, for $n \leq 3$, any hypercube is rearrangeable that is it can route any permutation on disjoint paths. Under packet-switched hypothesis he shows that routing can be made in 2n-1 steps result which has been then improved to 2n-3 in [10] by Shen et al under the assumption that each link is used at most twice. Under the single port MIMD communication model, Zhang in [11] proposes a routing in O(n) steps on a spanning tree of the hypercube. In [12] and [13] Hwang et al consider online oblivious routing under buffered all port MIMD communication models and prove that n steps routing is possible for $n \leq 12$ using local information. In [14] Vöcking proves that deterministic offline routing in buffered all port MIMD model can be done in $n+O(\sqrt{n\log n})$ steps while online oblivious randomized routing can be made in $n+O(n/\log n)$ steps.

For the more restrictive conditions, that is single-port, queueless, and MIMD communication model two classes of works can be distinguished: the class of works which tackle the problem for arbitrary permutations and the ones which tackle it for particular permutations among which the omega permutations.

For arbitrary permutations, the personal communication of Coperman to Ramaras reported by Ramras in [15] and the works of Ramras [15] constitute certainly the leading ones. Indeed while Coperman gives the computational proof that arbitrary permutations can be routed in 3D-hypercube in 3 steps, Ramras proves that if an arbitrary permutation can be routed in r steps in rD-hypercube, then for $n \geq r$ arbitrary permutations on nD-hypercubes can be routed in 2n–r steps. Thus, it follows that arbitrary permutations on nD-hypercubes can be routed in 2n-3 steps; improving so the 2n-1 routing steps of Gu and Tamaki [16]. Recently, Laing and Krumme in [7] have introduced an approach which simplifies the problem enough to permit a human verification of the possibility of routing in 3 steps arbitrary permutations on 3D-hypercube and computer verification for the 4 steps routing in 4D-hypercube. The approach is based on the concept of k-separability of a permutation which is the possibility to partition a permutation after k steps of routing into 2^k permutations on disjoints (n-k)D-hypercubes. For n=3 they identify three classes of 1-inseparable permutations for which they exhibit 3 steps routing. From experimental results based on the same paradigm they conjecture that in 4D-hypercube arbitrary permutations can be routed in 4 steps.

To our knowledge, routing of omega permutations has been first, implicitly, studied by Schwartz. Indeed he proved in [17] that it is possible to route packings that is monotonic permutations on nD-hypercubes in O(n) steps with O(n) bits storage at each node. The resulting algorithm consists, for dimension i = 0, …, n-1 and for each node detaining a message which wants to cross dimension i, in sending it across this dimension. Then Kuszmaul [18], which noticed that naive implementation of this algorithm requires $O(n^2)$ routing steps, on the basis that the same result applies as well for semi-contractions from a personal communication of J. Rose reported by Kuszmaul in [18], reduced it to O(n) steps with data pipelining technique. More recently, a study of the permutation capability of a binary hypercube under the commonly used dimension-order routing was presented by Veselovsky and Batovski [19]. They explored two modes of the basic routing algorithm based on non-skipping or skipping identical bits in source and destination addresses when adjusting a route as in algorithm proposed by Schwartz. The whole study was done computationally. It has been found that the skip mode in comparison with the non-skip one provides better permutation capability; its beneficial effect on low-dimensional hypercubes, especially concerned with the routing of the so called bit-permute-complement (BPC) permutations [20], is evident. The possibility of conflict-free routing of the most frequently used permutations under dimension-order routing was also tested.

In the remainder of this paper we address the problem of devising an algorithm for routing optimally omega permutations in a MIMD queueless communication models on nD-hypercubes. Unlike the others approaches, the class of algorithms we are looking for is the one which exploits the incremental construction of nD-hypercubes as an interconnection in some well suited dimension of two (n-1)D-hypercubes. The approach is based on the concept of k-separability from Laing and Krumme we call partitionability. However instead of being computational our approach is purely analytical. We first prove that any omega permutation is 1-separable. As the optimal routing of arbitrary permutations is still an open problem, we prove that only the partition in the dimension n-1 guarantees to any omega permutation to yield two distinct omega permutations on disjoint (n-1)D-hypercubes. So any omega permutation can be optimally routed recursively by successive partitions in the dimension n-1 and then requires at most n exchanges steps each of which is completely defined by the partition process.

4. Mathematical Foundations

The partition process which is the foundation of our routing algorithm seems like to the computation of a perfect matching in bipartite graph. So in this section we first recall some basic notions and results on bipartite graphs and maximum matching of the nodes of a bipartite graph. Then we describe the mathematical foundation of the partition process.

4.1 Definitions and notations

4.1.1 Bipartite graphs

A bipartite graph is a triplet $G = (V_1, V_2, E)$ where:

- V_1 and V_2 are disjoint set of nodes,
- E is a set of edges $\{u, v\} \in V_1 x V_2$

The bipartite graph associated to the nD-hypercube is the one where:

- V_1 and V_2 are two disjoint copies of the hypercube nodes
- E is the set of edges $\{u, v\} \in V_1 x V_2$ such that $u = v$ or $\{u, v\}$ is an edge of the hypercube.

4.1.2 Adjacency matrix

The adjacency matrix of a graph is the matrix M whose rows and columns are indexed by the graph nodes and each component M[u,v]=1 (resp. 0) if $\{u, v\} \in$(resp. $\notin$) E.

Two matrices A and B of the same order and whose rows and columns are indexed by the nodes of a nD-hypercube are said identical if there is a bijection ρ (resp. φ) of A rows (resp. columns) indexes on B rows (resp. columns) such that for any u belonging to A rows indexes and v belonging to A columns indexes, $A[u, v]= B[\rho(u), \varphi(v)]$.

Network Protocols and Algorithms
ISSN 1943-3581
2011, Vol. 3, No. 1

Let $D_{x,i}$ be the set of the hypercube nodes u such that u_i=x. We can observe that the adjacency matrix of the bipartite graph associated to the nD-hypercube can be viewed as the adjacency matrices couple ($M_{x,i}$, x=0,1) of the bipartite graphs couple ((V, $D_{x,i}$, E); x=0,1), for any dimension i. Table 2 illustrates, for n=4 and i=2 such a couple of adjacency matrices. Let's remark that $M_{0,i}$ and $M_{1,i}$, are identical; take ρ=φ and for any u, $\rho(u)=u\oplus 2^i$ where ⊕ is the bitwise XOR operator.

Table 2. Structure of The adjacency matrices of the bipartite graphs ((V, D0,2, E); x=0, 1) for the 4D-hypercube. The left upper corner number stands for the dimension

2	0	1	2	3	8	9	10	11
0	1	1	1		1			
1	1	1		1		1		
2	1		1	1			1	
3		1	1	1				1
8	1				1	1	1	
9		1			1	1		1
10			1		1		1	1
11				1		1	1	1
4	1							
5		1						
6			1					
7				1				
12					1			
13						1		
14							1	
15								1

2	4	5	6	7	12	13	14	15
4	1	1	1		1			
5	1	1		1		1		
6	1		1	1			1	
7		1	1	1				1
12	1				1	1	1	
13		1			1	1		1
14			1		1		1	1
15				1		1	1	1
0	1							
1		1						
2			1					
3				1				
8					1			
9						1		
10							1	
11								1

4.1.3 Matching of a bipartite graph

A matching of a bipartite graph G is one-to-one mapping Γ which associates each node u of a subset of V_1 with a node Γ(u) of V_2 such that the edges {u, Γ (u)} ∊ E and are pairwise non adjacent. A matching Γ is said to be maximum when its cardinality, |Γ|, is maximum. A matching Γ is said to be perfect when $|\Gamma|=|V_1|=|V_2|$.

The computation of a maximum matching is one of the main problems in the study of bipartite graphs. The main results about this computation are due to C. Berge characterisation of maximum matching [21] and König-Hall characterization of matching which saturates a subset of V_1 [22]. In the sequel we will use the simplest implementation of the C. Berge theorem due to Neiman [23] which proceeds by distinguishing one and only one "1" by row and by column in the adjacency matrix.

Network Protocols and Algorithms
ISSN 1943-3581
2011, Vol. 3, No. 1

4.1.4 Partitionable permutations

Given a permutation π on $H^{(n)}$, for x = 0, 1 and $0 \le i \le n-1$, let:

- n_opt (=n without any contraindication) be the minimal number of routing steps of an arbitrary permutation on a nD-hypercube,
- $S_{x,i}$ be the set of the nodes u of $H^{(n)}$ such that $\pi_i(u) = x$, that is $\pi(u) \in D_{\underline{x},i} = H^{(n)}_{x,i}$,
- G_{xi} be the bipartite graph ($S_{x,i}$, $D_{\underline{x},i}$, E),
- N_{xi} be the adjacency matrix of G_{xi}.

A permutation π on $H^{(n)}$ is said to be partitionable in a dimension i of $H^{(n)}$ if there is a permutation Γ on $H^{(n)}$ such that:

- Γ is routable in one step,
- the bijection α (resp. β) which associates π(u) with Γ(u) such that $\Gamma_i(u)=0$ (resp. 1) is a permutation on $H^{(n)}_{0,i}$ (resp. $H^{(n)}_{1,i}$)

Examples: (0,4,8,12,1,5,9,13,2,6,10,14,3,7,11,15) is partitionable in any dimension of the 4D-hypercube. For instance in dimension 0, the adjacency matrices of the induced bipartite graphs $G_{x,i}$ are given in Table 3.

Table 3. The adjacency matrices of the bipartite graphs induced by the permutation π = (0,4,8,12,1,5,9,13,2,6,10,14,3,7,11,15). The boldfaced "1" stand for the ones distinguished by the Neiman matching algorithm.

	0	2	4	6	8	10	12	14
0	1	1	**1**	0	1	0	0	0
2	1	1	0	**1**	0	1	0	0
8	1	0	0	0	1	1	**1**	0
10	0	1	0	0	1	1	0	**1**
1	**1**	0	0	0	0	0	0	0
3	0	**1**	0	0	0	0	0	0
9	0	0	0	0	**1**	0	0	0
11	0	0	0	0	0	**1**	0	0

	1	3	5	7	9	11	13	15
5	**1**	0	1	1	0	0	1	0
7	0	**1**	1	1	0	0	0	1
13	0	0	1	0	**1**	0	1	1
15	0	0	0	1	0	**1**	1	1
4	0	0	**1**	0	0	0	0	0
6	0	0	0	**1**	0	0	0	0
12	0	0	0	0	0	0	**1**	0
14	0	0	0	0	0	0	0	**1**

From where:

- Γ=(4, 0, 6, 2, 5, 1, 7, 3, 12, 8, 14, 10, 13, 9, 15, 11) defined on $H^{(4)}$ is routable in one step,
- the induced permutations α=(4, 12, 0, 8, 6, 14, 2, 10) on 3D-hypercube $H^{(4)}_{00}$ and β=(5, 13, 1, 9, 7, 15, 3, 11) defined on 3D-hypercube $H^{(4)}_{1,0}$ are both routable in at most 3 steps as any arbitrary permutation on $H^{(3)}$ is routable on $H^{(3)}$ in at most 3 steps [15, 7].

On the contrary (7,14,15,13,11,10,9,12,2,6,5,4,0,3,1,8) is not partitionable. Indeed there is no dimension in which the adjacency matrices of the induced bipartite graphs $G_{x,i}$ both admit a maximum matching. It can be easily verified that in dimension i=0 (resp. 1, 2 and 3) column 2 (resp. 0, 3 and 7) of $G_{0,i}$ adjacency matrix is null.

4.2 Characterization of omega and inverse-omega permutations

The prefix (resp. suffix) of two hypercube nodes is the longest common most (resp. least) significant bits of their binary addresses.

Let s(u, v) (resp. p(u, v)) the length of nodes u and v prefix (resp. suffix).

Proposition 1: A necessary and sufficient condition for a permutation π on $H^{(n)}$ to be an omega (resp. inverse omega) permutation is that for any couple (u, v) of distinct hypercube nodes, $s(u, v) + p(\pi(u), \pi(v))$ (resp. $p(u, v) + s(\pi(u), \pi(v)) < n$.

Proof: The condition is necessary. Let π be an omega permutation on the nD-hypercube and for any couple (u, v) of distinct omega inputs let:

- $\mathbb{S}$ (resp. $\mathbb{D}$) be the root of the smallest subtree whose leaves contain the inputs (resp. outputs) u and v (resp. $\pi(u)$ and $\pi(v)$),
- $h_{\mathbb{S}}$ (resp. $h_{\mathbb{D}}$) the height of the subtree rooted in $\mathbb{S}$ (resp. $\mathbb{D}$).

Let's observe that $\mathbb{S}$ (resp. $\mathbb{D}$) could not exist. In fact in such a case the corresponding root is out of the scope of the considered omega network and appropriately we set the height of the corresponding subtree to a some number, say nmax > n. On the contrary if $\mathbb{S}$ and $\mathbb{D}$ do exist then $h_{\mathbb{S}}$ and $h_{\mathbb{D}}$ are such that $h_{\mathbb{S}} + h_{\mathbb{D}} = n+1$ as the number of hops from u (resp. v) to $\pi(u)$ (resp. $\pi(v)$) in the omega network is n+1. $\mathbb{S}$ being a node of a balanced binary tree rooted at a 2x2-switch, say $\mathbb{R}_{\mathbb{S}}$, of the n^{th} stage, and according to its leaves labelling, s(u, v) is the length of the of the common edges labels from $\mathbb{S}$ to $\mathbb{R}_{\mathbb{S}}$. Necessarily $s(u, v) \leq \max(0, n-h_{\mathbb{S}})$. Similarly, $\mathbb{D}$ being a node of a balanced binary tree rooted at a 2x2-switch of the 1^{st} stage, say $\mathbb{R}_{\mathbb{D}}$ and according to its leaves labelling, $p(\pi(u), \pi(v))$ is the length of the common edges labels from $\mathbb{D}$ to $\mathbb{R}_{\mathbb{D}}$. Necessarily $p(\pi(u), \pi(v)) \leq \max(0, n-h_{\mathbb{D}})$. Thus we have $s(u, v) + p(\pi(u), \pi(v)) \leq \max(0, n-h_{\mathbb{S}})+ \max(0, n-h_{\mathbb{D}})$. By hypothesis π can be routed through the nD-Omega network in one pass without conflicts between its routes. Therefore each one of the 2x2-switches is set either straight or crossed and consequently four situations can arise.

<u>Case 1</u>: $\mathbb{S}$ and $\mathbb{D}$ do exist. $s(u, v) + p(\pi(u), \pi(v)) \leq n-h_{\mathbb{S}} + n-h_{\mathbb{D}} = n-1< n$.

<u>Case 2</u>: $\mathbb{S}$ does exist and $\mathbb{D}$ does not. $s(u, v) + p(\pi(u), \pi(v)) \leq n-h_{\mathbb{S}} < n$.

<u>Case 3</u>: $\mathbb{S}$ does not exist and $\mathbb{D}$ does. $s(u, v) + p(\pi(u), \pi(v)) \leq n-h_{\mathbb{D}} < n$.

<u>Case 4</u>: neither $\mathbb{S}$ nor $\mathbb{D}$ does exist. $s(u, v) + p(\pi(u), \pi(v)) = 0 < n$.

The condition is sufficient. Let's suppose that π is not an omega permutation. There are two distinct inputs u and v and a 2x2-switch, say $\mathbb{S}$, such that the routes which connect u to $\pi(u)$ and v to $\pi(v)$ share a same $\mathbb{S}$ output link. As π is a permutation $\pi(u)$ and $\pi(v)$ are distinct and then there is a distinct 2x2-switch, say $\mathbb{D}$ from which the two routes diverge. Again let $h_{\mathbb{S}}$ (resp. $h_{\mathbb{D}}$) the height of the subtree rooted in $\mathbb{S}$ (resp. $\mathbb{D}$). By the same reasoning as in the proof of the necessary condition, $s(u, v) = n-h_{\mathbb{S}}$ and $p(\pi(u), \pi(v)) = n-h_{\mathbb{D}}$ while $h_{\mathbb{S}} + h_{\mathbb{D}} \leq n$. Then $s(u, v) + p(\pi(u), \pi(v)) = 2n – (h_{\mathbb{S}} + h_{\mathbb{D}}) \geq n$.

4.3 Characterization of partitionable permutations

The partionability of a permutation on $H^{(n)}$ is a guaranty that it can be decomposed, after at most one exchange step, in two independent permutations each one on a distinct (n-1)D hypercube.

In [24] Jung and Sakho prove the following characterization of arbitrary partitionable permutations.

Proposition 2: A necessary and sufficient condition for a permutation on a nD-hypercube to be partitionable is that there is a dimension i such that for any x=0, 1, $N_{x,i}$ does contain no null column.

Some remarkable partitionable permutations are the ones for which $S_{x,i}=D_{x,i}$ or $S_{x,i}\cap D_{x,i}=\varnothing$. Therefore, in the sequel we will only consider permutations such that $S_{x,i}\subset D_{x,i}$.

5. Optimal routing of omega permutations

In this section we deal with the declination of the above characterization of partitionable permutations for omega permutations. In fact we show that any omega permutation is partitionable. Then we analyse the structure of omega permutations to show that there exists recursive partitioning which leads to an optimal routing algorithm. Before to do this, let's first characterize partitionable omega permutations.

5.1 Characterization of partitionable omega permutations

To characterize the partitionable omega permutations, let's consider the adjacency matrices $N_{x,i}$ for x=0,1 induced in a dimension i by an omega permutation π on a nD-hypercube $H^{(n)}$.

Lemma 1: For any dimension i, the adjacency matrices $N_{x,i}$ x = 0, 1 are identical.

Proof: It is straightforward. It comes from the fact that $M_{x,i}$ x=0,1 are identical. Indeed as π is an omega permutation, two (n-1)-suffixed sources nodes can not have 1-prefixed destinations nodes. So, for any node u of $H^{(n)}$, $\pi(u)\in D_{x,i}$ for x=0,1 if and only if $\pi(u\oplus 2^{n-1})\in D_{\underline{x},i}$. In others words, $u\in S_{x,i}$ if and only if $u\oplus 2^{n-1}\in S_{\underline{x},i}$. For any $u\in S_{x,i}$, and $v\in D_{x,i}$, $N_{x,i}[u, v] = M_{x,i}[u, v] = M_{\underline{x},i}[u\oplus 2^i, v\oplus 2^i] = N_{\underline{x},i}[u, v]$.

Proposition 3: Any omega permutation on a nD-hypercube is partitionable in any dimension.

Proof: According to Proposition 2 we have to prove that for any dimension i, the adjacency matrices $N_{x,i}$ x=0,1 do contain no null column. Given Lemma 1, we will consider only the matrix $N_{0,i}$, the case of $N_{\underline{1},i}$ being induced from the two matrices identity. We proceed by absurd. Let's suppose that $N_{0,i}$ contains null columns and let v the index of such a column. $N_{x,i}$ being constituted of $M_{x,i}$ rows, there is a permutation of its rows and a permutation of its columns such that it can be put in the form of Table 4 matrix where x=0, i=n-1.

Table 4. The structure of the adjacency matrix of the bipartite graph $G_{0,i}=(S_{0,i},D_{0,i},E)$.

O (resp. I) stands for the null (resp. identity) matrix.

<table>
<tr><td colspan="2"></td><td colspan="6">$D_{0,i}$</td></tr>
<tr><td colspan="2"></td><td>...</td><td>v</td><td>...</td><td>...</td><td>$u\oplus 2^i$</td><td>...</td></tr>
<tr><td>$S_{0,i}\cap D_{0,i}$</td><td>...</td><td colspan="3">A</td><td colspan="3">B</td></tr>
<tr><td rowspan="3">$S_{0,i}\cap D_{1,i}$</td><td>...</td><td colspan="3" rowspan="3">O</td><td colspan="3" rowspan="3">I</td></tr>
<tr><td>u</td></tr>
<tr><td>...</td></tr>
</table>

As v indexes a null column, necessarily this column belongs to the block $[A, O]^T$. Furthermore we have the following:

- $\pi_i(v)=1$. As v indexes a column of $[A, O]^T$, $N_{0,i}[u, v]=0$ for $u\in S_{0,i}\cap D_{0,i}$ and then $v\notin S_{0,i}\cap D_{0,i}$ because if this was the case, we would have $N_{0,i}[v, v]= M_{0,i}[v, v]=1$. As on the other side $v\in D_{0,i}$, it follows that $v\notin S_{0,i}$. On the contrary $v\in S_{1,i}$, from where $\pi(v)\in D_{1,n-1}$ and then $\pi_i(v)=1$.

- $\pi_i(v\oplus 2^{n-1})=1$. As v indexes a column of $[A, O]^T$, $v\notin\{u\oplus 2^i : u\in S_{0,i}\cap D_{1,i}\}$. Thus $v\neq u\oplus 2^i$ for any $u\in S_{0,i}\cap D_{1,i}$. From where $v\oplus 2^i\neq u$ for any $u\in S_{0,i}\cap D_{1,i}$ and then $v\oplus 2^i\notin S_{0,i}\cap D_{1,i}$. As $v\in D_{0,i}$, $v\oplus 2^i\in D_{1,i}$. So, at the same time, we have $v\oplus 2^i\in D_{1,i}$ and $v\oplus 2^i\notin S_{0,i}\cap D_{1,i}$ from where $v\oplus 2^i\notin S_{0,i}$. Necessarily $v\oplus 2^i\in S_{1,i}$, $\pi(v\oplus 2^i)\in D_{1,i}$ and then $\pi_i(v\oplus 2^i)=1$.

But $\pi_i(v) = 1$ and $\pi_i(v\oplus 2^i) = 1$ imply $s(v, v\oplus 2^i) + p(\pi(v), \pi(v\oplus 2^i)) = n$; this contradicts the fact that π is an omega permutation.

At this level of our study, a natural question concerns the routability and, if so, in how many steps of the permutations resulting from a partition of an omega permutation. Let π be an omega permutation on a nD-hypercube $H^{(n)}$, Γ be the partition of π in a dimension i, α (resp.β) be the permutation on $H^{(n)}{}_{0,i}$ (resp. $H^{(n)}{}_{1,i}$) induced by Γ. First of all, let's observe that $\Gamma = (\Gamma^0, \Gamma^1)$ where Γ^0 (resp. Γ^1) is a perfect matching of the bipartite graph $G_{0,i}$ (resp. $G_{1,i}$) and that it can be chosen, this will be the case in the sequel, such that Γ^0 and Γ^1 are identical as $N_{0,i}$ and $N_{1,i}$ are identical. So we again consider only Γ^0, the same reasoning applying to Γ^1. The structure of $N_{0,i}$, see Table 4, imposes $\Gamma^0=(\Gamma^{0,I}, \Gamma^{0,A})$ where $\Gamma^{0,I}$ (resp. $\Gamma^{0,A}$) is a perfect matching of the bipartite graph having I (resp. A) as adjacency matrix.

I being the identity matrix, $\Gamma^{0,I}$ is unique. It is the set of the couples (u, v) such that I[u, v]=1, that is the interconnection links of the nodes u of $S_{0,i}\cap D_{1,n-1}$ to the nodes $u\oplus 2^i$ of $D_{0,i}$. More formally, $\Gamma^{0,I}$ associates any $u\in S_{0,n-1}$ such that $u_{n-1}=1$, with $\Gamma^{0,I}(u)=u\oplus 2^i$. Unlike $\Gamma^{0,I}$, $\Gamma^{0,A}$ is not necessarily unique and then any of them suits to constitute with $\Gamma^{0,I}$ a perfect matching of $G_{0,i}$. However does any of them still suits to make α routable in n-1 steps on $H^{(n)}{}_{0,i}$? The response to this question remains related to the more general and open problem of routing arbitrary permutations on $H^{(n)}$ in at most n steps. To get round this difficulty, according to Proposition 3, an interesting alternative approach is the search of $\Gamma^{0,A}$ which assures α to be an omega permutation.

Within this purpose, let y and z be two nodes of $H^{(n)}{}_{0,i}$. There are two nodes u and $v\in H^{(n)}$

such that $y=\Gamma(u)$ and $z=\Gamma(v)$. Then

$$s(y, z)+p(\alpha(y), \alpha(z)) = s(\Gamma(u), \Gamma(v))+p(\alpha(\Gamma(u)), \alpha(\Gamma(v)))$$

where Γ is such that:

- $u=(u_{n-1}u_{n-2}\ldots u_i\ldots u_0) \rightarrow \Gamma(u)=(\Gamma_{n-1}(u)\Gamma_{n-2}(u)\ldots\Gamma_i(u)\ldots\Gamma_0(u))$
- $v=(v_{n-1}v_{n-2}\ldots v_i\ldots v_0) \rightarrow \Gamma(v)=(\Gamma_{n-1}(v)\Gamma_{n-2}(v)\ldots\Gamma_i(v)\ldots\Gamma_0(v))$

and α is such that:

- $\Gamma(u)=(\Gamma_{n-1}(u)\Gamma_{n-2}(u)\ldots\Gamma_{i+1}(u)0\Gamma_{i-1}(u)\ldots\Gamma_0(u))\rightarrow\alpha(\Gamma(u))=(\pi_{n-1}(u)\pi_{n-2}(u)\ldots\pi_{i+1}(u)0\pi_{i-1}(u)\ldots\pi_0(u))$
- $\Gamma(v)=(\Gamma_{n-1}(v)\Gamma_{n-2}(v)\ldots\Gamma_{i+1}(v)0\Gamma_{i-1}(v)\ldots\Gamma_0(v))\rightarrow\alpha(\Gamma(v))=(\pi_{n-1}(v)\pi_{n-2}(v)\ldots\pi_{i+1}(v)0\pi_{i-1}(v)\ldots\pi_0(v))$

As $\alpha(\Gamma(u))$ and $\alpha(\Gamma(v))\in H^{(n)}{}_{0,i}$,

$$p(\alpha(\Gamma(u)), \alpha(\Gamma(v))) = \begin{cases} p(\pi(u), \pi(v)) - 1 & \text{if } i = n\text{-}1 \\ p(\pi(u), \pi(v)) & \text{otherwise.} \end{cases}$$

Now let's consider $s(\Gamma(u), \Gamma(v))$. One of the following situations may happen.

- $s(\Gamma(u), \Gamma(v)) \leq s(u, v)$:

$$s(x, y)+p(\alpha(x), \alpha(y)) \leq s(u, v) + \begin{cases} p(\pi(u), \pi(v)) - 1 < n - 1 & \text{if } i = n\text{-}1 \\ p(\pi(u), \pi(v)) < n & \text{otherwise.} \end{cases}$$

  Only $i = n-1$ assures α to be an omega permutation.

- $s(\Gamma(u), \Gamma(v)) > s(u, v)$:

$$s(x, y)+p(\alpha(x), \alpha(y)) > s(u, v) + \begin{cases} p(\pi(u), \pi(v)) - 1 < n - 1 & \text{if } i = n\text{-}1 \\ p(\pi(u), \pi(v)) < n & \text{otherwise.} \end{cases}$$

No dimension assures α to be an omega permutation.

This analysis can be summarized in the following proposition.

Proposition 4: If Γ is a partition in dimension n-1 of an omega permutation such that $s(\Gamma(u), \Gamma(v)) \leq s(u, v)$ for any couple of nodes (u, v) then the permutation α (resp. β) which associates $\pi(u)$ with $\Gamma(u)$ such that $\Gamma_i(u)=0$ (resp. 1) is an omega permutation too.

5.2 Partition of omega permutations

Now, we deal with the problem of the existence of a partition Γ which satisfies the conditions of Proposition 4 and, if so, the computation of one of its instances. Again we will restrict our analysis to Γ^0 . It should be noticed that as $\Gamma^{0,I}$ is unique only $\Gamma^{0,A}$ is concerned. More formally we will discuss about the existence of $\Gamma^{0,A}$ such that:

- $s(\Gamma^{0,A}(u), \Gamma^{0,A}(v)) \leq s(u, v)$ for any couple (u, v) of nodes of $S_{0,n-1}\cap D_{0,n-1}$,

- $s(\Gamma^{0,A}(u), \Gamma^{0,I}(v)) \leq s(u, v)$ for any $u \in S_{0,n-1} \cap D_{0,n-1}$ and $v \in S_{0,n-1} \cap D_{1,n-1}$.

Within this purpose we will again refer to Table 4. Let $u \in S_{0,n-1} \cap D_{0,n-1}$. As π is an omega permutation on $H^{(n)}$, for any $v \in S_{0,n-1} \cap D_{1,n-1}$, $v \neq u \oplus 2^{n-1}$ therefore $u \neq v \oplus 2^{n-1}$. As $u \in D_{0,n-1}$, necessarily $u \in D_{0,n-1} - \Gamma^{0,I}(S_{0,n-1} \cap D_{1,n-1})$.

Inversely, let $u \in D_{0,n-1} - \Gamma^{0,I}(S_{0,n-1} \cap D_{1,n-1})$. In others words, $S_{0,n-1} \cap D_{0,n-1} \subseteq D_{0,n-1} - \Gamma^{0,I}(S_{0,n-1} \cap D_{1,n-1})$. As the two sets are of a same cardinality, $S_{0,n-1} \cap D_{0,n-1} = D_{0,n-1} - \Gamma^{0,I}(S_{0,n-1} \cap D_{1,n-1})$. A being a block of the adjacency matrix of $H^{(n)}$, it follows that $A[u, u] = M_{0,n-1}[u, u] = 1$ for any $u \in S_{0,n-1} \cap D_{0,n-1}$.

Then let's consider the perfect matching $\Gamma^{0,A}$ which associates any $u \in S_{0,n-1} \cap D_{0,n-1}$ with $\Gamma^{0,A}(u) = u$. By construction we have:

- for any couple (u, v) of nodes of $S_{0,n-1} \cap D_{0,n-1}$, $s(\Gamma^{0,A}(u), \Gamma^{0,A}(v)) = s(u, v)$,
- for any $u \in S_{0,n-1} \cap D_{0,n-1}$ and $v \in S_{0,n-1} \cap D_{1,n-1}$, $s(\Gamma^{0,A}(u), \Gamma^{0,I}(v)) = s(u, \Gamma^{I}(v)) = s(u, v \oplus 2^{n-1}) = s(u, v)$

from where the following theorem.

Theorem: The perfect matching Γ of a nD-hypercube which associates any node u with the node u (resp. $u \oplus 2^{n-1}$) if $u_{n-1} =$ (resp. $\neq$) $\pi_{n-1}(u)$ partitions any omega permutation on a nD-hypercube in two different omega permutations on two disjoint (n-1)D-hypercube.

5.3 Optimal routing algorithm of omega permutations

We are now ready to devise an optimal routing algorithm of omega permutations on nD-hypercubes. From the above study such an algorithm proceed recursively by partitioning, in the dimension n-1, the permutation in two different omega permutations on two disjoint (n-1)D-hypercubes. To do so, at each step, the messages located in the nodes which do not belong to the (n-1)D-hypercube of their destination move to this sub-hypercube in dimension n-1 while the messages, which are already located in a node of the (n-1)D-hypercube of their destination stay on this node. Doing so, at most n routing steps are required to route any omega permutation on a nD-hypercube. Then, given an omega permutation π on a nD-hypercube, we obtain the routing algorithm schematized in Fig. 3, 4 and 5 where **Par** is a constructor of parallel actions.

Algorithm Route
Input:
- n: the hypercube dimension,
- π: an omega permutation on a nD-hypercube

Local variables: k: an integer, routing step number
Output: R: a $n \times 2^n$ matrix where R[k, u] is the next location of the outgoing message from u at the step k
Begin

```
    k = 0
    For u = 0 To 2^n-1 Do
        R[k, u] = u
    Compute_route(n, π, k)
End.
```

Figure 3. The omega permutations routing algorithm.

Algorithm Compute_route
Input:
- n: the hypercube dimension,
- π: an omega permutation on a nD-hypercube
- k: routing step number

Local variables:
- α, β: permutations on a (n-1)D-hypercubes
- Γ: permutation on a nD-hypercube

Output: R: a nx2^n matrix
Begin
    If (u = π(u) for u = 0 **To** 2^n-1) **Then**
        Exit
    (Γ, α, β) = Partition(n, π)
    k = k +1
    Par u = 0 **To** 2^n-1 **Do**
        R[k, u] = Γ[u]
    Par
        Compute_route (n-1, α, k)
        Compute_route (n-1, β, k)
End.

Figure 4. The omega permutations routing table computation algorithm.

Algorithm Partition
Input: n: the hypercube dimension,
        π: an omega permutation on a nD-hypercube
Output: Γ: omega permutations on nD-hypercubes
        α, β: omega permutations on (n-1)D-hypercubes
Begin
    Par u = 0 **To** 2^n-1 **Do**
        If ($u_{n-1} \neq \pi_{n-1}(u)$) **Then**
            $\Gamma[u] = u \oplus 2^{n-1}$
        Else
            Γ[u] = u
            Set $\alpha:\Gamma[u]:\Gamma_{n-1}[u]=0 \rightarrow \pi(u)$
            Set $\beta:\Gamma[u]:\Gamma_{n-1}[u]=1 \rightarrow \pi(u)$
End.

Figure 5. The omega permutations partitioning algorithm.

<u>Illustration</u>: Let's consider the omega permutation π = (2,7,4,9,6,15,14,1,8,10,11,5,12,3,0,13) on $H^{(4)}$. The partition of π in dimension 3 is realised by the perfect matching Γ=(0,1,2,11,4,13,14,7,8,9,10,3,12,5,6,11) induced by the boldfaced "1" of Table 5.

Table 5. Perfect matching of the bipartite graphs $G_{x,3}$, x=0,1 induced by the permutation π= (2,7,4,9,6,15,14,1,8,10,11,5,12,3,0,13) on $H^{(4)}$. The boldfaced "1" are the distinguished ones by the Neiman algorithm while double crossed ones are the forbidden matching.

3	**15**	**14**	**13**	**12**	**11**	**10**	**9**	**8**
0	**1**	1	1		1			
1	1	**1**		~~1~~		~~1~~		
2	1		**1**	~~1~~			~~1~~	
4	1				**1**	~~1~~	~~1~~	
7				~~1~~		~~1~~	~~1~~	**1**
11				**1**				
13						**1**		
14							**1**	

3	**15**	**14**	**13**	**12**	**11**	**10**	**9**	**8**
15	**1**	1	1		1			
14	1	**1**		~~1~~		~~1~~		
13	1		**1**	~~1~~			~~1~~	
11	1				**1**	~~1~~	~~1~~	
8				~~1~~		~~1~~	~~1~~	**1**
4				**1**				
2						**1**		
1							**1**	

On the 4D-hypercube this partitioning process results in the data exchanges expressed by the arrows in Fig. 6.

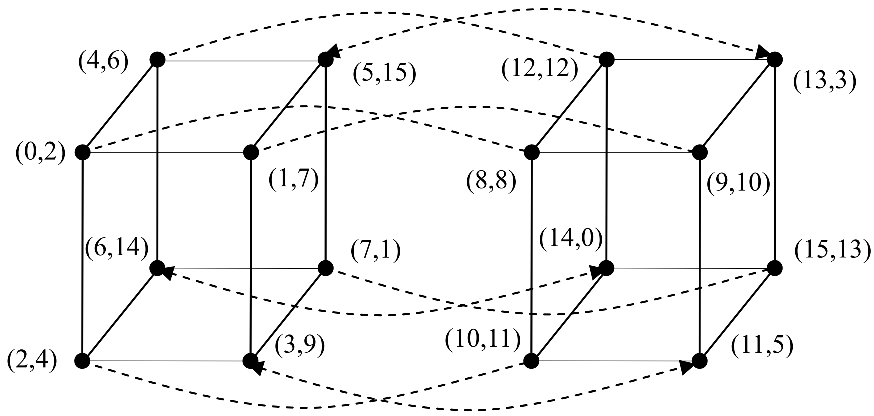

Figure 6. The first step routing of the permutation π=(2,7,4,9,6,15,14,1,8,10,11,5,12,3,0,13). Each node is labelled with the couple constituted of its address and the address of the message it holds. The dotted links are the ones which realise the partition.

We can observe that the resulting permutations α and β, illustrated in Fig. 7, are also omega permutations.

Permutation α   Permutation β

Figure 7. The permutations resulting from the partition induced by Γ=(0,1,2,11,4,13,14,7,8,9,10,3,12,5,6,11).

Again, we can partition α and β in dimension 2, each one in two independent omega permutations on two disjoint 2D-hypercubes, and so on. Table 6 shows the routes R resulting from the entire execution of the algorithm.

Table 6. Routing paths of π. X stands for a don't care transition node for a message which has sooner than the n^{th} routing step definitively attained its destination.

Source	14	7	0	13	2	11	4	1	8	3	9	10	12	15	6	5
Transition nodes	6	7	0	5	2	3	4	1	8	11	9	10	12	15	14	13
	2	3	0	1	6	7	4	5	8	9	11	10	12	13	14	15
	0	1	2	3	4	5	6	7	X	X	10	11	X	X	X	X
	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Destination	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

6. Optimal routing of inverse omega permutations

The routing of an inverse omega permutation π, may be deduced from the routing of the omega permutation π^{-1}. This can be done in two ways. The first way is obvious. Indeed, it suffices to compute the routing paths of π^{-1} and then to inverse them. For the second way, let's recall that an inverse omega permutation π on a nD-hypercube is characterized by the fact that for any couple of nodes (u, v) of the hypercube, p(u, v)+s(π(u), π(v)) < n which in turn implies that for x=0,1, u and $u\oplus 2^0$ can not both belong to $S_{x,0}$. On this basis, all the reasoning carried out in the previous sections on the routing of omega permutations can be repeated according to dimension 0 instead of dimension n-1. It then follows that any inverse omega permutation can be routed on a nD-hypercube by successive partitions in dimension 0. So, following the omega permutations case, each step of the routing algorithm of an inverse omega permutation consists in moving in the dimension 0 the messages which are not yet located on a node of their destination sub-hypercube while the messages located on a node of their destination sub-hypercube stay on this node. As in the case of the omega permutations this routing algorithm requires at most n routing steps.

To illustrate, let's consider π=(14,7,0,13,2,11,4,1,8,3,9,10,12,15,6,5) which is the inverse of the permutation (2,7,4,9,6,15,14,1,8,10,11,5,12,3,0,13). It can be easily verified that the routing paths of π by successive moves in dimension 0 of the suitable messages are the ones given in Table 7.

Table 7. Routing paths of π = (14,7,0,13,2,11,4,1,8,3,9,10,12,15,6,5)

Source	2	7	4	9	6	15	14	1	8	10	11	5	12	3	0	13
Transition nodes	2	7	4	9	6	15	14	1	8	11	10	5	12	3	0	13
	0	5	6	11	4	13	14	3	8	9	10	7	12	1	2	15
	X	1	2	11	X	13	6	7	X	9	10	3	X	5	6	15
	X	X	X	3	X	5	6	X	X	X	X	11	X	13	14	X
Destination	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

Let's remark that the paths of the omega permutation and those of its inverse are not inverse the one of the other.

7. Conclusion and perspectives

This paper has addressed the problem of the optimal routing of omega and inverse omega permutations on nD-hypercubes. The partitioning paradigm is the framework of the proposed routing algorithm. It consists in decomposing recursively a permutation by perfect matching of the bipartite graphs induced by the permutation in two independent permutations on two disjoint (n-1)D-hypercubes.

We have first proved that any omega (resp. inverse omega) permutation can be partitioned in any dimension of the hypercube. Then, given the absence of certainty about the routability of arbitrary permutation, we have proved that there are partitions which assure their induced permutations to be also omega (resp. inverse omega) permutations. The perfect matching which induces such partitions consists for an omega (resp. inverse omega) permutation in matching the nodes detaining messages which are not yet in the sub-hypercube of their destination nodes with their neighbours in dimension n-1 (resp. 0) and the others nodes with themselves. Then any omega (resp. inverse omega) permutation can be routed in at most n steps of data exchange. At each routing step, for any node, the exchange decision is the result of the comparison of the most significant bit of its address with the one of the destination nodes of the message it detains. Clearly the proposed routing algorithm is self routing.

As at each step of the routing a message which is already located in a node of the sub-hypercube of its destination node must stay on this node, one can hope to reduce the number of exchange steps to the necessary one by moving it towards suitable node. In a self routing perspective, this requires a self perfect matching of the host nodes of such messages accordingly to the theorem. In this context, our future works on routing omega permutations on hypercubes will concern the study of the way to perform such a matching.

Beyond omega and inverse omega permutations there are others subclasses of permutations which are also results of shuffle operations on the binary addresses of the nodes of nD-hypercubes and which are subject to great attention. It is for instance the ones which define the interconnection logic of the De Bruijn graphs [25]. So we also plan to test the applicability of the partitioning paradigm for these classes of permutations.

References

[1] Y. Saad and M. H. Schultz, "Topological properties of hypercubes", IEEE Trans. Comput., C-37, 1988, pp. 867-872. http://dx.doi.org/10.1109/12.2234

Y. Liu, J. Han and H. Du., "A Hypercube-based Scalable Interconnection Network for Massively Parallel Computing", Journal of Computers, 3(10) 2008, pp. 58-65. http://dx.doi.org/10.4304/jcp.3.10.58-65

[2] D.C.W. Pao and W.N. Chau, "Design of ATM switch using hypercube with distributed shared input buffers and dedicated output buffers", in Proc. of International Conference on Network Protocols, Tokyo, 7-10 Nov. 1995, pp. 92–99, http://dx.doi.org/10.1109/ICNP.1995.524823

[3] J. S. Park and N. J. Davis, "The Folded Hypercube ATM Switches", in Proc. of IEEE International Conference on Networking, July 2001, Part II, July 2001, pp. 370–379, http://dx.doi.org/10.1007/3-540-47734-9_37.

[4] Chang Wu, Yubai Li , Qicong Peng, Song Chai and Zhongming Yang, "Construction of a multidimensional plane network-on-chip architecture based on the hypercube structure", Progress in Natural Science, Volume 19, Issue 5, 10 May 2009, pp. 635-641, http://dx.doi.org/10.1016/j.pnsc.2008.10.003.

[5] N. Gopalakrishna Kini, M. Sathish Kumar, H.S. Mruthyunjaya, "Analysis and Comparison of Torus Embedded Hypercube Scalable Interconnection Network for Parallel Architecture", IJCSNS International Journal of Computer Science and Network Security, 9(1), January 2009, pp 242-247.

[6] A. K. Laing and D. W. Krumme, "Optimal Permutation Routing for Low-dimensional Hypercubes", Networks 55, 2010, pp. 149-167, http://dx.doi.org/10.1002/net.20325.

[7] J.T. Draper and J. Ghosh, "Multipath e-cube algorithms (MECA) for adaptive wormhole routing and broadcasting in k-ary n-cubes," in Proc. of International Parallel Processing Symposium, 1992, pp. 407.

[8] T. Szymanski, "On the permutation capability of a circuit switched hypercube", in Proc. of International Conference on Parallel Processing, 1989, pp. I-103-109.

[9] X. Shen, Q. Hu, and W. Liang, "Realization of arbitrary permutations on a hypercube," Inf. Pro. Lett., vol. 51(5), 1994, pp. 237-243.

[10] L. Zhang, "Optimal bounds for matching routing on trees," in Proc. of 8th Annual ACM-SIAM Symposium on Discrete Algorithms, 1997, p. 445.

[11] F. Hwang, Y. Yao, and M. Grammatikakis, "A d-move local permutation routing for d-cube," Disc. Appl. Maths., vol. 72, 1997, pp. 199-207. http://dx.doi.org/10.1016/S0166-218X(96)00019-4

[12] F. Hwang, Y. Yao, and B. Dasgupta., "Some permutation routing algorithms for low dimensional hypercubes," Theoretical Computer Science, vol. 270 (1), 2002, pp. 111-124. http://dx.doi.org/10.1016/S0304-3975(00)00279-6

[13] B. Vöcking, "Almost optimal permutation routing on hypercubes," in Proc. of the 33rd Annual ACM- Symposium on Theory of Computing, 2001, pp. 530-539. http://dx.doi.org/10.1145/380752.380848

[14] M. Ramras, "Routing permutations on a graph," Networks, vol 23, pp. 391-398, 1993. http://dx.doi.org/10.1002/net.3230230420

[15] Q-P. Gu and H. Tamaki, "Routing a permutation in the hypercube by two sets of edge disjoint paths," J. of Par. and Distr. comp., vol. 44, 1997, pp. 147-152. http://dx.doi.org/10.1006/jpdc.1997.1358

[16] J. T. Schwartz, "Ultracomputers" ACM Transactions on Programming Languages and Systems, vol. 2(4), 1980, pp. 484-521. http://dx.doi.org/10.1145/357114.357116

[17] B. C. Kuszmaul, "Fast Deterministic Routing on Hypercubes Using Small Buffers" IEEE Transactions on Computers, vol. 39, Nov. 1990, pp. 1390–1393. http://dx.doi.org/10.1109/12.61048

[18] G. Veselovsky and D. A. Batovski, "A study of the permutation capability of a binary hypercube under deterministic dimension-order routing Parallel, Distributed and Network-Based Processing, " in Proc. Eleventh Euromicro Conference 2003, vol. 5-7, p.173-177. http://dx.doi.org/10.1109/EMPDP.2003.1183584

[19] D. Nassimi, S. Sahni, "Optimal BPC permutations on a cube connected SIMD computer," IEEE Trans. Comput., C-31(4), 1982, pp. 338-341.

[20] C. Berge, Graphes, 3ème édition, Dunod, Paris, 1983.

[21] M. Sakarovitch. Techniques mathématiques de la recherche opérationnelle, Université Scientifique et Médicale et Institut National Polytechnique de Grenoble, 1982.

[22] V.I. Neiman. "Structures et commandes des réseaux sans blocage," Annales des Télécom., 1969.

[23] J. P. Jung and I. Sakho, "Graphs Partitioning: An Optimal MIMD Queueless Routing for BPC-Permutations on Hypercubes," in Proc. of PPAM 2009, Part I, LNCS 6067, 2010, pp. 21–30. http://dx.doi.org/10.1007/978-3-642-14390-8_3

[24] Z. F. Ji, and L. G. Ning, "On the de Bruijn-Good graphs," Acta Math. Sinica 30 (2), 1987, pp. 195–205.

Copyright Disclaimer