

Description

Many one dimensional (linear) networks have been proposed for a variety of purposes which have a number of parallel channels which are processed by several distinct stages in a pipelined manner. Each stage can be split into two parts: an interconnect stage where the lines are permuted followed by a layer of two-input two-output processing modules which operates on adjacent pairs of data channels.

Recently, efficient optical computers have been proposed which are networks where the interconnect part is performed using optics and the module part in some other medium, for example a chip of lithium niobate directional couplers (P.Granstrand et al, "Strictly non-blocking 8x8 integrated optical switch matrix" Electron. Lett. 22 No.15 (1986)) or an optoelectronic integrated circuit (J.E. Midwinter, "Novel approach to optically activated wideband switching matrices" IEEE Proc. J 134 261 (1987)). Such machines have the physical layout described above. The use of optics for the interconnect stage has the advantages of high bandwidth, zero time skew and low crosstalk which give the whole processor a high throughput of parallel data. The various interconnection patterns which the networks employ can be generated using bulk or holographic optical components. Such an arrangement does not, however, take full advantage of the parallelism possible with optical systems.

Shing-Hong Lin et al in an article entitled "2-D Optical Multistage Interconnection Networks", SPIE Vol 752 Digital Optical Computing (1987) pp 209-216, describe the use of 2-D networks employing 2-D perfect shuffles interconnects and 2-D four-input, four output processing modules. Lin et al give no indication of the control structure necessary to achieve any particular network interconnection but rather point out that full, 24 cross-bar switches can in principle achieve a desired configuration.

It is an object of the present invention to provide an optical interconnect network having less structural complexity than such known 2-D networks. Accordingly an optical interconnect network comprises at least one stage which has an optical interconnect stage connecting a two-dimensional array of interconnect input ports to a two-dimensional array of interconnect output ports; an array of optical processing modules each having a two-dimensional array of module input ports, optically coupled to a respective interconnect output port, and a two dimensional array of module output ports characterised in that each module is functionally identical to a first and a second pair of two-input, two-output processing sub-modules in which each input of each of the second pair of processing sub-modules is connected to a respective output of the first pair of processing sub-modules, and wherein the inputs of each second processing sub-module are connected to the outputs from different

first processing sub-modules.

The two dimensional network can be assembled using optics, with similar performance as before. There is a limit to the number of channels which can be accommodated which is proportional to either the maximum width of the module element or to the distance across which the optics can image faithfully. If this limit is N channels in the case of a one dimensional network then it becomes N^2 for a two dimensional one. A further advantage of a network according to the present invention is that a network of a given size can be built much more compactly in a two dimensional form.

Because the present invention requires only a sufficient number of processing elements as are necessary to carry out the same processing as the four sub-modules, the structural complexity is reduced from that required to provide a full 4x4 processing yet, as will be shown later, it can be functionally identical to a concatenation of the processing modules of a one dimensional network and have the same control structure.

The present invention therefore allows the use of one dimensional network control structure with two-dimensional optical networks.

An embodiment of the present invention will now be described, by way of example only, with reference to the accompanying drawings, of which:

Figure 1 is a schematic generalised diagram of prior art one-dimensional networks;

Figure 2 is a schematic diagram of an exchange and by-pass module used in an interconnection network as shown in Figure 1;

Figure 3 is a schematic generalised diagram of a two dimensional interconnect network according to the present invention;

Figure 4 is a schematic diagram of an exemplary one-dimensional network which can be configured as a two-dimensional network according to the present invention;

Figure 5 is a schematic diagram of one of the modules of the Figure 3 embodiment of the present invention when reproducing the functionality of the network of Figure 4;

Figure 6 is a schematic diagram to show the functional equivalence of a two dimensional processing module of the present invention to four one dimensional sub-modules; and

Figure 7 is a schematic diagram illustrating the condition for rich stages of a one-dimensional network can be concatenated to form a two-dimensional network according to the present invention.

Referring to Figure 1 is a generalised one-dimensional, prior art interconnect network comprises three stages $S_1 - S_3$ each having a one-dimensional interconnect stage 2 and a layer of two-input, two-output processing modules 4. The interconnections at each

stage can be different as can the functions performed by the modules.

Networks which fall into the general description above can be made to do many parallel processing tasks. The most common function is to have a number of data streams at the input, and to have the output be the same data lines in a different order, the order to be determined by the settings of the modules, keeping the interconnections fixed. Such switching networks usually have regular interconnections patterns, and modules which sit in one of two states, as illustrated in Figures 2a and 2b, in which the two input lines are either exchange or bypassed. A number of networks of this format exist each having a control structure to specify the setting the modules to achieve a certain overall permutation.

Other classes of computer are possible when the modules perform additional processing. If the modules are AND gates and OR gates then a programmable logic arrays can be built by using perfect shuffle or butterfly interconnections. A fast Fourier transform machine has been proposed where the interconnections are perfect shuffles and the modules perform a weighted sum and difference calculation. These and other processes can be achieved in an optical two dimensional network of the configuration of the present invention.

Referring now to Figure 3, a two-dimensional optical interconnect network according to the present invention is shown in generalised form. It comprises three stages S_4 to S_7 each comprising one, two dimensional interconnect stage 6 followed by a two dimensional array of modules 8 each module having a two-dimensional array of four inputs 10 and four outputs 12. The interconnections 6 permute the incoming lines in two dimensions.

Referring now to Figure 4 there is shown a known one-dimensional interconnect network (D.E.Knuth, "Sorting and Searching: Addison Wesley (1973)) which will be used by way of example to explain the steps by which a given one-dimensional network can be reconfigured to a two dimensional network according to the present invention.

The linear network depicted in figure 4 is a sorting network having interconnections 14 which are perfect shuffles and processing sub-modules 16 of which only exemplary ones are referenced for clarity. Several different numbers enter input lines 15; arrowed modules (16) output the higher number at the port pointed to, and unmarked modules 16 always bypass. The numbers emerging at outputs 17 are in numerical from left to right order.

A perfect shuffle interconnection comprises the splitting of the stack of lines into two halves and interleaving them, as shown in Figure 5.

If the 2^4 ports labelled by the four binary bits abcd are shuffled it can be considered as the barrel rolling of the binary addresses of the ports that is the input

port labelled abcd is shuffled to the output port addressed bcda. When a channel arrives at a module at port abcd it can leave from either abc0 or abc1 where the port it leaves depends on what the type of module is and on the destination address of the other channel arriving at the same module.

Consider now the routing effect of two consecutive stages of the sorting network : a line arriving at abcd is shuffled to bcda, exchanged or bypassed in a module (16 or 18) to bcdA where A can be 0 or 1, shuffled again to cdAb and finally switched to cdAB, again where B can be 0 or 1.

abcd
bcda
bcdA
cdAb
cdAB

The line has a choice of four output addresses AB = 00, 01, 10 and 11. Also there are four incoming lines which might emerge from these four ports - those arriving from ab = 00, 01, 10 and 11.

Consider now that the lines in the linear array are arranged in space into a square array according to the rule

abcd $\rightarrow$ (column bd, row ac).

The function of the two stages of the 1D network can be performed by one of the stages of the 2D network shown in figure 3. The new interconnection is a 2D perfect shuffle which is a horizontal perfect shuffle followed by a vertical perfect shuffle. One of these permutations has the effect

(bd, ac) $\rightarrow$ (db, ca)

which corresponds to two 1D shuffles in succession. The effect of the modules in the two consecutive stages of the linear network is to fix the two bits a and b in the address. A module in the two dimensional network operates on the four ports differing in the final bits of the row and column addresses, a and b again, so one 4x4 module can be made to do the same operations as all the relevant modules in the two stages of the old network. That is, by mapping the lines from a 1D to a 2D array according to the rule above, the work of two stages of the original network is done by one stage of the new network incorporating a 2D perfect shuffle, and hence a two dimensional network can be built up using this rule.

Considering now the operation of a 2D module in more detail. Each links ports (do, co), (do, cl), (dl, co) and (dl, cl). First bit a is fixed, then bit b, so the module can be regarded as being split in to four two-input, two-output sub-modules as in figure 6. These sub-modules are four of the modules in figure 4 moved in space, and it is readily deducible which four are associated in each module by using the mapping rule. The 2D module need not be physically separated into the four 2x2 sub-modules as long as the module as whole is functionally identical to such a concatenation of sub-modules as shown in Figure 6.

The exchange/bypass decisions of the original network generate a sorted output of the network, but the output of the equivalent 2D network will have the channel with the nth destination address at the port which is the mapping of port n in the linear array. This scrambling of lines is wholly determined in advance and is functionally unimportant; if it is required to sort into a more visible order then the destination addresses can be modified at the start, or equivalently the modules can be set according to the reverse-mapped addresses.

A single exemplary two dimensional network has been described above. The starting point was a one dimensional network. Consideration will now be given to whether and how any two dimensional network according to the present invention can be constructed from a linear one, independent of the function of the modules and the interconnection pattern.

In analogy to the treatment of the perfect shuffle sorting network above, two consecutive stages can be concatenated into one stage of a 2D network by regarding the modules as being moved around in space. It may not be possible for one interconnection followed by a layer of 4x4 modules to do the work of two, one dimensional stages comprising an interconnection, a 2x2 module as an interconnection and a final 2x2 module sequence. Whatever the function of the modules in figure 2 we know (referring now to figure 7) that the intermediate module output X' is a function of X and Y .

$$X' = X'(X, Y)$$

and similarly for three of the other intermediate outputs

$$Y' = Y'(X, Y)$$

$$W' = W'(W, Z)$$

$$Z' = Z'(W, Z)$$

The final module outputs X'' and Z'' which are functions of X' and Z'

$$X'' = X''(X', Z')$$

$$Z'' = Z''(X', Z')$$

so,

$$X'' = X''(X, Y, Z, W)$$

$$Z'' = Z''(X, Y, Z, W)$$

In order that the functions X'' and Z'' be calculated by one two dimensional, four-input four-output processing module, it is necessary for the other two outputs, Y'' and W'' , to be functions of the same four inputs

$$Y'' = Y''(X, Y, Z, W)$$

$$W'' = W''(X, Y, Z, W)$$

which can only be achieved if intermediate results Y' and W' meet at a common module.

$$Y'' = Y''(Y', W')$$

$$W'' = W''(Y', W')$$

This imposes a restriction on the second of the two interconnects of figure n, which can be summarised as follows -

The complement of a line emerging from a module is defined to be the other line coming out of

it. For any two lines which are the two inputs to a module, the complements of these two lines must also meet at a common module.

If an interconnect does not satisfy this condition then it is not possible to concatenate the consecutive stages of Figure 7 into the functionally identical single stage of two dimensional 4x4 modules according to the present invention.

In principle, for a network satisfying the interconnect condition as stated above, it is always possible to pick up the lines in the linear arrangement of Figure 1 and move them into a two dimensional array as in Figure 3. There will be many ways to do this mapping, but what is sought is a 2D interconnection which can readily be achieved using a simple optical arrangement. An example is the special mapping applied to the perfect shuffle network. If a successful mapping scheme can be found then the task of generating a 2D network out of its linear counterpart is complete.

Claims

1. An optical interconnect network comprising: at least one stage which has an optical interconnect stage (6) connecting a two-dimensional array of interconnect input ports to a two-dimensional array of interconnect output ports; an array of optical processing modules (8) each having a two-dimensional array of module input ports, optically coupled to a respective interconnect output port, and a two dimensional array of module output ports characterised in that each module (8) is functionally identical to a first and a second pair of two-input, two-output processing sub-modules in which each input of each of the second pair of processing sub-modules is connected to a respective output of the first pair of processing sub-modules, and wherein the inputs of each second processing sub-module are connected to the outputs from different first processing sub-modules.
2. A network as claimed in claim 1 in which each processing sub-module of the sub-modules to which each module is functionally identical has two states, one state in which the inputs are exchanged, and a second state in which they are not exchanged.
3. A network as claimed in claim 1 in which each processing sub-module of the sub-modules to which each module (8) is functionally identical performs weighted sum and difference calculations, the network being configured to perform a fast Fourier transform.
4. A network as claimed in either claim 1 or claim 2 in which the interconnect stage (6) performs a

The exchange/bypass decisions of the original network generate a sorted output of the network, but the output of the equivalent 2D network will have the channel with the n th destination address at the port which is the mapping of port n in the linear array. This scrambling of lines is wholly determined in advance and is functionally unimportant; if it is required to sort into a more visible order then the destination addresses can be modified at the start, or equivalently the modules can be set according to the reverse-mapped addresses.

A single exemplary two dimensional network has been described above. The starting point was a one dimensional network. Consideration will now be given to whether and how any two dimensional network according to the present invention can be constructed from a linear one, independent of the function of the modules and the interconnection pattern.

In analogy to the treatment of the perfect shuffle sorting network above, two consecutive stages can be concatenated into one stage of a 2D network by regarding the modules as being moved around in space. It may not be possible for one interconnection followed by a layer of 4x4 modules to do the work of two, one dimensional stages comprising an interconnection, a 2x2 module as an interconnection and a final 2x2 module sequence. Whatever the function of the modules in figure 2 we know (referring now to figure 7) that the intermediate module output X' is a function of X and Y .

$$X' = X'(X, Y)$$

and similarly for three of the other intermediate outputs

$$Y' = Y'(X, Y)$$

$$W' = W'(W, Z)$$

$$Z' = Z'(W, Z)$$

The final module outputs X'' and Z'' which are functions of X' and Z'

$$X'' = X''(X', Z')$$

$$Z'' = Z''(X', Z')$$

so,

$$X'' = X''(X, Y, Z, W)$$

$$Z'' = Z''(X, Y, Z, W)$$

In order that the functions X'' and Z'' be calculated by one two dimensional, four-input four-output processing module, it is necessary for the other two outputs, Y'' and W'' , to be functions of the same four inputs

$$Y'' = Y''(X, Y, Z, W)$$

$$W'' = W''(X, Y, Z, W)$$

which can only be achieved if intermediate results Y' and W' meet at a common module.

$$Y'' = Y''(Y', W')$$

$$W'' = W''(Y', W')$$

This imposes a restriction on the second of the two interconnects of figure n, which can be summarised as follows -

The complement of a line emerging from a module is defined to be the other line coming out of

it. For any two lines which are the two inputs to a module, the complements of these two lines must also meet at a common module.

If an interconnect does not satisfy this condition then it is not possible to concatenate the consecutive stages of Figure 7 into the functionally identical single stage of two dimensional 4x4 modules according to the present invention.

In principle, for a network satisfying the interconnect condition as stated above, it is always possible to pick up the lines in the linear arrangement of Figure 1 and move them into a two dimensional array as in Figure 3. There will be many ways to do this mapping, but what is sought is a 2D interconnection which can readily be achieved using a simple optical arrangement. An example is the special mapping applied to the perfect shuffle network. If a successful mapping scheme can be found then the task of generating a 2D network out of its linear counterpart is complete.

Claims

1. An optical interconnect network comprising:
at least one stage which has an optical interconnect stage (6) connecting a two-dimensional array of interconnect input ports to a two-dimensional array of interconnect output ports; an array of optical processing modules (8) each having a two-dimensional array of module input ports, optically coupled to a respective interconnect output port, and a two dimensional array of module output ports characterised in that each module (8) is functionally identical to a first and a second pair of two-input, two-output processing sub-modules in which each input of each of the second pair of processing sub-modules is connected to a respective output of the first pair of processing sub-modules, and wherein the inputs of each second processing sub-module are connected to the outputs from different first processing sub-modules.
2. A network as claimed in claim 1 in which each processing sub-module of the sub-modules to which each module is functionally identical has two states, one state in which the inputs are exchanged, and a second state in which they are not exchanged.
3. A network as claimed in claim 1 in which each processing sub-module of the sub-modules to which each module (8) is functionally identical performs weighted sum and difference calculations, the network being configured to perform a fast Fourier transform.
4. A network as claimed in either claim 1 or claim 2 in which the interconnect stage (6) performs a



perfect 2-D shuffle between the interconnect input ports and interconnect output ports.

5. A network as claimed in claim 1 in which each sub-module defines an AND or an OR operation, the network comprising a programmable logic array.

Patentansprüche

1. Optisches Vermaschungsnetzwerk mit: zumindest einer Stufe, die eine optische Vermaschungsstufe (6) hat, die ein zweidimensionales Feld von Vermaschungseingangsanschlüssen mit einem zweidimensionalen Feld von Vermaschungsausgangsanschlüssen verbindet; einem Feld von optischen Verarbeitungsmodule (8), von denen jedes ein zweidimensionales Feld von Moduleingangsanschlüssen hat, die optisch jeweils mit Vermaschungsausgangsanschlüssen verbunden sind, und einem zweidimensionalen Feld von Modulausgangsanschlüssen, dadurch gekennzeichnet, daß jedes Modul (8) funktional identisch ist mit einem ersten und einem zweiten Paar von Verarbeitungssubmodulen mit zwei Eingängen und zwei Ausgängen, wobei jeder Eingang eines jeden zweiten Pairs von Verarbeitungssubmodulen mit einem jeweiligen Ausgang des ersten Pairs von Verarbeitungssubmodulen verbunden ist, und wobei die Eingänge eines jeden zweiten Verarbeitungssubmoduls mit den Ausgängen aus verschiedenen ersten Verarbeitungssubmodulen verbunden sind.
2. Netzwerk nach Anspruch 1, bei dem jedes Verarbeitungssubmodul der Submodule, mit denen jedes Modul funktional identisch ist, zwei Zustände hat, wobei in einem Zustand die Eingänge vertauscht und im zweiten Zustand nicht vertauscht sind.
3. Netzwerk nach Anspruch 1, bei dem jedes Verarbeitungssubmodul der Submodule, mit denen jedes Modul (8) funktional identisch ist, eine Berechnung der gewichteten Summe und Differenz durchführt, wobei das Netzwerk zur Durchführung einer schnellen Fourier-Transformation konfiguriert ist.
4. Netzwerk nach Anspruch 1 oder 2, bei dem die Vermaschungsstufe (6) eine vollständige 2-D-Umordnung zwischen den Vermaschungseingangsanschlüssen und den Vermaschungsausgangsanschlüssen vornimmt.

5. Netzwerk nach Anspruch 1, bei dem jedes Submodul eine UND- oder eine ODER-Operation definiert, wobei das Netzwerk ein programmierbares Logikfeld enthält.

Revendications

1. Un réseau de commutation optique comportant: au moins un étage comprenant un étage de commutation optique (6) reliant une rangée bidimensionnelle de ports d'entrée de commutation à une rangée bidimensionnelle de ports de sortie de commutation; une rangée de modules de traitement optique (8), chacun d'eux comprenant une rangée bidimensionnelle de ports d'entrée de module couplés optiquement à un port de sortie de commutation respectif et une rangée bidimensionnelle de ports de sortie de module, caractérisé en ce que chaque module (8) est fonctionnellement identique à une première et à une seconde paire de sous-modules de traitement à deux entrées, deux sorties, dans lesquels chaque entrée de chacun des sous-modules de traitement de la seconde paire est reliée à une sortie respective de la première paire de sous-modules de traitement, et où les entrées de chaque second sous-module de traitement sont reliées aux sorties des premiers sous-modules de traitement différents.
2. Un réseau selon la revendication 1, dans lequel chaque sous-module de traitement des sous-modules avec lesquels chaque module est fonctionnellement identique comprend deux états, un état dans lequel les entrées sont commutées, et un second état dans lequel elles ne sont pas commutées.
3. Un réseau selon la revendication 1, dans lequel chaque sous-module de traitement des sous-modules avec lesquels chaque module (8) est fonctionnellement identique effectue des calculs de sommation pondérés et différentiels, la configuration du réseau étant agencée pour effectuer une transformée de Fourier rapide.
4. Un réseau selon l'une quelconque des revendications 1 ou 2, dans lequel l'étage de commutation (6) effectue un brassage bidimensionnel parfait entre les ports d'entrée de commutation et les ports de sortie de commutation.
5. Un réseau selon la revendication 1, dans lequel chaque sous-module définit une opération ET ou OU, le réseau comprenant une rangée logique programmable.

