Interconnection networks in computer architecture pdf

Go to the main content Go to the Content Table Reference EntryDOI: High-performance computing architectures use connection networks as a link between processor cores, memory modules, and sub-part devices. In their most common form, accession networks are a central component of all computing and communication systems, from internal relationships of embedded chip architectures to geographic scale systems such as broad-scale networks and the Internet. This section focuses on joining networks because they are used in multiprocessing and multi-core systems. In particular, the section discusses the basic concepts and principles of systems for interconnected processors, cores and memory modules. These concepts and principles are not exclusive to this domain application and have in fact also found application in the networks used in storage systems, area system networks, and local area networks, albeit with different engineering... This is a preview of the content of the subscription, log in to check access. The comments of Mitchell Rascinia and Jeffrey Young are grateful. Anjan KV, PinkstonTM (1995) Effective fully adaptive gridlock recovery scheme: DISHA. In: Proceedings of the 22nd International Symposium on Computer Architecture, IEEE Computer Society, Silver Spring, MD, page 201-210Google ScholarAkers SB, Krishnamurti B (1989) Group-theoretical model for symmetrical communication networks. IEEE Trans Comput C-38 (4):555-566CrossRefMathSciNetGoogle ScholarClos C (1953) Study of non-blocking switching networks. Bell Syst Tech J 32 (5):406-424Google ScholarDally WJ (1990) K-ary n-cube Network Performance Analysis. IEEE Trans Comput C-39 (6):775-785Google ScholarDally WJ, Towles B (2004) Principles and Practices of Interconnected Networks. Morgan Kaufman, San Francisco, CAGoogle ScholarDally WJ, Seitz CL (1987) Route messages without deadlock in multiprocessing joining networks. IEEE Trans Comput C-36 (5):547-553Google ScholarDuato J (1993) New theory of non-block adaptive routing in wormhole networks. IEEE Trans Parallel Distrib Syst 4 (12):1320-1331CrossRefGoogle ScholarDuato J (1995) Necessary and sufficient condition for non-block adaptive routing in wormhole networks. IEEE Trans Parallel Distrib Syst 6 (10):1055-1067CrossRefGoogle ScholarDuato J (1995) Theory of non-block adaptive multi-gas routing in wormhole networks. IEEE Trans Parallel Distrib Syst 6 (9):976-987CrossRefGoogle ScholarDuato J (1996) Necessary and sufficient condition for blockless routing in end-to-end and shop and forward networks. IEEE Trans Parallel Distrib Syst 7(8):841-854CrossRefGoogle ScholarDuato J, Yalamanchili S, Mi L (2003) Networks Morgan Kaufmann, San Francisco, CAGoogle ScholarGlass ScholarGlass Ni LM (1992) Turn model for adaptive routing. Proceedings of the 19th International Symposium on Computer Architecture, queensland, Australia, page 278-287Google ScholarPatterson D, Hennessey J (2004) Computer Architecture: A quantitative approach. Morgan Kaufmann, Palo Alto, CAGoogle ScholarLeiserson CE (1985) Fat-trees: Universal Networks for Hardware-Efficient Supercomputer. IEEE Trans Comput C-34:892-901Google ScholarPreparata FP, Vuillemin J (1981) Cubic cycles: a universal network of parallel computing. Commun ACM 24(5):300-309CrossRefMathsciNetGoogle ScholarSamatham MR, Pradhan DK (1989) De Bruijn Multiprocessor network: a universal parallel processing and sorting network for VLSI. IEEE Trans Comput C-38 (4):567-581Google ScholarSchroeder MD, Birrell AD, Burrows M, Murray H, Needham RM, Rodeheffer TL (1991) Autonet: high-speed, self-adjusting local network using point links. IEEE J Select area Commun 9 (8):1318-1335CrossRefGoogle ScholarSong YH, PinkstonTM (2003) Progressive approach to message processing depends on gridlock in parallel computer systems. IEEE Trans Parallel Distrib Syst 14(3):259-275CrossGoogle ScholarWu CL, F TYeng (1980) About the class of multi-stage joining networks. IEEE Trans Comput C-29:694-702Google ScholarValiant LG (1982) Rapid Parallel Communication Scheme. SIAM J Comput 11:350-361CrossRefzbMATHMathSciNetGoogle Scholar© Springer Science-Business Media, LLC 20111.The School of Electrical and Computer EngineeringGeorgia Institute of TechnologyAtlantaUSA In order to continue to use our site, we ask you to confirm your identity as a person. Thank you so much for your cooperation. Jose Duato, ... Lionel Ni, at Interconnection Networks, the 2003Multistage Network Of Accession (MINs) connect input devices to display devices through a number of switch stages where each switch is the crossbar of the network. The number of steps and the connection patterns between the stages determine how networks can be routed. MIN were originally proposed for phone networks and then for processor arrays. In these cases, the central controller sets the path from input to exit. Where the number of inputs is equal to the number of outputs, each input simultaneously sends the message to one exit, and each output receives a message from exactly one input. These single-tissue patterns can be presented as a permutation of entry addresses. For this application, MINs have been popular as alignment networks for storage and access to arrays parallel to memory banks. Array storage is usually distorted to allow conflict-free access, and the network is used to decrypt arrays during access. These networks can also be configured with the number of inputs, more than the number of outlets (hubs) and vice versa (extenders). On on the other hand, in asynchronous multiprocessors, centralized management routing and permutations are not feasible. In this case, the routing algorithm is necessary to create a path at the MIN stages. Depending on the accession scheme used between the two adjacent stages and the number of stages, different MINs were proposed. MINs are good for building parallel computers with hundreds of processors and have been used in some commercial machines. William J. Dalley, of Interconnection Networks, 2003Interconnection Networks have a long history. Scheme networks have long been used in telephony. In the 1950s, there were many proposals to combine computers and cellular machines, but several prototypes. They came in the 1960s. Solomon in 1962 was the first of a long line of grid-related, bit-serial multicomputers. STARAN with its flip network, C.mmp with the crossbar, and Illiac-IV with a wider 2-D grid followed in the early 1970s. During this period, there were also several indirect networks used in vector and massive processors to connect multiple processors to multiple memory banks. To address this problem, the academic community has developed several options for multi-stage accession networks. BBN Butterfly in 1982 was one of the first multiprocessors to use an indirect network. The n-cube or hypercube binary network was introduced in 1978 and was first implemented at Caltech Cosmic Cube in 1981. In the early 1980s, the academic community focused on the mathematical properties of these networks and increasingly separated from the practical problems associated with connecting real systems. Dan C. Marinescu at Cloud Computing(Second Edition), 2018Intercoding Networks for Multiprocessor Systems, Supercomputers, and Cloud Computing is discussed in the following sections. Computing and communication are deeply intertwined, as we've seen in Chapters 3 and 4, and connectivity networks are critical to the performance of computer clouds and supercomputers. A number of concepts important for understanding interconnected networks will be introduced in the following. The network consists of nodes and links or communication channels. The degree of a node is the number of links a node is connected to. The nodes of the join network can be processors, memory units, or servers. The site's network interface is the hardware that connects it to the network. Switches and communication channels are elements of the compound tissue. Switches receive data packets, look inside each package to determine the IP address of the destination, and then use the routing tables to rewind the package to the next move to the final destination. The n-way switch has n ports that can be connected to links to link n. The network of accession may not be blocking if it is possible to connect any permutation of sources and directions to any The accession network is blocked if the requirement has not been met. While CPU and memory technology followed Moore's law, the accession networks evolved at a slower pace and became an important factor in determining the overall performance and cost of the system. For example, from 1997 to 2010, the speed of the ubiquitous Ethernet network increased from 1 to 100 Gbps. This increase is a little slower than Moore's Law for Traffic (354), which predicted 1 Tbps Ethernet by 2013.Interconnection networks are distinguished by their topology, routing and flow control. Network topology is determined by how nodes are interconnected, routing decides how a message gets from source to destination, and thread management determines how buffer space is allocated. There are two main types of network topology: Static networks where there are direct connections between servers; The networks where switches are used to connect servers are switched. The topology of the network of accession determines the diameter of the network, the average distance between all pairs of nodes, the width of the bisection, the minimum number of links cut into two halves of the network section, the bandwidth, and the cost and energy consumption. When a network is divided into two networks of the same size, bisection bandwidth measures the bandwidth between the two. The full bandwidth of the bi-enabled allows one half of the network nodes to communicate simultaneously with the other half of the nodes. Let's say that half of the nodes inject data into the network at a B Mbps speed. Switching fabric should have sufficient dual-direction bandwidth for cloud computing. Some of the most popular topologies with static connectivity: Bus, simple and economical network, see Figure 5.8A. It doesn't scale, but it's easy to implement cache consistency by spying on distributed memory systems. The bus is often used in multiprocessing systems of general memory. Figure 5.8. Static connection networks. (A) Bus; (B) Hypercube; (C) 2D grid: (D) 2D-torus. The hypercube has good bandwidth, the number of nodes is N-2n, the n'logN degree, and the average distance between O/N nodes. Example: SGI Origin 2000.2-D grid, see Figure 5.8C. N×n 2D grid has many ways to connect nodes, has an O (n) cost, and an average O/n latency. Thus, the grid is not symmetrical at the edges, its performance is sensitive to placing communication nodes at the edges compared to the middle. An example: the 1990s Intel Paragon supercomputer. Thor is good for apps that use a connection with your nearest neighbor, see Figure It is common for proprietary relationships. Example: 6-D grid/torus of fujitsu K. Switched networks have multiple levels of switches connecting nodes, as shown in figure 5.9:Figure 5.9. Networks have been switched. (A) 8 × 8 crossbar switch. The 16 nodes are connected by 49 switches represented by dark circles; (B) 8 × 8 Omega Switch. The 16 nodes are connected by 12 switches represented by white rectangles. Cost of O (NlogN) and Delay O (journalN). Omega networks are low-diameter networks. Cloud-based joining networks. Cloud infrastructure consists of one or more storage scale computers (WSCs) discussed in Section 8.2. WSC has a hierarchical organization with a large number of servers, interconnected high-speed networks. Cloud network infrastructure must meet several requirements, including scalability, cost, and performance. The network should provide low latency, high communication speed and, at the same time, provide transparent communication between servers. In other words, each server should be able to communicate with each other server at a similar speed and delay. This requirement ensures that applications do not have to be location-aware and, at the same time, reduce the complexity of managing the system. As a rule, network infrastructure is organized hierarchically. WSC servers are packed in racks and connected by a rack router. The rack routers are then connected to cluster routers, which in turn are interconnected by local communication. Finally, intercenter networks connect multiple WSCs. Cost considerations ultimately decide the actual organization and performance of the communication tissue. Resubscription is a particularly useful indicator of how the network is fit to join a large-scale cluster. Ressubscription is defined as the relationship of the worst achievable aggregate bandwidth between servers to the total bandwidth of the connection. A one-to-one resubscription indicates that any host can communicate with arbitrary hosts throughout the connection bandwidth. A 4 to 1 resubscription value means that only 25% of the bandwidth available to servers can be achieved for some communication models. Typical resubscription rates are in the 2.5 to 1 and 8 to 1 range. The cost of routers and the number of cables associated with routers are the main components of the total cost of the network of accession. The density of the wires increased at a slower rate than the processor speed, and the wire delay remained constant over time, so performance and lower costs can only be achieved through an innovative router architecture. This motivates you to take a closer look at the actual design of the routers. Thick trees. A special copy of Clos's topology, discussed in section 5.7, fat trees, are optimal relationships for large-scale clusters and, as an extension, for WSCs. Thick trees have additional links to increase bandwidth near the root of the tree. Some set of paths in a thick tree will saturate all the bandwidth available for end hosts for arbitrary communication models. The oily wood architecture can be built with cheap product parts, as all the commuting elements of the thick tree are identical. Figure 5.10 shows 192 nodes built of two 96 switches and twelve 24 switched switches. Two 96-way breakers are fundamentally connected through 48 links. Each 24-way switch has a 6×8 connection to the root and a 6×16 down connection with 16 servers. Another 192 knot of fat tree is associated with two 96-way and twelve 24-way switches shown in figure 5.11.Figure 5.10. Thick trees. (above) Fat in nature. No, no, no. The connection network is 192 nodes with two 96th and twelve 24-way switches in a computer cloud. Figure 5.11. The 192 knot of fat tree is connected to two 96-way and twelve 24-way switches. Table 5.2 from No.228 summarizes the performance/cost of a 2D mesh, 2D torus, hypercubes of about 7, fat wood and a fully connected network. Two measurements of the network's connection performance, the bandwidth and the average and maximum number of transitions are shown, as well as three elements that affect the cost, the number of ports per switch, the number of switches, and the total number of links. The thick tree has a large bandwidth with a small number of I/O ports per switch, while a fully connected connection has an extraordinarily large number of links. Table 5.2. A one-by-one comparison of the performance and costs of several topologies of the network of accession for 64 nodes. Property2D mesh2D torusHypercubeFat-treeFully connectedBW in no links81632321024Max/Avg hop count14/78/46/311/91/1I /O ports on the switch557464Number switches64646419264Total number of links1761922563842080Farhad Mehdipur, ... Bahman Javadi in Advances in Computers, 2016 The Data Center Accession Network is the core for supporting big data and is the most important infrastructure that requires urgent attention. Big data applications transmit large-scale data to processing and analysis centers So we need to look at two types of relationships: data center connectivity and intra-sponsored communication centers. Communication between data centers: This is a connection from the outside world to a data center, data, usually based on the existing network infrastructure (i.e. the Internet). Modern physical infrastructure for most countries is a high bandwidth of fiber optic relationships that can handle the rapid growth of data size. The common architecture in these compounds is the IP-based wavelength multiplexing (WDM), which is based on the multiplexion of multiple optical media at different wavelengths. In this technology, each optical fiber can carry multiple signals with different wavelengths and therefore increase the bandwidth of the network. While most WDM-based networks are deployed with a bandwidth of 40 Gbps, a new 100Gbps standard has recently been introduced to meet the demand for high-speed inter-datacenter connectivity. Because WDM technology is limited by the bandwidth of an electronic bottleneck, a new technology called orthogonal frequency-mapping (OFDM) is being introduced. OFDM is a multi-carrier parallel transmission technology that divides high-speed data flow into low-stage sub-data streams to transmit it through several orthogonal sub-careers. This makes OFDM more flexible and efficient than WDM technology and more promising for inter-datacenter connections, which deals with big data transmission. (ii) Connection within the data center: Connecting an intra-data center is used to transmit data in a data center. As mentioned earlier, most current data centers consist of high-performance servers associated with product switches in the multi-level topology of fat trees. While the servers in each rack are connected through 1 Gbps, all rack servers are interconnected through a set of 10Gbps switches. Although this architecture is scalable and error-prone, energy consumption is a major problem. Another problem in this architecture is the high delay due to multiple stores and forwards in the switch hierarchy. For big data applications, we need network technology to connect intracenter data centers that can provide high bandwidth, low latency, and low energy consumption. To address these problems, optical networks have recently received a great deal of attention as an alternative to joining data centers to meet the needs of big data analytics. Currently, optical networks are only used for point communications in data centers, where links are based on low-cost multimodal fibers for short communication. The desired future for interconcation of intra-data centers there will be all optical connections with switching in the optical domain, as shown in the pic. 5. Electric and optical (E/O) and optical-electric (O/E) transceivers, which are the main source of the source of the source of the energy consumption in the current architecture. Thus, all-op connection networks can provide bandwidth on a low-energy throp (Tbps) scale for future DCNs.Figure 5. Point to point versus all-optical relationships. Dimitrios Serpanos, Tilman Wolff, in the architecture of network systems, 2011Multiprocessor network accession are attractive candidates for networking on chip designs because of their characteristics in terms of performance and reliability. Multiprocessor connections are generally seen as a class of networks with closely related end-of-systems that are colocated, with short latency, high bandwidth, and extremely low bit error rates. These characteristics differ from those of long-distance networks, which emphasize high bandwidth rather than low latency, and which show higher bit error rates. In addition, NOC finds power and space constraints that go beyond the constraints and requirements of multiprocessing relationships. Multiprocessor accession networks are generally classified into two categories: point and switch networks. Figure 13-2 shows examples of both categories of networks. In point networks, the processing elements are connected directly to each other. In switch-based networks, processing elements connect to switches that are organized in some topology. Thus, in switch-based networks, processing elements communicate with each other through a connection that includes one or more switches. In point networks, processing elements communicate with each other either directly (if they are connected directly) or through a path that includes one or more processing elements. In any case, each processing site requires an I/O component that allows you to communicate with a connected processing element or a connected switch. Figure 13-2. Networks are characterized by a number of structural parameters, such as the degree of the node (the number of links to the node), the diameter (the maximum length of the shortest length paths, interconnected between any pair of nodes), etc. Each parameter affects one or more physical parameters of the NoC chain, i.e. bandwidth, latency, power and area. Network topology can be used for both point and switching networks. A wide range of literature explores the implementation of various networks such as NoC, such as mesh nets and torsos, octagon and thick wood. Grid, torus and octagon networks are examples of point nets, and the fat tree network is an example of a switch-based network because it consists of switches as internal nodes connecting the processing elements on the leaves. Despite the similarity of multiprocessing networks NoC, there are also important differences due to technological characteristics. These differences are related to wiring and buffering, which affect performance characteristics such as bandwidth, latency, and energy consumption. In typical multiprocessor networks with a chip-to-chip connection and from board to board, the width of the data path has always been limited to a small number of wires, while the density of wires in NoC is about hundreds, which greatly increases the bandwidth. In multiprocessor networks and LAN and WAN networks, low wire density leads to widespread use of buffers where data is temporarily stored in the event of strife or overload. In NoC, the widespread use of buffers is prohibitive because buffers introduce latency and memory on the chip is costly in terms of power and silicon area. Cheng-Ta Chan, ... Mi Lu, in advances in parallel computing, 1998Working Network Accession has long been seen as one of the key issues in improving system performance for multiprocessor systems and telephone switching systems. Among the popular class of multi-stage joining networks (MIN) are Gamma Network, REGIN, CGIN and our recently proposed Palindrome (PIN) connection network consisting of 3 x 3 switch elements. The PIN is an unsustainable network. Completely disparate paths exist in the network between any source/couple of destinations. A failure of a specific internal switch element will not disable the connection between the source and the destination. Redirecting can be done effectively when locked (due to congestion or failures). Compared to its counterparts, the PIN achieves error tolerance without increasing the complexity of the hardware. With the routing tag in the PIN, you can calculate an alternative routing tag by changing only two bits of the routing tag, regardless of network size. Thus, the redirection logic is extremely effective and easy to implement, which is conveniently included in the composite switching element. Analyzing the performance of redundant path networks, such as a PIN, can be quite complex because of the inconsistency of paths between the source/destination pair. However, you can evaluate performance using a simplified routing scheme that doesn't take full advantage of the network connection. This way, you can find a lower network performance boundary. Conflict resolution strategies for a simplified routing algorithm will fall or withhold service requests depending on where the PIN is at. The PIN architecture and routing and redirection schemes are introduced in Section 2. To analyze the performance of the PIN, this document in section 3 offers a simplified routing algorithm (SR). Performance analysis with the proposed SR algorithm Section 4. Conclusions are given in the section BASU, in soft computing and intelligent systems, 2000 The simplest network of accession for parallel processing is a linear array. Here we have N processors about numbers P(1), P(2),..., P (N), each P processor (i) linked by communication paths with P processors (i No. 1) and P (i No. 1), where 1 qlt; i lt; N, without other references available, and P(1) and P/N processors each has only one nearest neighbor. The P(1) processor is connected to P(2), and the P (N) processor is associated with P(N No. 1). Obviously, the number of O (N) and diameter links is also O(N). In general, input and outputs are limited to the end processors P/1 and P (N) because doing I/O through all N processors in parallel is not technologically feasible for large N values, as I/O pads occupy large areas on VLSI chips. Linking the connections between the two neighboring PEs on a linear array results in two smaller linear arrays and therefore it has a linear connection of 1.William J. Dally, in Interconnection Networks, 2003Point-to-point network joining replaced buses in an ever-expanding range of applications, which includes connection to the chip, switches and routers, and VIO systems. The new Infiniband and RapidIO standards are based on the concept of using a network to connect one or more processors with an array of devices in the I/O area. Companies such as Sonics build on a network connection chip, sometimes referred to as micronetworks. The use of join networks as router or switch tissue was first used by Avici, Pluris and others. Pramod Pandya, in the Computer and Information Security Handbook (Third Edition), 20131.Discover network connectivity and configuration, and look for network vulnerabilities:A.DoSB.SniffingC.SYN floodingD.ReconnaissanceE.All of the above2. Removing and/or altering data, installing backdoors and hiding traces of attacks are known as A.EnumerationB.ScanningC.DoSD.Operational attacksE.All of the above3. Port scanning techniques are used to detect open ports.A.TCPB.NetBIOSC.PDPD.HTTPE.All of the above4. The three-headed TCP handshake is installed during which of the TCP scan sessions? A.TCP Connect ()B.TCP SYNC. TCP FIND. TCP OpenE.All of the above5. TCP SYN scanning is also known as A.full openB.half openC.full closeD.half closeE.All of the above interconnection networks in computer architecture slideshare. interconnection networks in computer architecture pdf. various interconnection networks in computer architecture