CHAPTER 10.

DISCUSSION AND CONCLUSIONS

10.1. General

The concept of a synchronous optical hypergraph has potential for many applications, of which the metropolitan area network is only one. The hypergraph architecture may be extended in size by combining satellite networks, or it may be implemented, in a very small area, as a machine for processing data rather than as a system for transferring data. Another possible use of the synchronous hypergraph is in real–time control of industrial automation. In general, almost any real–time system which is spread over a large area can be realized by the synchronous hypergraph.

The two most important attributes of this system are

(1) The ability to transfer efficiently small and large messages at high bandwidth by using the conservative code. Thus, it is possible to exchange control messages efficiently, which enables the integration of various network functions in a uniform manner.

(2) Efficient global synchronization and total event ordering, enables the straightforward integration of voice and simplifies the design of distributed parallel algorithms.

The core of this system is a centralized, passive, optical star coupler, which is very reliable, and which splits the energy among the receiving nodes in an optimal way: if there are n nodes, each node receives $\frac{1}{n}$ of the energy. The links connecting the nodes to the centralized switch are single-mode, very low-loss, optical fibers. Thus, longer links do not impose any significant reduction to the energy received by each node, and it is better to let the signal travel a longer distance and to split equally the signal energy, rather than have shorter links but without an equal division of the energy, as in the cases of linear bus and passive ring.

10.2. Thesis Discussion

The following discussion shows that the thesis of this dissertation "*A low-dimension, high width (the number of ports on a net), globally-synchronized, optical hypergraph, can be efficiently (with low overhead) managed and controlled in a distributed manner,*" is an immediate result of the design and analysis of the system, as presented in the previous chapters. This design is not arbitrary, but is rather a logical outcome of the technological advances in optical communication and semiconductors. Furthermore, the design gives specific solutions to problems in the areas of communication and computation.

The argument is illustrated in Figure 10.1. The different attributes of the optical hypergraph are circularly related to one another. In the discussion only the major relevant arguments are mentioned.

The immediate consequence of recent technological advances is that communication bandwidths about 100 times greater than at present are feasible. Therefore, a medium with such high bandwidth can support multiple nodes. Note that as the bandwidth increases the width (number of nodes) of the net can be increased as well. The high

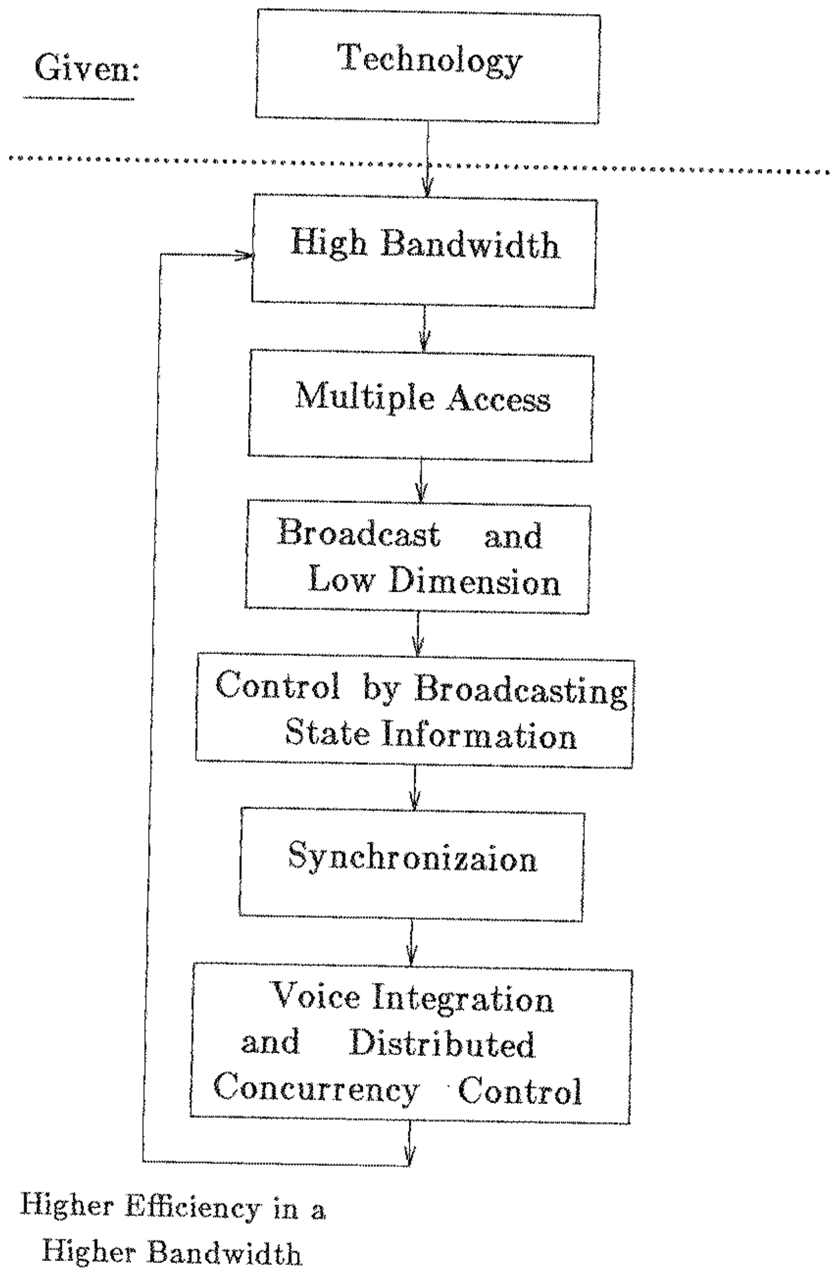

Figure 10.1: Thesis Discussion

bandwidth is well suited for the realization of the conservative code, which enables nodes to efficiently transmit and receive very short messages for the control of the multiple-access medium.

For a given total number of nodes, p, the dimensionality, D, of the network is related to its width, W, by

$$p = W^D \quad or \quad D = \log_W p$$

The broadcast nature of the transmission, together with a low dimensionality hypergraph, enable state information to propagate rapidly through the network, which is essential for the periodic exchange of state information and the implementation of global synchronization, so as to maintain total event ordering. Using this basic construction, the efficient integration of voice becomes very simple. The efficiency of voice integration increases with bandwidth. With this argument, the loop is closed. The period of exchanging state information gets shorter as the bandwidth gets higher. Thus, distributed algorithms are more efficient, since they use more recent state information.

10.3. On the Virtue of an Optical Hypergraph

The design principles of this system are applicable to different topologies; the design of 2D–R and 2D–P can be easily extended to an arbitrary hypergraph. This is important because of

(1) Graceful degradation – in principle, any hypergraph can degrade its operation to another arbitrary hypergraph, which is smaller in size. If a node or a port ceases to exist, only its neighbors on its connected nets need be aware of this and make sure not to route any message to it. They may use an alternative route or send a

message to the originating node.

(2) Graceful upgrading – the opposite of the previous property. In many practical cases the system is not implemented at once, either because the requirements are not known or due to lack of resources. The optical hypergraph can be extended gradually, the basic configuration being a single net, and more nets can then be added later as needed.

10.4. Synchronous and Asynchronous Computations

The system is synchronous in its macrolevel (slot duration) and is asynchronous in its microlevel. The computations that are distributedly performed on this system can be viewed as asynchronous computations with bounded time, where the bound is the time slot or frame. It is interesting to note that if the slot size is determined by a typical page or file in the system (e.g., 4K bytes), then the bound on asynchronous computation is inversely proportional to the baud rate. This is another reason why increasingly higher bandwidth is significant.

For example, with a slot length of 8k bytes and a baud rate of 1 gigabit/second, the slot duration is 64 microseconds, which is shorter than the time of most software processes, even for fast, general purpose machines. Thus, if the basic computational step is not smaller than the bound asynchronous computation, then it is reasonable to consider that the distributed execution is synchronous.

10.5. Reduction of Software Complexity

One of the major design objectives is to reduce the system's software complexity. In the optical hypergraph design this objective has been achieved by several means:

(1) Open mode concurrency control, which operates without an explicit acknowledgement. This alleviates the necessity for special acknowledgement messages. Moreover, the open–mode concurrency control requests are performed by the interface, independent of the distributed operating system software.

(2) Total event ordering, simplifies the execution of distributed algorithms and the scheduling mechanism of distributed operating systems. With total event ordering it is also easier to maintain the bookkeeping of events in the system. The temporal order is the natural order in any system. Thus, with total event ordering it is simpler to ensure the serializability of a distributed procedure.

(3) Implementing "traditional" operating system functions (like routing or buffer management) by the interface in real time, by a dedicated hardware or firmware.

10.6. Further Considerations

The principles for constructing the synchronous hypergraph can be extended in two ways:

(1) Higher dimensionality, such as 3D regular hypergraph, with three ports per node. The network then can have 100,000 to 1,000,000 nodes. Increasing the dimensionality is done only when it is not possible to accommodate all the nodes on a lower–dimension hypergraph.

(2) Satellite network, another interesting extension, can be done by combining several optical hypergraphs via satellite networks. The satellite network has an **active star topology**, so that synchronization among its nodes should be as simple as for an optical star. Thus, coherent integration with the optical nets is possible but not simple. The major difference is that the satellite network exhibits much longer delays (about 1000 times).