Justifying Distributed Computing

Brian Goff, M.Eng., MBA

Introduction

Justifying an information system upgrade is as difficult today as it ever was. Once you get past the intangibles such as increased flexibility, improved service to customers and improved systems control, it gets very difficult to get a handle on the financial savings that everyone expects these systems to provide. However, it is possible to approach the cost/benefit analysis of a distributed system even if you don't have a crystal ball. In this article we'll review a methodical technique pioneered by **GOFF Tech Consulting** that can help you justify your next system upgrade.

No one will deny that the development of smaller and more powerful processors, as well as the Internet itself, is fueling the drive toward distributed computing. Ever since the mini/microcomputer became a serious contender in the business arena, new computers have been appearing on the market that perpetuate the trend toward using less expensive CPU's for the same tasks. Add to this the increased capability of communications and networking, derived primarily from the Internet, and it becomes easy to see why managers are clamoring for budgets for new and ambitious undertakings. It's ironic that, historically, as hardware costs go down, IT budgets continue to rise (recent macro-economic difficulties notwithstanding). One reason why is that more and more can now be done. Consequently, the issue of cost justification hasn't gone away, rather it has become ever more important.

There are many ways in which to define "distributed" systems. We want a definition that allows us to include everything from Internet applications through PC front-ends to LAN computing. By summarizing the essential requisites of distributed computing we get the following definition:

1. Computer processing at more than one site, and application of that ability toward user-oriented tasks
2. Interconnection of these locations, hence some data transmission facility
3. Some element of data handling and/or data storage at each site
4. Conformity to common standards for operation

There are many motivations for examining distributed computing. For example, many organizations look to distributed systems to host mission-critical applications, others to provide better customer service. But while every organization's motivations are different, one of the common major motivations is the possibility of providing solutions at a lower cost than comparable centralized systems. The cost of a new system or a system upgrade is always important and it must be seen to be reasonable in relation to the system's benefits.

Let's look more closely at the major cost elements in distributed systems. These include:

1. The system components - central and local hardware and software, and communications
2. System development - design, coding and implementation
3. System operation

The total cost of the system and the proportions attributable to each of the system's components will depend on the degree of distribution in the system. Increasing the degree of distribution will increase certain costs and decrease others. It is necessary to consider these trade-offs in order to decide the optimum solution. Understanding the degree of distribution helps us develop a general, practical model for the discussion of costs.

The degree of distribution is a way of comparing local and non-local activities. Since computer systems operate in different application areas, it is useful to calculate the degree of distribution in a way appropriate to the particular situation. In general, the degree of distribution compares the functions performed locally with the functions of the total system as follows:

DOD = L / T where DOD = degree of distribution

L = a measure of functions performed locally

T = a measure of functions performed by the total system

There are a number of different measures that can be used for measuring system functionality (see inset next page). Clearly in choosing any one of these criteria there will be a situation where the calculated degree of distribution for a given system will appear unrealistic. In order to minimize this effect it is important to select the method that is appropriate for the type of work performed. In back office work, file storage and data access are predominant factors, whereas processing may be the dominant factor in systems that run management sciences applications.

We have already identified the major cost elements in a distributed system so now let's look at the effects of distribution on these elements. In terms of system components, the least distributed case (DOD = 0) is represented by the simple terminal, or VDU. There is no cost for special equipment, environment, operators, programmers, or for the development of local software. As the degree of distribution increases, we must provide for the cost of local software to utilize the processing power of what must be a smarter (and more expensive) terminal. Hardware costs increase as more local processing capability is provided.

For the most distributed case (DOD = 1), we have local processing with no data transmission or central processing costs. However, depending on the volume of work to be performed, the application may have outgrown a PC, a workstation, or even a minicomputer. The user may now be involved in the world of environment-controlled computer rooms, specialized operators and systems programmers, and possibly specialized applications developers.

<u>Input Comparison</u>. Measuring the number of data characters entered and used locally, by taking the total number of characters entered locally and subtracting the number transferred to the central computer, and comparing it to the number entered in the whole system.

<u>Output Comparison</u>. Subtract the number of characters sent from the central computer from the total number of characters printed or displayed

locally. Compare this with the number of characters printed or displayed by the total system.

Processing Comparisons. Compare the number of instructions executed locally to the total number of instructions executed.

File Volume Comparisons. On this basis comparison is made between the volume of data held for local use with the volume of data held by the total system. Exclude any redundant data and any data that exists solely for backup or security reasons. One disadvantage of this method is that in cases where files are created centrally, but stored locally the degree of distribution will be over-emphasized.

Data Accesses. This comparison is made on the basis of the volume of data that is accessed locally to the volume of data that is accessed in the total system. Again, redundant data should be excluded as should data such as index files or non-accessed fields in a database table.

Degree of Distribution Measures

The effects of distribution on transmission facilities are a little more complicated. In the least distributed case (DOD = 0), all data is transmitted to the central installation. The number of terminals supported, their type, the kinds of lines (leased or dialed), and the peak volumes are all inputs in the determination of the speed of the line and the modems. As DOD increases, the costs of communications decreases as a step function as leased lines are used, and as a succession of straight line slopes as usage charges come into play. When communication distances are short, the largest proportion of communication charges will come from the modems, controllers, or other hardware used for data transfer. At the extreme DOD of 1, the cost of communications drops to zero.

Charges for central processing are at a maximum for the least distributed case (DOD = 0). These costs include on-line connection charges, use of core storage, channels, disk or tape drives, and on-line and off-line storage. As DOD increases, central processing costs decrease as determined by the application, and when DOD = 1, they could reach 0 (even at DOD = 1, some network control is required and thus, the costs may never reach zero).

When the major cost elements are brought together, the resulting total cost curves looks like

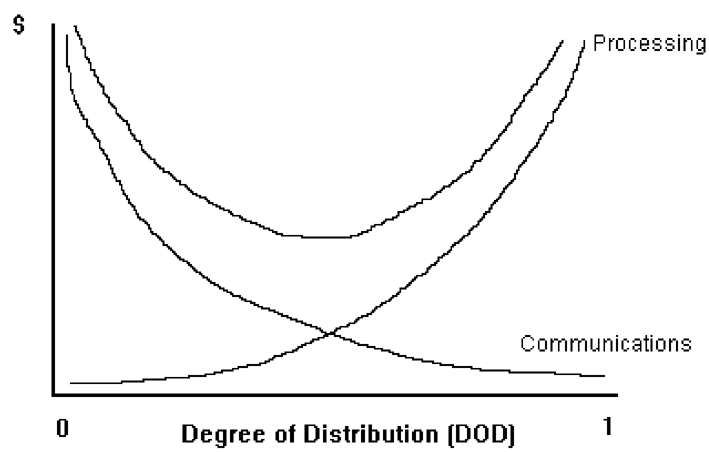

What we see here is the not so obvious! The notion that the economies of scale for computing power offered by large centralized systems offsets data communications costs is not accurate. Even though this figure is only a general representation and the cost curve of your particular configuration may be significantly different the upper curve suggests that under certain circumstances a small change from extreme centralization to extreme decentralization will increase costs. Conversely, a significant change can be required to achieve a more cost effective solution. *All* the curves indicate that a partially distributed one is more cost effective.

The idea of economies of scale is embodied in Grosch's Law which states that the power of a computer is proportional to the square of its costs. If you pay twice as much, you get four times the power. This law is expressed as

$$\text{Cost} = K \; \nabla \; \text{Power}$$

where K is a constant. This implies that computing power will be more expensive if a computer of power 4 (cost of 2K) is replaced by four computers of power 1 (at a cost of 4K). While this is an approximation, there is evidence that the rule holds when comparing computers in a particular range. It is unclear how this relationship holds when comparing computers in different ranges. Specifically, benchmarks show that a powerful minicomputer has a "power" equivalent to a mainframe costing at least ten times as much. And now, PC's are performing at rates equivalent to existing workstations. The rationalization of this is that, at least in the micro/mini/workstation ranges, performance is rising at somewhat comparable rates. Also, Grosch' s law deals with performance as measured in terms of processor power. There are other references that are used for measuring economies of scale for other computer facilities.

Let's consider a business example. In this case, we'll examine a small banking office that deals in foreign exchange and large loan and deposit accounts. The computing requirements in this example are driven by these business requirements (see Figure 2)

Foreign Exchange

The bank handles high levels of forward deals which demand access to up-to-date information on exchange rates and existing agreements. Decisions are made as to whether or not to accept the deal and at what exchange rate. Details of arrangements are stored for monitoring and payment.

Loan and Deposit Accounts

Loan and deposit accounts are offered to customers. Payments and withdrawals are recorded and the system computes interest and charges. A link to the foreign exchange facilities is required so that deals that mature can be paid into or drawn from a specified loan or deposit account.

Fig 2. Computing Requirements (Example)

These requirements imply that there is a mixture of heavy on-line work (inquiries to decide whether to proceed with a deal) and a significant amount of batch work (interest calculations and complex statistical analyses of future deals). The stand-alone workstation solution was rejected in this case because the bank did not want to employ non-banking staff and they did not want to run such a machine during non-banking hours (i.e. full batch mode). A mainframe solution was not appropriate because of the requirement for high availability during banking hours. The solution that was chosen was therefore formulated on the type of processing required for on-line work on the workstation, and the batch work on the mainframe.

Degree of Distribution			
Local Processing Costs Foreign Exchange Loan Accounts Deposit Accounts Ledger Management			

	Information			
	TOTAL LOCAL COSTS			
	TOTAL COMM. LINE COSTS			
	Mainframe Processing Costs Foreign Exchange Loan Accounts Deposit Accounts			

Ledger Management Information			
TOTAL MAINFRAME COSTS			
TOTAL COSTS			

In this case and in others that have been examined, the partially distributed solution proved to be cheaper than either the least or the most distributed solution. Client-server systems present clear examples of this concept. In client-server systems, functionality is apportioned in a partially distributed configuration. Graphic information processing is performed locally, and data maintenance is centralized on the host (or hosts).

With this understanding, project managers will be able to bring an improved financial basis to their systems justifications. These methods are best applied during the planning phase. Methodologies typically go only so far as to describe the requirement for financial planning of systems. A sound understanding of the cost dynamics of distributed computing can augment even those methodologies supported by sophisticated CASE tools.