LDP RED BOOK

COMPANY CONFIDENTIAL

MARCH 1977

CONTACT:

OOD SECTION - BILL MCBRIDE

LDP DEVELOPMENT SECTION - BILL AVERY

LDP RED BOOK INDEX

I. INTRODUCTION

This document is the first edition of the LDP Red Book. It is a compilation of inputs from both the marketing groups and the development groups within LDP. As such, it reflects the official LDP position on product development strategies at the date of issue.

The LDP Red Book is divided into two sections. The first section represents inputs to the OOD product strategy and a response to their Red Book. It is meant to be more marketing oriented than technically oriented. It represents the needs of the LDP marketplace for products which are currently under development or under consideration for development within OOD. Because market needs change, these needs also change over time and while these inputs are as accurate as we can make them at the current time they will be modified as time passes. The second section represents the LDP development plan and strategy. It is also based on inputs from the LDP marketing groups and represents a product strategy intended to create the LDP systems required over the next two to five-year period.

The LDP Red Book is meant to be used in conjunction with the LDP five-year plan. The Red Book describes product strategies which will be required to meet the financial goals and growth goals shown in the five-year plan and its implementation is required before we can achieve those goals.

The OOD portion of the LDP Red Book is divided such that it gives detailed inputs to each of the major development areas within OOD. The sum of these inputs represents the total marketing needs of the LDP organization from OOD developments. However, we have several major overall concerns which are listed below.

1. There is a lack of a coherent, overall low-end strategy for our low-end systems. The Krypton system, as it is currently defined, does not meet the very low-end needs of the LDP marketplace. In order to meet this we will need lower cost systems with cheaper, more adaptable packaging techniques. In addition, we will need low-cost peripherals which are attached to the system via some very low cost interface bus. In conjunction with this bus, we will need to be able to do on-board integration of the CPU and limited numbers of specific peripherals. For example, in the future it must be possible for us to integrate both the CPU chips and A to D conversion chips on the same board, or on an IEEE Bus.

2. We must have a coordinated strategy for replacement and coexistence of the VAX and 11 families. We must plan our replacement timing for VAX machines and 11's as well as the product positioning of each member of the VAX family versus the members of the 11 family. We must quickly determine how far down in the VAX family we will extend the 11 emulation mode. We must also commit to and determine the specifications of a hot-box VAX product.

3. We need a higher performance set of DECNET products which includes support for a high-speed multidrop serial bus. These tools are needed for our development of distributed lab systems. An essential part of this program is the need for a high-speed serial link between processors and between nodes in lab networks.

4. We need competitively priced disks. The RK05 and RK06 are too expensive and cause us to have non-

competitive systems. We are also concerned that the RL01 may be too small by the time it finally comes to market to be competitive with the disk offerings from DG and HP.

5. We have a long-term need for data and program transportability between our various operating systems and languages. We need a compatible single-user and multi-user pair of real-time operating systems. These should be compatible right down to the interface to the operating system so that we can have both user programs and LDP application programs be usable without change on each system. This implies compatibility between our BASIC and our FORTRAN. We also need data and program transportability between our real-time operating systems and our timesharing operating systems.

A. OOD PRODUCT STRATEGIES

This section contains a description of the products and product strategies LDP requires from OOD to carry out its five-year business plan. The LDP systems strategy, described in the first section, demonstrates the importance of OOD products in our LDP systems. We will continue to need products in the future on which we can base our LDP systems and which correctly address the needs of our marketplace. Since our strategies change over a time as market needs change, we need to work together to develop a flexible long-range strategy which maintains our competitive position in terms of price and performance over a wide range of products. This requires close communication on product strategy issues and this LDP Red Book is meant to be an important part of that process. The chart at the end of this section shows the projected spending in OOD by LDP over the next five years.

Each of the specific subsections in this section is split into a long-range strategy and a short-range strategy. Long-range strategies are meant to be product independent and to describe directions that we feel are required for our markets. The short-range strategies are specific to the issues which we will face over the next six to eighteen months and relate to decisions that have to be made within that period.

B. OOD SYSTEM STRATEGY

The LDP marketplace has learned to expect that any processor and/or system which they purchase will be a member of a family of such processors and systems. For example, when a customer purchases an 11/34 system, part of the reason he purchases the system is that it is a member of the PDP-11 family, and as such he can expand or contract his needs and purchase other systems which will handle them.

Compatibility with other systems in the family is an important buying parameter. Likewise, the user expects families of disks, tapes, communication equipment, etc. to make up systems at various levels of functionality and performance. To operate these systems he expects software with various kinds of functionality and performance such as small Real Time single-user systems and multi-user multi-task Real Time systems. The customer also expects multiple language support under each of these operating systems.

In the current range of price and performance we can reasonably expect from a single processor based system, the LDP marketplace requires four separate processors to make up a complete family. Three of these processors would be used to make a medium, high performance, and low performance system. The fourth processor will be used to create a "hotbox" or high performance FORTRAN oriented system. For example, within the 11 family the 55 can be considered the "hotbox" machine and the 70 is considered the high performance machine. Likewise, the 11/34 fulfills the need of the medium performance CPU and the LSI-11 based 11V03 serves the need of the low performance, low-priced CPU. Each of these processors requires an appropriate set of peripherals and software to create a range of low, medium and high performance systems.

The long-range strategy for specific product developments within this family over a time also relates to the various sized and priced CPU's. Typically, new functionality and new performance levels should be introduced at the high end of the family. Over a time this functionality and performance should be moved down from the high performance to the medium performance systems and from the medium performance to the low performance systems. Examples of this within the 11 family are the FPP, memory management, cache memories, physical address space extensions, multi-user multi-task Real Time systems (such as RSX-11M to RSX-11S) and disks such as RK and RSL.

These guidelines have been used in developing the specific product strategies shown in the following sections. These sections are split up into CPU's and Systems, Hardware and Peripherals, and Software. These are further split up by the specific product areas and each section is meant to be somewhat independent in that they can be without reference to the rest of the LDP Red Book.

<u>LDP Spending for Central Engineering</u>

FY77	$4.9M
FY78	$7.2M
FY79	$9.5M
FY80	$12.1M
FY81	$15.7M
TOTAL	$49.4M

II. CPU AND SYSTEMS STRATEGY

A. Small CPU's

Long-Range Strategy. Our long-range strategy for small CPU's and systems ($<$ \$25K) is to develop a family of LSI products based on Q-Bus technology. Systems which use these processors should achieve greater functionality and performance for approximately a constant cost over a time. This functionality and performance should be achieved by both the availability of faster, more sophisticated chips and by utilizing other features of the system such as WCS and ECS. One specific type of functionality and performance we require is to make an FPP compatible floating point unit available down to the lowest processor level to allow FORTRAN IV + to run on all systems.

With regard to the costs of these systems greater emphasis should be placed on system packaging. Packages are needed which are both lower cost than currently available and which provide a commonality of appearance that makes them distinctly a DIGITAL product. Another necessary cost reduction must be made by designing products with lower FA&T costs.

The maintainability and serviceability of these products will also become more important in the future. A $5K to $10K system must be self-installable with regard to both hardware and software. This implies easy to use systems and high quality user documentation. It also means simple, modular construction which can be easily configured and put together by our customers on site.

Short-Range Strategy. Overall, we need a firm plan for a set of LSI processors. We do not currently understand the OOD strategy in this area (i.e. Fonz-11, SC-11, Annie Oakley, etc.)

The Krypton system is too expensive to be used for our very low-range systems. However, it is a good start at repackaging our systems such that we can sell modular instrument type systems into the laboratory marketplace. These systems will sell at a minimum of $10,000 and do not address our need for a product with an entry level price of $5,000 to compete with programmable calculators.

We, therefore, require a product below Krypton, such as VT100 which is upward compatible to Krypton and has enough expandability to allow LDP to create a very low-cost system.

Another aspect of the packaging of Krypton which is important is to make the system easy to use and approachable. By approachable we mean that an unsophisticated computer user should feel confident about his ability to use the knobs, switches and buttons etc. so that he feels comfortable operating the system. It must also be a portable system which can be carried or wheeled from place to place in a laboratory.

The inclusion of Graphics within Krypton and the VT100 is an important feature for the lab market. We plan to provide several different kinds of Graphics capabilities for our low-end lab systems.

Specific goals of our short-term strategy should include:

* Hard disks on Krypton and 11V03
* Increase the expandability of the 11V03
* Develop a <u>very</u> small stand-alone system
* Graphics should be included with Krypton FCS
* Begin development of an upward compatible version of Krypton based on the Fonz-11 which has all the functionality of the current 11/34 (i.e. FPP, EIS, etc.)

B. Medium Systems

Long-Range Strategy. Medium systems for the LDP market systems which range in price from $25,000 to $75,000. In general, our strategic need for CPU's for our medium systems is to bring the functionality and performance of our current 11/70 systems down into this price range. The key requirements here are the ability to create large systems and systems with large throughput capabilities in this price range. This means a need for large physical address space, a high-speed I/O Bus with a direct port to memory, parallel I/O and CPU operations, and support of large peripherals and software operating systems.

We also need a member of the VAX family in this price range. This will provide the kind of large system capabilities outlined in the first paragraph in addition to fulfilling the need for large program address space required for many LDP applications. In addition, 32 bits is becoming a lockout spec in our market, and we are losing substantial business due to a lack of a 32 bit offering.

A further need in our market is for a "hotbox" processor which is clearly the fastest FORTRAN machine offered. Today this need is being filled by the 55 as a successor to the 11/45. For the VAX family, this requirement should be met by a version of STAR with the high performance I/O, LSI console, and some RAS features removed and features added to enhance the CPU speed and performance. The COMET then becomes the medium price system CPU with priority given to cost rather than performance.

Within the next five years we see a need for the FPP to be enhanced and evolve into an array processor. Such an array processor would have capability to execute the current set of floating point instructions and to execute additional instructions for such array processing needs as the fast Fourier transform. These added instructions could also include fixed point arithmetic for doing indexing operations and special instructions for doing such calculations as the FFT butterfly calculation.

Short-Range Strategy. We need the 11/60 as soon as possible. It should be introduced with emphasis on the performance and speed of the basic system. It should be offered as a machine with 55 like performance at 11/40 prices. It should be announced with real time enhancements using WCS to improve the performance of the system when running RSX-11M. Mid-life kickers should minimally indlude the addition of physical address space extensions.

Also included in this short-range strategy is a requirement for 11/34 enhancements. These are:

* Cache memory
* Improved packaging and power supplies
* Improved expandability
* Low cost RK06 controller

C. Large CPU's

Long-Range Strategy. Our long-range strategy for large CPU's is to move entirely to a VAX architecture over the next two to five years. The key parameters of these systems include large system capacity, large virtual and physical address space, high throughput through use of a high speed I/O Bus, and concurrent CPU and I/O operations. Heavy emphasis must be given to their reliability, availability and serviceability. Compatibility over a broad range of networking configurations is also important for these systems.

This level of systems should be used to introduce new kinds of functionality and performance. For example, this should be the system to first offer the array processor as described in the previous section. We should also look to introduce other kinds of new functionality and enhanced performance on these systems which can later be reduced in cost to be offered on our lower-priced systems.

Short-Range Strategy. It is vital to the LDP strategy that the first VAX product, STAR, be kept on its current schedule for first customer ship in Q2 of FY78. This product should be the most sophisticated, highest performance 32 bit mini-computer system on the market. Emphasis should be placed on large systems capability. High system throughput in the style of the 11/70 is the most important feature. The performance of the system as a FORTRAN machine will also be important. We must be able to offer a system with higher throughput performance than the 11/70 at a competitive price.

This machine will be used within LDP as a host for our large lab systems. Therefore, DCOPS must be supported at FCS to support RSX-11M/S, RT11, and RSTS nodes. The communications options available at FCS must include a high-speed MUX for synchronous and asynchronous operation; support for a high-speed multidrop link is desired as early as reasonable after FCS.

III. MEMORIES STRATEGY

Long-Range Strategy. We would like to have as few a number of options in the memory area as possible. As a minimum, these should be at the performance level of most of our competitors. Also, they should be priced such that we can provide systems at a competitive price relative to our major minicomputer competitors.

We would like to see a memory strategy which addresses CCD and bubble memory utilization.

Short-Range Strategy. We would like to see us move towards small MOS memories (i.e. small in terms of overall system size) using the latest MOS technology. On these MOS memories ECC is mandatory. We should be designing for the 16K chips as soon as possible.

We need a multiport memory on the Q-Bus to allow us to create clusters of LSI-11's in a multiprocessor configuration and to use the Unibus/Q-Bus bridge module with the LSI-11 to create an intelligent I/O controller for larger machines.

IV. HARDWARE AND PERIPHERALS STRATEGY

A. Disks

Long-Range Strategy. LDP is in substantial agreement with what we perceive to be the long-range disk strategy for the Corporation. Our perception is that OOD is developing a strong in-house capability to design and manufacture our own disks. The specific implementation of this strategy is to design and build the small disks in-house and buyout the large disks. It is also our perception that the size of the disks which we develop and build in-house will increase over the next five years to the point where we design and build our own disks over the entire range of our needs.

At the low end we must develop a family of disks which are both low in cost and competitive in performance. The packaging of these disks is particularly important and we must be able to package them economically and compactly in our small system offerings. Family compatibility of disks and controllers across the range of these products is essential.

In the mid-range we require a family of disks which span the performance of 15 megabytes up to 100 megabytes. These must also be a compatible family of products which are price and performance competitive with other disk offerings.

At the high end, we should continue to buyout large disks similar to our current RP series until our in-house capabilities allow us to produce them more economically. These disks must be performance competitive and offer high reliability, availability and serviceability. We should have the largest disk offering of any minivendor. We must also become price competitive at the high end with IBM, Burroughs, NCR etc. At the high end we also need better multiport support for our large disks.

Overall, we must integrate our disk strategy with our other mass storage products. This includes floppies and CCD's as they become available. Also, some large, mass storage devices, such as the data cell offered by CDC, should be proposed as a member of this family of products.

With regard to I/O Buses which support the various kinds of controllers and disks, the following is a statement of our needs. For the low-end disks it will be sufficient to put these only on the Q-Buses. Mid-range disks should be available on both the Q-Bus and on the Unibus or some other VAX family I/O Port. The high-end disks will only be used with the Mass Bus or VAX Bus, whichever is used there.

Short-Range Strategy. The RL01 should be a complete product replacement for our current RK05's. The RL01 must be enhanced with new product offerings which create a family of low cost and competitive performance low-end disk products. This family of products should have a compatible set of controllers and disk formats. It should be possible to take a disk from any member of the family and use it on any other member of the family.

Fixed and removable systems like the RK05-F should also be developed based on the RL01 technology. These products offer a very price competitive way to offer reasonable storage and performance capabilities.

The priorities for the low-end disks should be first price and second performance.

For our family of RK products we have several requirements. We need a mass Bus controller to provide more I/O speed with concurrent processing for this family of products. This will bring the functionality and performance of our larger high throughput systems down to the level of our medium-priced systems.

We need more emphasis on maintainability and serviceability for our disk products. The disk is a very key element in the systems we sell and when it goes down the entire system is down.

The possibility of a bigger RK type drive should also be examined. We should continue to look at these larger RK's drives and develop new products up to the point where we can buyout a drive in the RP family for less money than we can produce it in-house.

For our larger sized RP buyout disks, we need to develop a mechanism between the hardware and the software to handle bad blocks in a more efficient manner.

The reliability, availability and serviceability are key elements in any disk in this family. The combination of controller and disk should be as reliable as we can afford to make it. More hardware should be added to enable Field Service to better service these products.

We need to reduce the price of our large disk offerings and in particular we must make them price competitive with the mainframers. We should begin a program to buyout a disk with 200-300 megabyte storage capability.

B. Tapes

LDP is in concurrence with the current two tape strategy. However, we have several specific needs in our marketplace relative to both the low performance and high performance drives.

The current costs of tape drives are too high; the design of the low performance drive should give top priority to keeping the costs to a minimum. This includes the costs of packaging the drive in a system such as rack space and power. The drive must be rack mountable and preferably two per rack with the extra space usable for other options.

The high performance drive should balance cost and performance. It should have options for 800/1600 BPI and 6250 BPI at a tape speed of 125 IPS. The cost should be $5,000 or less.

For VAX there is a need for some extremely large storage device. Investigations of data cells and BEAMOS devices should be made and product proposals generated for consideration by the Product Lines.

Floppies

Long-Term Strategy. The floppy fills the hole below the RSL type disk in LDP systems. We need floppies with up to approximately 1 megabyte capacity or where the RSL becomes a cheaper solution. The speed of access of these devices is the main performance requirement for our systems.

We would like to see a renewed focus on reducing the cost of a dual-drive floppy, hopefully to the $500 level. This is needed in order to have more price-competitive low-end systems, and to provide a better "fit" vs. RSL.

One of the key factors for using floppies in our LDP systems is the type of packaging available. We need both a small table-top package and a rack-mount package as well.

IBM capability is also a desirable feature.

In addition to that, we need to have an integrated strategy for including other kinds of devices in our low-end systems at lower cost than floppies. Such devices as the 3M tape cartridge and CCD's can play an important role in a family of low-cost mass storage devices.

Short-Term Strategy. We should have a DMA type interface for the Q-Bus for our existing in-house design. This is needed to allow sufficient time for the processor to do A to D conversions and other parallel tasks while still maintaining Input/Output from the floppy. This is particularly important in those applications in which we are acquiring data from some analog process and storing it on the floppy drive. In such systems, both the transfer rate and the storage capacity become important parameters.

C. CRT Terminals

Long-Range Strategy. There are several features we would like to see in our future family of CRT terminals. We would like to see better human interfacing than was present in the VT52, especially in terms of the feel of the keys to the operator and making the keyboard portable. There should be a standard DEC keyboard for all members of the family with the option of providing a special purpose keyboard as well. Another feature which is desirable is to have a tactile interface on the screen such that the user can point with his finger and use it like a light pen.

Short-Range Strategy. LDP would like the ability to create applications specific CRT's from the corporate family of

CRT's. This would be done by using special purpose keyboards and the tactile interface with special software.

The DK series should replace the VT52 with a less expensive, higher performance unit. The DK will also be used as a low-end basis for both inexpensive graphics terminals and small lab terminals.

D. Printers

Long-Range Strategy. LDP long-term needs for our printer family include the need to provide remoting capability via a modem. We also need to be completely EIA/CCITT compatible here and in Europe. Multidrop capability is also an important feature.

For the large line printers we feel the need to have the fastest and largest available from any mini-computer vendor. This is particularly important for our Medical Information Systems Group.

We also need some form of hard copy to go with the VIDEO option on KRYPTON. It is not clear whether there should be a follow on product to the current FAX Copier or whether some form of copy using the printhead from an LA00 would be better.

Short-Range Strategy. We need a 600 line per minute printer which is very inexpensive. We need portability for the LA120. We need to be able to use the LA120 with new microcode for graphics plotting capability.

E. Communications

Long-Range Strategy. Our interpretation of the corporate communications long-range strategy is for the development of fewer, more comprehensive options on the Unibus but with a few options for the Q-Bus. The addition of microprocessors for throughput enhancement is also an objective. Overall there is a move from asynchronous towards synchronous communication options, and toward a high-speed Unibus controller for multidrop operation.

LDP is not in complete agreement with this strategy. We would like to see more emphasis on Q-Bus (as opposed to Unibus) options, with Unibus capability through the bridge module currently in development. This conserves engineering resources and simplifies communication product space by eliminating the need for both Q-Bus and Unibus versions of each option.

We would like to see the hardware strategy closely coordinated with our DECNET strategy and we would like to see improved capability with IBM and other vendors both over low speed and high speed lines. We feel that moving toward synchronous communication is a good thing to do if it does not increase cost for a given data rate. This cost should be relative to the total system cost, including the cost of the communication lines.

LDP also needs the ability to have a multiplexed adaptor for large machines which can communicate over high speed lines to small machines. In particular, we require 1M baud on at least eight lines on 11/70- and STAR-class machines. This requirement may be mutually exclusive not only with the Q-Bus option strategy but also with a Unibus implementation and may require bypassing the Unibus and going direct to memory via an 11/70 or STAR Mass Bus port.

We also need to be able to integrate in a high speed multidrop serial bus to our communication strategy. This will be important in our Krypton and Minc plans for creating clusters of analog front-ends tied to either a large VAX machine or a large Unibus machine.

Short-Range Strategy. We need a lower cost DH-11 with a greater number of lines.

We need a multidrop Q-Bus adaptor for hosting small processors on large Unibus machines.

V. SOFTWARE STRATEGY

A. Real-Time Operating Systems

Long-Term Strategy. LDP's long-term requirements are for two real-time operating systems. One is a small single-user system and the other is a multi-task, multi-user, multi-function operating system. Both systems should be compatible at both the user interface and the application program interface. The systems should appear as identical as possible to the user of the small system and any one of the users on the larger system. A single kernel operating system should be used as the base for both.

The single-user system should emphasize real-time performance. The user should be able to fully utilize the capabilities and speed of our analog front ends. A particularly important parameter is the throughput speed from an A/D to disk.

Ease of use is also important on the single-user system. User commands, NCR's, documentation, the editor, etc. must be easily understood and easy to use.

The large MU/MT operating system requires high functionality and high performance. The real-time performance available to the sophisticated user should be equal to or better than the single-user system. Increased functionality should be offered with this system to make it equivalent to RSX-11M. Batch and limited timesharing should be offered to provide multi-function capability. The increased functionality should be optional at sysgen time to enable the configuration of small as well as large systems.

Both of these types of operating systems will be required for the VAX machines. The larger, more sophisticated system will be used on the STAR class machine, and both types of operating systems will be used on COMET. Therefore, it is important that they be able to take advantages of the real-time aspects of the VAX architecture and that they utilize the hardware in such a way as to enable high-speed data acquisition and high throughput to disk.

Short-Term Strategy. RT11 should be kept as a small, high speed, single-user system. It should be made easier to use so that it can be sold to less sophisticated users. Specifically, LDP has no requirement for the PLAS option. The current level of performance of RT should be moved to the Krypton systems via the kernel operating system.

The real-time performance of RSX-11M needs substantial improvement. We should be able to at least acquire data at a rate in excess of 50 KHz instead of the current level of 1 KHz. In addition, the system should be made more highly interactive through the use of an improved terminal handler. The system must be made easier to use (i.e. the current Sysgen time is 8 hours).

We also require a smaller version of RSX-11M which limits functionality in favor of performance and reduces the gap between RT and M. Our desire is to move RT-11 based products to an RSX like operating system.

In the near future we should plan specific WCS support for M using the 11/60 WCS option. These WCS options should emphasize:

* Real-time performance
* Context switching
* Fortran library execution speed
* Interrupt processing
* FFT signal processing

B. Time Sharing and Multipurpose Operating Systems

Long-Term Strategy.
The LDP long-term strategy for these products is to increase the current level of functionality and performance on VAX. It is assumed that this capability will be provided within one year after FCS of the STAR system using a version of STARLET.

With regard to other strategic requirements the most important input from LDP is from the Medical Products Group. For this pass of the Red Book they have decided to concur with the current commercial product line inputs.

Short-Term Strategy. With regard to RSX-11D and IAS we would like to have better support in terms of selling aids and market descriptions. We would also like to support DECNET, APL, and PLAS. We also need a better batch capability. The number of users should also be increased to a reasonable limit and performance measures given to the field for various levels of user operation.

We also need some performance measures that give accurate statistics with regard to performance versus number of users. Some transaction processing improvements are also desirable.

C. Languages and Utilities

Long-Term Strategy. Long-term language strategy should include a single high performance real-time oriented FORTRAN for the VAX machine. Likewise, it should include a DEC standard BASIC which has enhanced real-time capabilities and high performance for normal interactive programming. Our data base requirements dictate that DBMS and RMS be transported to VAX in native mode to maximize performance. These need to be integrated with some distributed data base solution under VAX and DECNET. These data based programs should also maintain a CODASYL standard.

With regard to utilities, there should be a standard editor which is CRT oriented for both Krypton and VAX. Emphasis here should be on ease of use and consistent in human interfacing across all operating systems. We also need a symbolic FORTRAN debugger for all our FORTRANS.

Short-Term Strategy. To stabilize our current 11 based FORTRANS down to two. Namely, a single high performance FORTRAN in which performance is the key parameter which is emphasized. The second would be a smaller, easy to use FORTRAN for our smaller RT11 system. We should also include the ability to handle large arrays on the FORTRAN IV Plus where we currently support IMS. We should do a rewrite of the arithmetic library to improve its performance to be competitive with that of the Eclipse. We also need WCS enhancements for the 11/60 so that it can be marketed as a real-time FORTRAN machine. With regard to BASIC we should begin immediately to provide upward compatible migration path from BASIC to BASIC PLUS to BASIC PLUS II. We also want real-time enhancements in BASIC which are extremely important in our marketplace. We also need to maintain the old BASIC because of its capability with all the existing LDP application packages and much of the software written by the LDP customer base.

D. DECNET

Long-Range Strategy. Our long-term strategy should be to improve DECNET performance with particular emphasis on terminal oriented networks. This improved performance should include distributed data base systems, and better support of some other large main frames protocols. These large main frames should particularly include DECsystem 10 and 20 host support for both 11's, Kryptons, and VAX machines.

Another important long-term consideration is the definition of a high-speed multidrop serial bus which is integrated in with the rest of the DECNET architecture. This is particularly important for LDP's plans with regard to MINC and other Krypton based systems we may develop. We need the capability of supporting clusters of Mincs up to 10,000 feet from a host processor. Specifically with regard to VAX, the DCOPS option for STAR should be the primary method of implementing DECNET. However, for the COMET system we need an integrated DCOPS capability in order to maintain low-cost systems.

In general, LDP strongly feels we need more emphasis on low-end system support under DECNET instead of the current direction which emphasizes larger and larger systems. We also need to have a measure of DECNET performance for various configurations. This should be an ongoing program with some quantitative measurements which can be explained to our customer base. We also need support for a Graphics protocol under the DECNET architecture.

Short-Range Strategy. We need to have support for RT-11 DECNET with improved performance and a clean up of the current system. We also need LSI support for RT-11 and RSX-11S under DECNET.

We need to freeze our current protocols so that we can establish a firm base for our customers and convince them that we indeed have a viable protocol. We should also emphasize maintaining our current announcement and delivery schedules.

We need to improve the RSX-11M DECNET performance and throughput particularly for RSX-11M to RSX-11M systems and RSX-11M to other operating system nodes.

VI. LDP DEVELOPMENT GROUP PRODUCT STRATEGY

INTRODUCTION AND OVERALL STRATEGY

The LDP product development strategy has been divided into two segments; small systems, and medium and large systems. Each of these segments will be treated independently in terms of long range and short range strategies; each of these segments is briefly described in the following paragraphs.

Small systems: Small systems include modular system products which consist of a processor, signal conditioning options, and easy to use software. These products will be mobile for individual usage to connect to laboratory experiments, etc. These products will be produced in high volume and will be capable of either hosted or standalone operation. In essence, small systems require combining signal conditioning, processor, and software into an economical, easy to use product for utilization by persons at the calculator programming skill level.

Medium and large systems: These systems include options and software which can be utilized with Unibus 11 systems to make them applicable to the real-time laboratory market. Medium systems have been the traditional product area for LDP development. Products offered here previously include the LPS A/D and D/A options for Unibus systems. This product area will include I/O options such as A/D and D/A converters, clocks, signal connection panels, plus options for new areas such as array and signal processing, options for the VAX family, and interface standard options.

LDP STRATEGY ASSUMPTIONS

The overall LDP development strategy assumes the following points:

* It is possible to develop a single family of signal conditioning and conversion options which can be effectively marketed for both LSI-11 bus and Unibus systems.

* The LSI-11 bus will become the dominant program I/O bus for options in the small and medium systems area; this bus will be the standard option bus for small processors, and will be available at low cost on medium scale systems through bridge modules.

* The corporation will continue to develop processor peripherals and software kernel components at the low end of the systems market, and that these tools will be better suited for developing low cost small systems than current tools. For instance, future low end systems will provide higher levels of interface integration, reduced power consumption, and require lower cost packages.

* It is practical to use I/O processors or "demons" with modular program I/O signal conditioning options to achieve high performance.

* We will develop some market specific products when necessary to achieve a cost/performance goal, but will try to satisfy market needs with general purpose laboratory products.

* That the corporation will develop better base programming tools for implementing easy-to-use software products for unskilled users.

A. Small Systems Strategy

Five Year Strategy. Products for this segment of the LDP market will be developed in accordance with the following priorities:

* Low Cost Relative to Computer Products
* Approachable, Easy to Use
* Functional

The goal in this product area is to develop products which allow persons with minimal or no previous computer experience to install, start-up and apply these systems without outside assistance. These products will bring computer technology to persons without the budget and/or capability to purchase and utilize DEC's current computer systems. This type of system product will have to compete against the upper end of the sophisticated calculator and associated real-time peripheral products offered by Tektronix and Hewlett-Packard, as well as against "smart" instruments and intelligent terminals. Our cost goal in these products is to be reasonably cost competitive with the calculator based products, but offer better performance and functionality to the user. Small systems products will consist of the following components:

* Processor and memory
* Real-time I/O options (A/D, D/A, clock, digital in, digital out, etc.)
* Easy to use, self-diagnosing software
* Program load/mass storage via a standalone mass storage option, or a communication line to a host system
* Keyboard and video display
* Hard copy option

Hardware Strategy. The product will have to be mobile, accommodates the processor, mass store, communicates, and I/O options, and provide an integrated appearance and convenience to the customer. This product will initially be offered as several hardware modules packaged together within a roll around stand, but must trend over the next five years to a lower cost portable package. Initial products will probably be offered with 32K bytes of memory; this will probably grow to 128K bytes of memory over the next five years with this additional memory being utilized to make the product easier to use. The product will require mass storage for program load when used standalone. The product MLP should be in the ten to fifteen thousand dollar range with mass store when initially introduced and this should reduce to approximately five thousand dollars over the next five years.

The product will also have to be connectable to a host computer. Connection to the host eliminates the requirement for local mass storage and thereby reduces the system cost for customers with multiple systems. The host system will also provide computational and data base capability not possible with a standalone small system. The product should be hostable either via a local high-speed serial line with multi-drop capability, and possibly via common carrier. It is essential that this capability be achievable with minimum satellite memory overhead so that a minimum user programmable system can remain under 32K bytes of memory.

Early products will utilize a single processor with program I/O options. The performance of this system will be enhanced by the addition of a simple DMA I/O control option. Later systems (3 to 5 years) may include an I/O processor option for greater I/O performance. This I/O processor will be instruction set compatible with the main processor to minimize the software tools required for its support. To achieve this, it is important that future processors for small systems be integratable into the products at both the chip and board level.

The initial product (MINC) will look like an instrument package with signal conditioning options having indicators and switches to be set up by the user. As systems size and cost is reduced, it is anticipated that a video graphics capability will replace these module indicators and controls, and all module functions will be performed under processor control; this will allow system size to approach terminal size and achieve higher user portability.

The initial product will offer the DCG IEC interface bus option (IEEE 488) to allow the system to operate with IEC instrument interfaces under program control. Approximately two years out it is anticipated that a more sophisticated IEC interface will be developed to allow DMA interfacing for high throughput, and to allow the ability to rapidly pass and receive control from other devices on the bus.

The price/performance trade-off for a cartridge tape drive/ CCD (or bubble)/MOS RAM memory hierarchy for laboratory small systems will be investigated so that we can better optimize small system cost/performance in 2 to 5 years.

Software Strategy. Software for this segment of the market requires development of real-time BASIC along with self-instruction and diagnostic capability such that an inexperienced user can perform experiments with reasonably high throughput after one or two hours of self-instruction. This will require intensive development of our Basic software over the next few years. The software for this product must support resource sharing from a RSX-11M host system via either dedicated or common carrier line. This resource sharing should require no additional skills on the part of the user, but should allow the user to utilize the host system for loading, mass storage, and high performance processing.

The base operating system for small system products should move from the KRYPTON RT Kernel to a Kernel more compatible with RSX like host system operating systems.

Two Year Strategy. The initial small system product offering is the MINC (Modular Instrumentation Computer). The MINC utilizes the current LSI-11 modules, and the same electrical backplane arrangement as the KRYPTON and Industrial PCS. The MINC package accommodates eight standard LSI-11 and KRYPTON modules and up to eight signal conditioning modules. The signal conditioning modules have an attached instrument front panel which contains signal connections, and switches and indicators for the user to set up the modules for experiments. The initial MINC product will consist of the following:

* MINC box with LSI-11, KRYPTON, and signal conditioning options
* VT100 with VT55 like graphics
* Dual Floppy Disk

The above configuration appears to offer the lowest cost product for the immediate future. MINC modules being developed include the following:

* A/D Converter
* Clock
* D/A
* Preamp
* MUX
* Digital Out
* Digital In
* DMA Controller

Additional options being considered for development over the next 12 months include the following:

* High Speed A/D
* High Resolution A/D
* Combination Digital I/O
* Programmable Gain/MUX Modules
* Parallel Sample and Hold
* Analog Filter/MUX Option
* Synchronous Sampling Option

MINC software will be based on the KRYPTON Kernel and Basic, and will be able to operate in two modes:

* Non-programming mode which is self installing
* BASIC language processor with Real Time application software

This software will provide for user installation and self-learning in the first mode.

A lower cost, entry level product for this market is being investigated. This product would run a subset of the MINC software, and would consist of:

* VT100 with VT55 like graphics
* LSI-11 based system board
* Cartridge tape
* Signal condition board
* IEC Bus port

Front panel functions would be achieved via interaction with the graphics and keyboard.

B. Medium and Large Systems Strategy

Five Year Strategy. The goal of development in the medium and large system area is to adapt standard DEC products to LDP markets by adding additional software and/or hardware options to provide a system which more directly serves the specific requirements of the LDP market. Current LDP products for these markets have been developed with the following priorities:

* Performance/Features
* Cost
* Ease of Use

Future LDP product developments for the medium and large system markets will have the following priorities:

* Ease of Use
* Performance/Features at the system level
* Cost

The following are the major areas of future development for medium and large systems:

I/O Options: LDP will continue to develop and support a range of I/O options from low cost low performance to higher cost high performance for analog and digital signal interfacing; these options should apply to both Unibus and Q-Bus machines.

I/O options can be divided into two categories: Those intended for instrument applications, and those intended for on line lab applications. Options for instrument applications are typically limited to a few I/O channels, a large number of features, and relatively high transfer rates. These options are typically used for single user laboratory experiments with the system being reconfigured and reconnected between successive experiments. Options developed by LDP during the last year have been directed towards this market. The second major option application area has begun to emerge in the last few years, that of on-line laboratory applications. In on-line applications the I/O options tend to be permanently connected to the lab instrumentation or experimental setup. I/O options for this area tend to require higher numbers of input channels, less features, and are more price sensitive. I/O options will be developed which address both of these application areas; these options will be supported under applicable operating systems with high level language calls.

A/D converter products will be developed for both of the above market segments; separate products will be developed to provide a low cost, medium performance conversion capability and to provide a high performance capability. Reduced costs A/D products will be characterized by twelve bit converters, more input points per module, and reduced costs. Increased functionality will be provided by programmable gain options which will increase the signal dynamic input range, input signal conditioning and filtering options, and higher conversion rates at 12 bits. LDP has traditionally concentrated on low numbers of multiplexer channels per converter, i.e., 32 differential channels per converter maximum. Higher input point counts will be supported on future converters, and will allow implementation of multiplexers for both the instrumentation and lab automation markets.

Digital to analog converters will retain their current functionality, but will undergo cost reduction while maintaining current performance levels. Digital I/O options such as clocks will probably maintain functionality and be cost reduced; increased functionality will be handled by I/O processors in the future. Digital signal conditioning modules will be characterized by more I/O points per module to further reduce module costs.

High speed digital interface options for large systems will be developed to provide a 32 bit I/O path with very high bandwidths for connection of high performance laboratory instrumentation.

Distributed System and I/O Capability: LDP requires three different capabilities for front ending medium and large systems. These are as follows:

* Programmable Nodes (LSI-11 or F-11) which connect to a host system via a shared high speed link over which these nodes can communicate data to and from the host, be down-line loaded, and remotely diagnosed. Typical nodes sell for $10,000 and up each.
* High Performance Remote I/O: A remote I/O box with analog and digital capability which requires no programming by the user. This capability is currently provided by the ICR, but at a premium cost due to the lack of a multi-drop controller. This product should sell for approximately $4,000, and provide a high performance capability to compete with other vendors remote I/O products.
* Low Cost Remote I/O: A fixed function, low I/O rate interface device for connecting instruments to a host computer. Functionality for this device would include one to four analog input channels and approximately 16 TTL digital I/O signals. This product would provide an effective method of interfacing laboratory instrumentation such as GC's, blood analyzers, etc. to a single host computer. A product similar to this is currently marketed by HP, and sells for $1,200.

The means for implementing these products is unclear at the current time, but all require a multi-drop communication capability to offer competitive adaptability and flexibility to the user. Whether this capability is obtained from the corporate Uniplex project, COMM, or IPG is unclear at the current time. It is our expectation that these products can be compatible with a corporately supported communication product.

Application Software: The current laboratory software packages such as lab apps will be improved by providing better documentation, user guides, application notes, and improved subroutines. More high level functions such as signal processing will be added to these packages so that lower skill level users can more easily utilize sophisticated processing subroutines with laboratory systems.

Much of the current LDP software offering is based on RT-11; it is anticipated that RT based products will be moved to an RSX compatible operating system over the next few years to reduce the number of unique operating systems supported by LDP, and to ease the problem of implementing distributed systems with incompatible operating systems. LDP desires to support RSX compatible satellites on an RSX compatible host operating system.

Signal Processing: LDP signal processing capability will be developed by offering products which utilize either WCS or a processor option to provide an attractive level of performance on standard laboratory systems. The goal in this area will be to develop a system which addresses both acquisition, signal processing, and storage or display capability.

Interface Standards: Interface standards for both parallel and serial interconnect will be tracked to maintain awareness of their capability and to assess when and if they are to be significant in the market place. CAMAC has both parallel and serial I/O standards which have been adopted by IEEE. These standards will be followed and probably supported only through CSS products, or by OEM. Follow-ons to these serial and parallel standards are being developed; these will be followed.

The IEEE 488 bus standard will be implemented by LDP with two levels of cost/performance. A program I/O version (or bit banging version), and a DMA version for higher throughput; appropriate high level language calls for these interfaces will be developed. The program I/O version will be the current DCG offering; the DMA version will be developed.

Two Year Strategy. The two year strategy for medium and large systems described in the following paragraphs represents those specific development projects that are either currently under way or planned for startup within the next twelve months to implement the five year strategy described earlier.

I/O Options: The LPA-11 (Lab Peripheral Accelerator) is being developed to provide high performance I/O for Unibus systems. This option is based on a KMC and modified DMC option configured to create a second Unibus for real time I/O options. The KMC interprets operating system commands, controls multiple buffers for multiple users, and controls the modified DMC to drive analog and digital I/O options. The LPA-11 will provide analog conversion rates approaching 140 KHz for single users and 70 KHz for multiple user tasks. This option will be supported under both RSX-11M and RT-11. This option will also be offered on the 11/60 along with ECS support to provide higher performance signal processing.

Planned projects include the following:

* Offering the MINC package and I/O options on Unibus machines as a DECLAB system. This would be a complete replacement for the LPS, and would provide additional features such as DMA, preamplifiers, and multiple sample-and-hold options. The MINC would be connected to Unibus systems via the Unibus/Q Bus bridge module being developed by Small 11 Engineering. Software support for the options would be developed under both RSX-11M and RT-11.

* 32 Bit High Bandwidth Digital I/O Option for STAR to satisfy the need for a medium cost, high performance digital interface for laboratory instrumentation.

* A large point count analog I/O capability to reduce current cost per input for applications requiring more than 32 analog input channels, and to provide input filter options. This product would provide low and high speed analog conversion rates to satisfy industrial R&D markets.

* A follow on product to the LPA-11, an I/O processor, which will offer similar functionality to the LPA-11 at reduced cost.

Distributed System and I/O Capability: The Remote-11 product for connecting LSI-11 systems to an RT-11 host is being completed. A limited number of enhancements to Remote-11 are planned for both Unibus and LSI-11 nodes to enhance diagnostic capability, and to support servicing of nodes without a local load device. Support of Remote nodes without a local terminal is not anticipated.

Planned projects include the following:

* A remote instrument interface for both PEAK-11 and PDL to interface lab instrumentation to a host computer. This will be a low data rate, fixed function device with one or two analog input channels and approximately 16 TTL digital I/O signals. This device will use dual slope A/D converters, and have a logarithmic or programmable gain amplifier to provide the wide dynamic input range necessary for GC's, etc.

* Multi-drop serial interface for hosting LSI-11s with an 11/34 or larger system. The market requirements for this product are currently being defined to determine required data rates, distances, and response time.

Application Software: Planned application software projects include the conversion of PEAK-11 to run under RSX-11M rather than RT-11, and a general upgrading of the lab application software package for both FORTRAN and BASIC.

Signal Processing: Current plans are to define and develop signal processing support which utilizes ECS on the 11/60. This will be offered along with the LPA-11 to provide a high performance data acquisition and signal processing package. Further work will be done to define an array processor option which would offer extended signal processing capability.

Interface Standards: Software support under RSX-11M and RT-11 will be developed for the DCG IEC-11 IEEE 488 interface option. This option will be offered on both Unibus and Q-Bus machines; the Unibus/Q-Bus bridge module will be used to offer this option on Unibus systems. A high performance, DMA IEEE 488 interface option will be developed for Unibus machines; this will be necessary to remain competitive with Hewlett Packard's offering on mid-range systems. The DMA IEEE 488 interface will be supported under both RSX-11M and RT-11. There are no plans to develop CAMAC compatible options by LDP; however, the need for a CAMAC interface with software support is being reviewed.

MEDICAL MARKET

Long Range Strategy:

Long range medical market strategy is concentrated in three areas; image analysis, medical laboratories computer automation, and medical information. Product development in each of these areas will be enhancing standard DEC products for data acquisition, and adding application software for the medical market.

Image Analysis: The current Gamma-11 System will be used as a base on which to develop a range of Gamma image processing products and from which to expand into other image analysis areas. Image analysis systems are expected to range from low cost portable systems on up. The low-end could be used as a standalone with portable scanners, or as satellites supported by host systems. Mid-range systems will be disc based and will be capable of simultaneously supporting multiple scanners if the required disc throughput rates can be achieved.

Image analysis areas for future growth include ultra sound, computed tomography, thermography, etc.

Laboratory Computer Automation: Currently the Programmable Data Logger (PDL). Future PDL systems will shift from RT11 to RSX-11M as a development base. A range of PDL systems is anticipated with lower end systems being based on small CPUs with minimal mass storage and I/O. These will be either used standalone, or hosted in networks by larger systems. Medium range PDL systems will be characterized by one or more discs, the ability to support large numbers of instruments, and the ability to support small PDL satellite nodes. It is anticipated that future intelligent laboratory instruments will tend to offer ASCII interfaces, and that it is essential to be able to support a large number of ASCII port instruments for data acquisition.

Medical Information Systems: Strategy is to grow from current installed base by developing "90% - type" solutions in certain areas. (Initial product will be for ambulatory care).

Short Range Strategy:

Image Anaylsis: The current Gamma-11 product will be extended downward to low cost systems via the NCV11 high speed A/D Converter which will be interfaced to the Q-Bus. This option will allow the introduction of a low cost Gamma system based around Krypton/Minc in FY78. This interface will also offer a lower cost mid-range system based on the F-11 chip set replacement for the 11/34 in FY79. The NCV11 will support Gamma options on mid-range systems via the Unibus/Q-bus bridge module. Exploration of other image analysis areas is anticipated in FY78.

PDL: A PDL system based on the MINC is planned in FY78. This system will provide a low cost PDL capability, and will also provide the basis for a PDL satellite to larger PDL hosts. Support of ASCII instrument interfaces will be developed around standard LSI-11 communication options.

Medical Information Systems: Current strategy is based on two areas: Mumps development and a 90% solution for ambulatory care. We will be preparing an ANSI MUMPS and placing it under RSX-11M to complement the other corporate products in the tools market and will start selling this in FY78. We are also developing a Doctor's Office System based on the Mumps Standard as a joint project with NIH and Mass General, called COSTAR. Test site FCS will be Q4 FY77. Volume will be Q1-Q2 FY78.