CPG MEMO #29

Report on the Numerix Presentation (22 June 84)

Don Wells & Bill Cotton

National Radio Astronomy Observatory
Edgemont Road, Charlottesville, VA

25 June 1984

Two representatives of Numerix Corporation, Peter Alexander and Tom Meckley, visited our Charlottesville offices on Friday, 22 May 1984, to discuss their MARS-432 "super" array processor. Peter Alexander (President of Numerix Corp.), who did most of the talking on Friday, said that he was one of the founders of the company, and that he and another fellow were the designers of the AP. Meckley is the sales rep for Numerix in the Southeast. The purpose of this memo is to summarize what they told us, and to discuss its implications for NRAO computer planning.

The following NRAO personnel attended the presentation: R. Burns, W. Cotton, E. Greisen, K. Hilldrup, F. Schwab, and D. Wells. R. Duquet was unable to participate by telephone link. Alexander seemed interested in making a presentation in New Mexico and G. Hunt, R. Duquet, and K. Sowinski were suggested to him as contacts.

Some relevant names and addresses:

Numerix Corporation
320 Needham Street
Newton, MA 02164-1594

617-964-2500

VP Sales & Marketing: Harold Matthias

Southeast Sales Rep: Thomas A. Meckley
Melbourne, FL: 305-725-7066

Alexander recommends that persons desiring additional information about the MARS-432 should talk with Matthias. DCW talked with Matthias in September 1983 and found him quite knowledgeable technically.

1.0 AN OVERVIEW: THE SIGNIFICANCE OF THE MARS-432 FOR NRAO

The Numerix MARS-432 is architecturally very similar to the FPS AP-120b in both hardware and software, thus greatly easing the usual problem of porting our existing 120b application code to a new AP (Alexander said that the 432 was specifically designed to be a replacement for the 120b). It has two advantages over the 120b:
(1) its floating point pipelines are rated at 30 MFlops, 2.5x faster than those of the 120b, and
(2) it can utilize a much larger memory size than the 120b.
Because of the high degree of architectural similarity to the 120b, we judge that an AIPS implementation for the MARS-432 would involve a minimum of risk and effort for NRAO. But what fraction of the increased pipeline performance could we achieve? From our experience with the 120b we have learned that obtaining higher performance in our AP applications will depend at least as much on getting higher I/O bandwidths as it will on getting higher pipeline speeds. The MARS-432 currently uses a Unibus interface, just like the 120b. We therefore expect that a MARS-432 on a VAX would deliver less than a 2.5x improvement in overall performance. Alexander said that Numerix is currently implementing a new 3 MB/sec SBI interface for VAXes (about 3x faster than the Unibus interface), and that they will have a 3 MB/sec interface for IBM machines in a few months. This means that they will have an I/O performance increase comparable to their pipeline performance increase, and therefore should be able to achieve a 2-3x improvement over the 120b in overall throughput. Unfortunately, it is not at all obvious that present VAXes can properly exploit the improved I/O rates, because their disk drives are still limited by Unibus interfaces. An IBM machine with multiple high speed channels, and IBM disks with performance to match, would probably make a better host for the MARS-432 than would a VAX. Therefore,

> **we recommend that NRAO consider a MARS-432 "add-on" to our existing IBM-4341 as yet another candidate for the "ModComp replacement" procurement planned for FY85.**

Obviously this recommendation must be contingent upon demonstration that AIPS functions on the IBM. Therefore, we are pleased to be able to report that in a test made late last week, UVMAP appeared to work correctly. This means that finally we are able to read visibilities (UVLOD), make dirty maps (UVMAP) and clean them (APCLN) with AIPS on the IBM under UTS. It now seems reasonable to assume that the rest of AIPS will also be made to function soon on the IBM. We further assume that there is no fundamental barrier to installing a new device driver in UTS which would interface to the AP on the IBM channel.

2.0 MISCELLANEOUS NOTES ON NUMERIX AND THE MARS-432

Alexander said that Numerix Corp. was formed with capital from Analog Devices Corp. and Standard Oil of Indiana, and that their stock is not currently publicly traded. While preparing this memo DCW reviewed some old files on various APs, and discovered that the "MARS" line of APs has been around for some years. In January of 1981 a company called "CNR, Inc.-Computer Products Division" in Needham, MA, was offering an AP called the MARS-232. Their glossy brochure, dated 1980, says: "CNR, Inc. was formed in 1972 to undertake advanced R&D in signal processing and communications systems". It also indicates that MARS is an acronym for "modular array processor system". Another brochure, dated January 1983, from Numerix (before the announcement of their MARS-432 model) also described the MARS-232, and it appears to be the same machine, and, in addition, it also mentioned a stripped version called the MARS-132 which was being offered to OEM customers. In summary, it appears that Numerix is probably a new company formed in 1981 or 1982 by ex-CNR people to exploit technology which CNR no longer wanted to develop.

Numerix now has 35-40 MARS-432 machines installed (Alexander didn't know the exact number).

The main buss of the MARS-432 is capable of I/O transfer rates of 20 MB/sec in burst mode. It is supervised by a microcoded processor called the IP ("Interface Processor"), which runs with a 5 MHz clock. "A repertoire of IP service routines in ROM...will satisfy the needs of most users", but "for those users who must write their own custom IP code, a set of macroinstructions and macros are provided for use with the assembler". The IP apparently can supervise up to eight DMA transfers simultaneously.

Alexander insisted that the MARS-432 can reliably transfer data between host and AP data memory while the AP is computing (the old 120b's have never been able to do this consistently). Here is a quote from the Programmer's Reference Manual: "a data memory page-loading feature provides the user with the option of zero overhead background loading of data during time critical program execution. With this form of I/O, no DMA cycle stealing overhead is incurred".

The best present implementation of the MARS-432 is for VAXes under VMS. Numerix also has an interface to the ELXSI 64-bit machines. They are currently working on an interface to the Sperry/IBM byte-channel interface (3 MB/sec), which apparently is a Federal standard. The MARS-432 uses DEC floating point format in its pipes (24-bit mantissa with hidden bit). We assume that floating point data format conversion occurs in the host-AP interface for non-DEC hosts, as in FPS 120b APs, but we forgot to ask about this detail. The Programmer's Reference Manual says: "format conversion is generally implemented on the host interface board, although for some hosts it may be implemented by DP software conversion routines" (i.e., in the floating adder pipes).

Numerix is interested in using dual channel interfaces to host computers to increase I/O throughput. This concept is especially relevant in IBM computers, and presumably depends on the ability of the IP to execute multiple simultaneous DMA transfers.

Alexander said that the Aptec (Aptec Computer Systems, Portland, OR) I/O system is now supported for the MARS-432 as well as for the FPS AP-120b. This is a controller which contains a large RAM buffer memory in order to interface one or more APs to multiple disk drives and to a host computer (generally a VAX). The Aptec has a bandwidth of 24 MB/sec, and can deliver 3 MB/sec to each of up to 8 APs. Alexander says that this scheme "eliminates VAX I/O bottlenecks". Numerix has done an interface to Aptec for one customer.

Numerix has a sophisticated interactive debugger for VAX/VMS systems which facilitates debugging microcode for both the main AP and the IP. Because it talks to the AP through the host-resident AP-executive code Alexander expects the debugger to be moderately portable. Alexander was quite interested in our experiences in porting our code to Unix because he is under some pressure to begin to support Unix.

Alexander said that Numerix intends to offer support for multiuser timeshared operation of the MARS-432 within a few months (this will not involve multiple hosts). The idea is that multiple host processes can each have a part of the AP (with memory protection) and can overlap I/O with each other's computing, switching on a time scale of seconds. Just as in the FPS 120b, the basic supervisory control of the MARS-432 resides in the "AP executive" code in the single host.

The MARS-432 has 8-deep FIFOs on the outputs of the arithmetic pipes and the memory. This greatly facilitates the optimization of microcode and certainly must be a vital element in the success of the Fortran compiler for the 432. The latency of both the add and the multiply pipes is either 4 or 5 clock cycles (400 or 500 nsec), depending on how you count (results are available at the input to the FIFO on the 4th cycle).

Numerix has a software tool which they call the "Loom". The idea is that the programmer writes out the microcode as essentially a scalar algorithm, but with knowledge of the optimization process. The Loom then analyzes the code and transforms it to a "rolled-up" loop with startup code. Alexander said that the Numerix programmers make extensive use of the Loom to produce optimum vectorized microcode in complicated algorithms. He showed us a number of assembly listings of the input to the Loom and the output it produced.

Alexander said that although the MARS-432 is not really a 64-bit AP, it can do 64-bit accumulation of partial sums, and that Numerix will soon provide microcode routines which exploit this ability. Double precision accumulation minimizes loss of precision in dot products and other linear algebra algorithms.

In the fall of 1983 we examined the manuals on the vector subroutine library for the 432. We compared them with the corresponding pages in the FPS 120b library manuals line for line, and found that they were identical. That is, the microcode routines in the two libraries have exactly the same names, exactly the same arguments, and exactly the same functional descriptions. The known exceptions are that the integer arguments have 32 bits precision rather than 16 bits, and that the AP open routine has two additional arguments in the 432

library. This high degree of compatibility means that the vector library calls in AIPS tasks are immediately usable with the 432. All that remains to implement AIPS on the 432 is to code the handful of special microcode routines programmed by NRAO personnel (APGRD4, CLNSUB, etc.). The "Loom" should aid in this process.

Alexander indicated that it should be possible for NRAO personnel to visit the Numerix factory in Massachusetts in order to make a trial installation of AIPS for the MARS-432 on a VAX under VMS. This would permit benchmarks, although without obtaining the full potential I/O bandwidth of the AP.

3.0 NUMERIX'S FORTRAN COMPILER FOR THE MARS-432

The language is ANSI standard Fortran 77 with the following exceptions:

1. Permanent Language Restrictions (i.e., things that are unlikely to be fixed):

   - CHARACTER data type and all associated operations and features are not supported. Tasks for the 432 will always exist in two parts, the part in the host and the part in the AP. The CHARACTER restriction obliges the AP part of the code to communicate error messages and other information to the host part of the code using binary codes and values rather than character strings. This restriction is merely a minor nuisance for AIPS tasks.

   - DOUBLE PRECISION data type and all associated operations and features are not supported, because the 432 hardware does not directly support 64-bit data (note that DP sum accumulation of products is possible). Lack of DP might be a slight problem for AIPS in some cases.

   - EXTERNAL, INTRINSIC, ASSIGN, assigned-GO-TO, input/output, FORMAT, and ENTRY statements are not supported. Nor is the alternate RETURN feature.

   It is obvious that Numerix does not intend to build a full Fortran support library for the 432.

2. Temporary Language Restrictions (as of 30 April 84):

   - Identifiers may not include $ or _ (underscore) characters; ! is not recognized as a comment statement identifier.

   - Logical assignments not yet implemented (e.g., A=B.LT.C not supported)

   - EQUIVALENCE statements require array element specification, not just names {e.g., one must say "EQUIVALENCE (A(1),B(1))", rather than "EQUIVALENCE (A,B)"}

- COMPLEX arithmetic expressions are not yet implemented although COMPLEX array declarations are permitted. (This might be a nuisance for AIPS implementation.)

3. Hardware Limitations:
   - probable hardware limit of 15 levels on subroutine calls. This might be a problem for AIPS in a few cases.

4.0 COMPARISON WITH THE STAR TECHNOLOGIES ST-100

Star Technologies and Numerix are direct competitors: the ST-100 is about twice as fast as the MARS-432 for about twice the price. Both APs are conspicuously faster than the previous generation of APs. Alexander asserted that Star has had some reliability problems with their unusual air-cooled ECL design. He said Numerix has used "safe technology" (Schottky TTL plus VLSI and some bit-slice chips). He also said that all designers are now planning to use CMOS in future AP systems (Star told us in May that they are developing a 50 MFlop CMOS AP for OEM customers). Meckley and Alexander say that they recently won a large seismic contract from a customer who had tried the ST-100 (but apparently Star won a large medical imaging contract competition last year). Of course, comments made by one company about another must be independently verified before being given full credibility. It would be interesting to get an analogous appraisal of Numerix from the Star Technologies people. When we asked the Star people in Minneapolis in May about this, they effectively claimed ignorance about Numerix!

Some comparisons are given below (data for the old 120b are also included in some cases for comparison):

- Prices:
  (for minimum acceptable system) ST-100 about $250K; MARS-432 about $125K; AP-120b (now called 5205) about $60K. Both the ST-100 and MARS-432 have 512 KW {kiloword} in this comparison (the minimum MARS-432 has only 64 KW, same as the AP-120b). Additional memory for the MARS-432 currently sells for about $28K/MW {megaword}, or $7K/MB, a price which is quite low compared to 120b memory. A large MARS-432 with 4 MW of memory (enough to hold a 2048-square map) would currently sell for about $200K. Memory prices are sure to decrease with time. [NOTE: above prices are list, before GSA or educational discounts]

- Clock speeds:
  The ST-100 has 25 MHz (40 ns); the MARS-432 has 10 MHz (100 nsec); the AP-120b has 6 Mhz (167 ns).

- Pipes:
  The ST-100 has two adders and two multipliers (4*25= 100 MFlops); the MARS-432 has two adders and one multiplier (3*10= 30 Mflops); the AP-120b has one adder and one multiplier (2*6= 12 MFlops).

- Timings for some common operations:

Function	ST-100 (*)	MARS-432	AP-120b
1K complex FFT	0.86 msec	1.7 msec	6.0 msec
VADD	0.04 microsec	0.2 microsec	0.8 microsec (@)
VDIV	0.6 "	0.5 "	1.7 "
VSQRT	0.56 "	0.8 "	1.8 "
VATAN	1.3 "	1.1 "	9.7 "
VSIN	0.5 "	0.9 "	1.3 "
VDOTP	0.04 "	0.1 "	0.5 "

* The ST-100 times do not include transfers to and from the cache. In practice, the VADD and VDOTP times probably need to be increased to account for this factor. Transfer times can probably be overlapped for the other operations.
@ times are microseconds per point.

The table shows that the theoretical speed advantage of the ST-100 over the MARS-432 (3.3x) is not fully realized in practice. For both FFT algorithms and vector sine/cosine calculations the ST-100 has less than a factor of 2 advantage over the MARS-432, probably due to differences in ability to optimize complicated loops. Note that the 432 actually beats the ST-100 for division and arctangents. Another notable detail is the surprisingly good performance of the 120b on vector sine/cosine. This is due to the table memory of the 120b, a feature which both the ST-100 and MARS-432 lack. On the whole, it appears that a fair summary is that the ST-100 delivers about twice the performance of the MARS-432 for about twice the price. Therefore,

> it follows that <u>two</u> MARS-432 APs attached to a suitable host (e.g. an IBM with enough data channels) would be about equal to <u>one</u> ST-100 in both total price and total performance.

- Memory Heirarchy:
The ST-100 has two kinds of memory, with separate address spaces: main memory (1 GW {gigaword} limit, 8 MW with current chips) and a "data cache" (192 KW limit, but only 48 KW with current chips); the MARS-432 has a single homogeneous address space (16 MW limit, 4 MW with current chips, and with the lowest 128 KW implemented in faster chips); the AP-120b has a small homogeneous space (only 64 KW, with a memory mapped extension to 1 MW).

- Data Formatting:
The ST-100 has a separate microcoded "storage move processor" (SMP) to handle data movement between main memory and cache with pipelined formatting; the MARS-432 depends on its floating adder pipelines (64 opcodes) to perform these operations. The power of the SMP is a major strength of the ST-100. It is also the major weakness of the design, because it puts a barrier between the arithmetic pipes and the main memory, thereby making data dependent addressing operations in main memory difficult.

- Resident CPUs:
  The ST-100 has two resident Motorola 68000s; neither the MARS-432 nor the AP-120b has such a CPU.

- Compilers:
  The ST-100 has a compiler for the 68000 which controls it, but not for the microcode of the SMP and its arithmetic pipelines; The MARS-432 has a Fortran compiler and linker for its microcode engine. Alexander says that compiled random scalar Fortran code typically runs at 3-5 MFlops, and a conservative statement is that it always delivers at least 1 MFlop. This scalar compiled code ability is a striking advantage for the MARS-432. In general, it seems to be fair to say that the MARS-432 is a much better scalar machine than is the ST-100.

- Disk Drives:
  The ST-100 will have high performance disks (CDC "Hydra" drives) by sometime in 85 or 86; the MARS-432 might have such disks (probably "IBIS" drives) eventually--- development of its disk drive capability will depend on customer interest. Both the CDC and the IBIS drives are expected to deliver about 12 MB/sec. Numerix displays less interest than Star in building the necessary operating system software for support of disk systems, but Alexander said he was willing to consider the idea if we really want it. Note: an AIPS implementation probably doesn't need very much, if any, vendor software support to implement disks because of the Z-routine interface. Alexander showed some interest in the idea of the host and the AP sharing access to a disk (watch out for different FP number formats!).

- Multihost/Multuser Capability:
  The ST-100 supports multiple hosts (and multiple users in each host?); the MARS-432 is expected to support multiple users in a single host soon. The multiple host capability of the ST-100 is certainly a plus for it.

It is difficult to draw a final, firm conclusion about the relative merits of the ST-100 and MARS-432 from the details of the comparison above, and especially difficult to draw a conclusion appropriate for all applications. But for the case of implementing AIPS in a standard host-plus-AP context, the fact that the basic architecture of the MARS-432 is so similar to that of the 120b (homogeneous address space, single totally synchronous microcode engine, easy data dependent addressing) means that the MARS-432 has a special advantage. The dual memory scheme and SMP of the ST-100 amount to a major conceptual and technical barrier which inhibits easy porting of 120b applications and 120b programming experience to the ST-100. On the whole, the small cache memory and difficulty of accomplishing data dependent addressing, which depends on coordination of two processors (SMP and arithmetic processor), probably outweigh the advantages which the ST-100 should have with its higher clock speed and the powerful SMP. Also, it appears that the compiler capability for the Numerix is substantially superior to that for the ST-100. The availability and performance of disks for either or both of these APs will be important in future evaluations of these APs for our applications.