Under the Hood of a Superchip: The NOVIX FORTH Engine

Charles H. Moore
Consultant

Robert W. Murphy
Design Engineer

NOVIX
10590 N. Tantau Avenue
Cupertino, CA 95014

Abstract

The NC4000P is a single chip FORTH Engine based upon minimum hardware concepts developed by Mr. Charles H. Moore. This highly parallel machine architecture directly executes FORTH primitives in a single clock cycle. The initial implementation of this device is based upon a 4000 gate CMOS semi-custom integrated circuit operating at an 8 MHz clock rate.

Background

The NOVIX FORTH Engine project was formally initiated in March of 1984 to develop a prototype printed circuit board to validate FORTH engine concepts. Upon closer examination of the architecture, it became apparent that a CMOS gate array would offer the proper context for validation. The prototype component resulting from this initial effort is the NC4000P FORTH Engine that is detailed herein.

Introduction

The NOVIX FORTH Engine is composed of a minimum number of hardware elements organized within the constraints of a two stack architecture to support the direct execution of FORTH. Figure 1 depicts the organization of the NC4000P microprocessor at the register transfer level. The following definitions describe the function of each element within the NC4000P.

A The 16-bit address multiplexer (A) points to the appropriate main memory locations to permit instruction fetches, data fetches and data stores.

M The 16-bit main memory data port (M) passes data between internal registers and main memory.

L The instruction register (L) determines the action of the processor from the information contained within the 16-bit word latched as a result of an instruction fetch.

P The 15-bit program counter (P) maintains the address of the next instruction to be fetched form main memory.

I The 16-bit index register (I) is the top of the return stack.

R The 16-bit return stack data port (R) passes data to and from return stack memory.

J The 8-bit return stack address counter (J) points to the appropriate location within return stack memory.

T The 16-bit ALU/register (T) is the top of the parameter stack.

N The 16-bit next register (N) is the first parameter under the top (T) of the parameter stack.

S The 16-bit parameter stack data port (S) passes data to and from parameter stack memory.

K The 8-bit parameter stack address counter (K) points to the appropriate location within parameter stack memory.

All registers (I,J,K,L,P and T) described above are positive edge-triggered D-type structures with multiple input sources as shown in Figure 1. The ALU functions required to execute single precision arithmetic and logical operations are a function of inputs to the T register. The specific design details of each processor element will not be addressed herein as such details are of minor importance compared to the improvements in FORTH system performance afforded by parallel data flow on the chip.

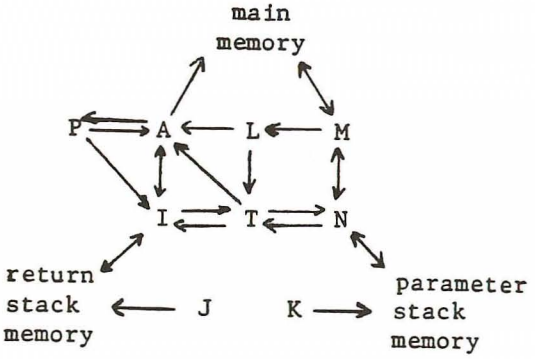

Figure 1

Instruction Execution

The parallel structure of the NC4000P permits the update of any or all internal registers and one or both stack memories during a single clock cycle. The instruction set of the NC4000P is composed of two types; namely, single-cycle zero address operations and two-cycle memory reference instructions. Single-cycle instructions force P or L onto A to fetch the next instruction in sequence (P) or branch (L). Two-cycle instructions force T or L onto A during the first cycle to accomplish a fetch or store, while executing a single-cycle operation during the last cycle of the instruction.

The concurrent character of the NC4000P is best described through examples of single-cycle data flow. The following single-cycle instructions are presented to convey the intent of the NC4000P design.

:	Force the embedded address latched in L onto A to fetch the first instruction of a FORTH word, push P onto I to maintain a return address, latch A + 1 into P to maintain a sequential reference to the next instruction of the called word.
;	Force the return address latched in I onto A to exit a FORTH word, pop the return stack into I, and latch A + 1 into P.
SWAP	Fetch the next instruction (P to A), pass N to T, pass T to N, and latch A + 1 into P.
OVER	Fetch the next instruction (P to A), pass N to T, pass T to N, push N onto the data stack memory, and latch A + 1 into P.
SWAP;	Fetch the next instruction from the return address (I to A), pass N to T, pass T to N, and latch A + 1 into P.
+	Fetch the next instruction (P to A), add N into T, pop the data stack into N, and latch A + 1 into P.
OVER + ;	Fetch the next instruction from the return address (I to A), add N into T, and latch A + 1 into P.

Through the above examples, it is obvious that the instruction decode performs two basic functions; enabling the specific input to a register and then resolving the question of latching that input as required to accomplish the desired result.

The last example above (OVER + ;) is representative of the incremental nature of the NC4000P, as often a grouping of FORTH words can be incrementally compiled into a single-cycle instruction. This incremental compression of FORTH source leads to a machine language that requires fewer instructions be assembled than FORTH source words necessary to describe an algorithm. NC4000P incremental compression of FORTH source represents a significant breakthrough in FORTH systems as the instruction ratio of source code to machine code is now less than one. (Typical instruction ratios for conventional systems can exceed 20 to 1). Incremental compression is a direct consequence of one specific design goal; namely reduce the overhead associated with calling a FORTH word to one cycle if possible.

Therefore, a new measure of FORTH engine performance can be defined in "automotive" terms. (Look out A.J. Foyt)

COMPRESSION RATIO: The ratio of FORTH source words to machine clock cycles required to implement a function.

Given this definition, compilers for the NC4000P can incrementally compile machine instructions by simple "look-back" techniques to compress whenever it is possible.

Summary

The NC4000P is a prototype engine presently operating at a machine instruction rate of 8 MIPS while the basic instruction set has a compression ratio greater than one (single instruction compression ratios range from 1 to 5 at present). This source to machine code compression permits FORTH language execution at averaged rates approaching 24 MIPS.