several from the DARPA Image Understanding Benchmark suite. Discusses the performance of Abacus on these algorithms.

Chapter 7: Architectural Tradeoffs. Evaluates architectural tradeoffs such as local memory size, ALU width, off-chip memory bandwidth and network bandwidth based on instruction traces from the parallel algorithms.

Chapter 8: The Next Generation: Abacus-2. Describes a set of modifications that allow a redesigned Abacus element to function in multiple-SIMD mode, as a systolic processor, or to emulate random logic circuits effectively.

Chapter 9: Conclusions. The contents of this chapter will depend (somewhat) on the results of the actual research.

Schedule

Task	Time	Start	End
Test board fabrication	21	Sep 11	Oct 1
Test board assembly	3	Oct 1	Oct 4
Test board/C30 interface	3	Oct 5	Oct 7
Test vector design	5	Oct 8	Oct 12
Chip testing	7	Oct 13	Oct 19
Area Exam	30	Oct 15	Nov 15
Simulator coding	14	Nov 17	Dec 1
Architectural analysis	21	Dec 1	Dec 22
Text writing	30	Dec 22	Jan 12
Thesis defense	7	Jan 13	Jan 20

1 Introduction and Motivation

Computer vision, the understanding of images by a computer, is both an exciting component of artificial intelligence and a daunthing computational task considered as one of the Grand Challenge problems of parallel processing. The design of a parallel computer for vision is an attractive problem for a computer architect, as the problem domain lends itself to optimizations that promise orders of magnitude performance over general-purpose parallel computer.

This thesis proposes to design, implement, and analyze a VLSI-based architecture tailored to the computing demands of early vision. Although the field has received a great deal of academic attention, this work targets a considerably higher speed design than previous efforts, and as a result encounter system-level consequences of a high speed parallel computer. In the body of this section we will the describe the problem domain and the detailed thesis approach.