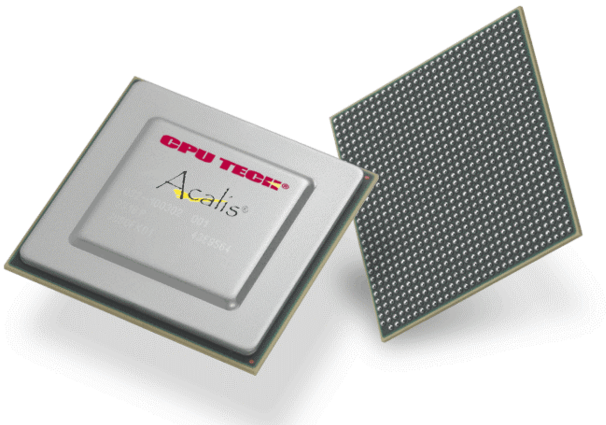

High-Performance, Parallel Embedded Architectures Using Acalis® CPU872 PowerPC® Multicore

John Swensen, Gail Walters
CPU Technology, Inc.

CONTENTS

- **Acalis Overview**
- **MPI Benchmark Results**
- **Parallel Granularity Benchmark Results**
- **Pipelined Computation Benchmark Results**
- **Summary of Performance and Features**

Acalis Overview

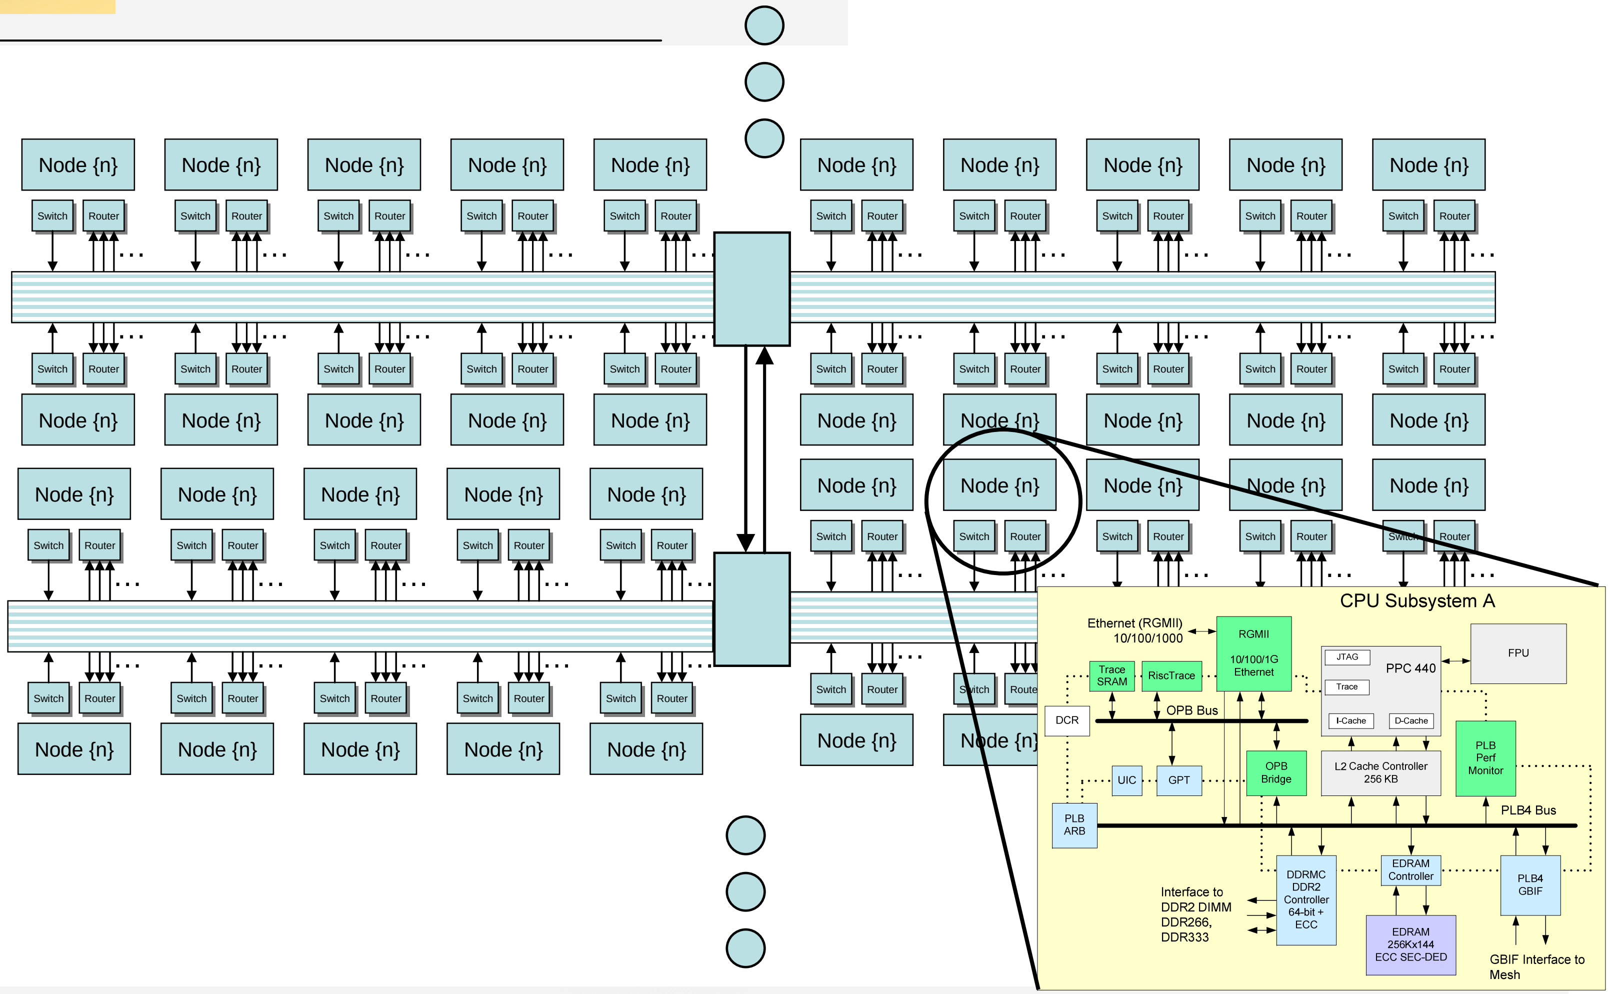

Acalis CPU872 Overview

- **Secure Processor SoC containing 2 complete PPC compute nodes**
- **AT protection**
- **Supports Red/Black processing**
- **Protects software without degrading performance**
- **Highly efficient (MIPS/Watt)**
- **Balanced scalable performance**
  (security, throughput, memory, I/O, communications, synchronization)
- **Scalable without additional devices**
- **Industrial grade with minimum 10-year lifespan**

Acalis CPU872 Features & Specifications

Security Functions	Scalable and programmable key structures Programmable On-chip Firewall Anti-Tamper, Anti-Reverse Engineering
Performance	Up to 3.6G Operations (DMIPS, FLOPS)
Typical Power	8 Watts
Process Technology	90 nm CMOS (IBM Trusted Foundry)
General Purpose Processors	Dual PowerPC® 440 Processor Cores Dual PowerPC Floating Point Units
L1 Cache	32KB data; 32KB instruction per PPC440 core
L2 Cache	256KB per PPC440 core
Embedded DRAM	4MB with ECC per PPC440 core
Communication Processors	Dual Quintillium™ MPI Processors
Streaming Processors	Dual Streaming Processors – 1024 Transactions
Package	31 x 31 mm, 899-pin BGA package

Acalis CPU872 Features & Specifications

Inter-processor Interface	Five 10 Gbs Express interfaces
Ethernet Interface	10/100/1000 Ethernet controller
Serial Protocol Interfaces	Secure Boot Interface and I2C
Memory Controller to External DRAM	Dedicated 32/64-bit DDR2 with ECC per PPC440 Core
DMA Controller	Dual Multichannel with 64-bit addressing
Multiprocessor Interrupt Controller	One Universal Controller per PPC440 Core
Parallel Barrier Synchronization	Scalable to massive arrays
Mesh Router	64-bit Addressing
Programmable I/O Interface & Timers	Industry Standard Interface
Software and Tools	Compatible with PowerPC-based operating environments
Availability	Samples Now to Qualified Customers Production 4Q08

Multi-Processor Benchmark Results

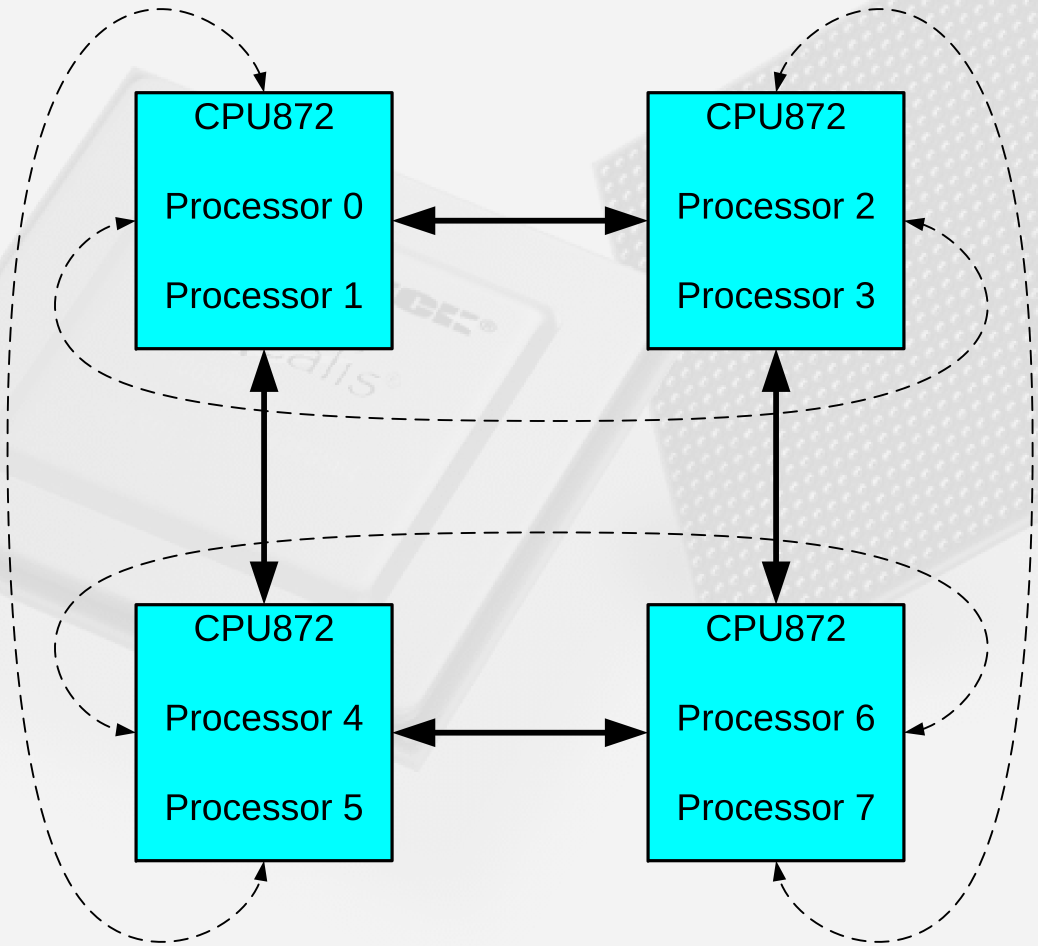

Multi-Processor MPI Transfers

- Transfer *n* bytes to processors on same chip, one-hop distance, and two-hop distance away

```
Start Timer
do {
        MPI_Send to X
        MPI_Recv from X
} 100 times
Stop Timer
```

- Elapsed time/200 = average one-way transfer time

Multi-Processor MPI Transfers

Transfer Distance	Transfer Latency		Transfer Rate
	8 Bytes	40 Bytes	1024 Bytes
On-chip	2.8 µs	3.8 µs	81.9 MBytes/sec
One-hop	3.9 µs	6.3 µs	68.9 MBytes/sec
Two-hop	4.8 µs	8.7 µs	58.9 MBytes/sec

Multi-Processor MPI Transfers

- **Compare Latency*Power, Bandwidth/Power with other multiprocessors[1]:**
  - **Acalis: 667 MHz PPC440, 8 W**
  - **MCR: 2.4 GHz Pentium Xeon, 65 W**
  - **ALC: 2.4 GHz Pentium Xeon, 65 W**
  - **Thunder: 1.4 GHz Itanium Tiger 4, 105 W**
  - **BlueGene/L: 700 MHz PPC440, 12 W (est.)**

[1] **(MPI at LLNL: MPI Performance Measurements, *https://computing.llnl.gov/mpi/mpi_benchmarks.html*)**

Multi-Processor MPI Transfers

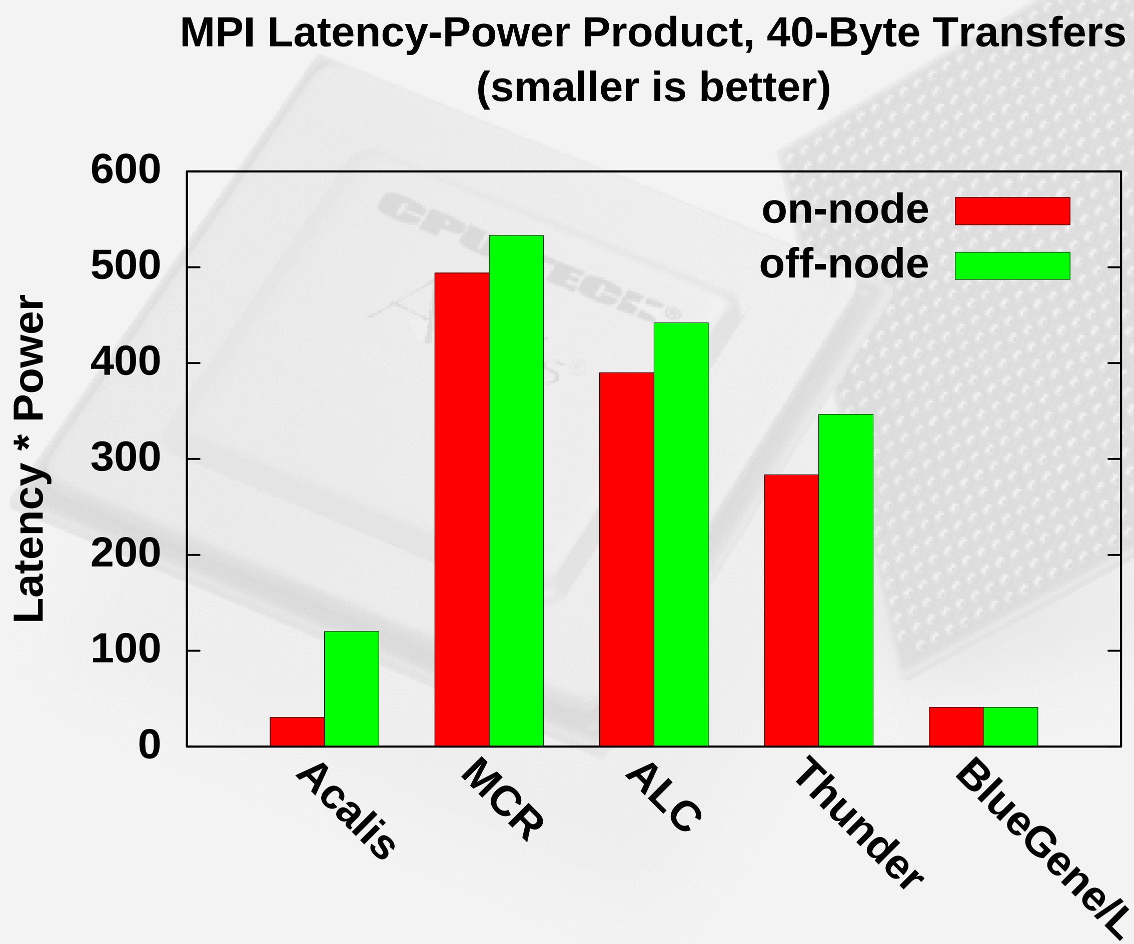

Multi-Processor MPI Transfers

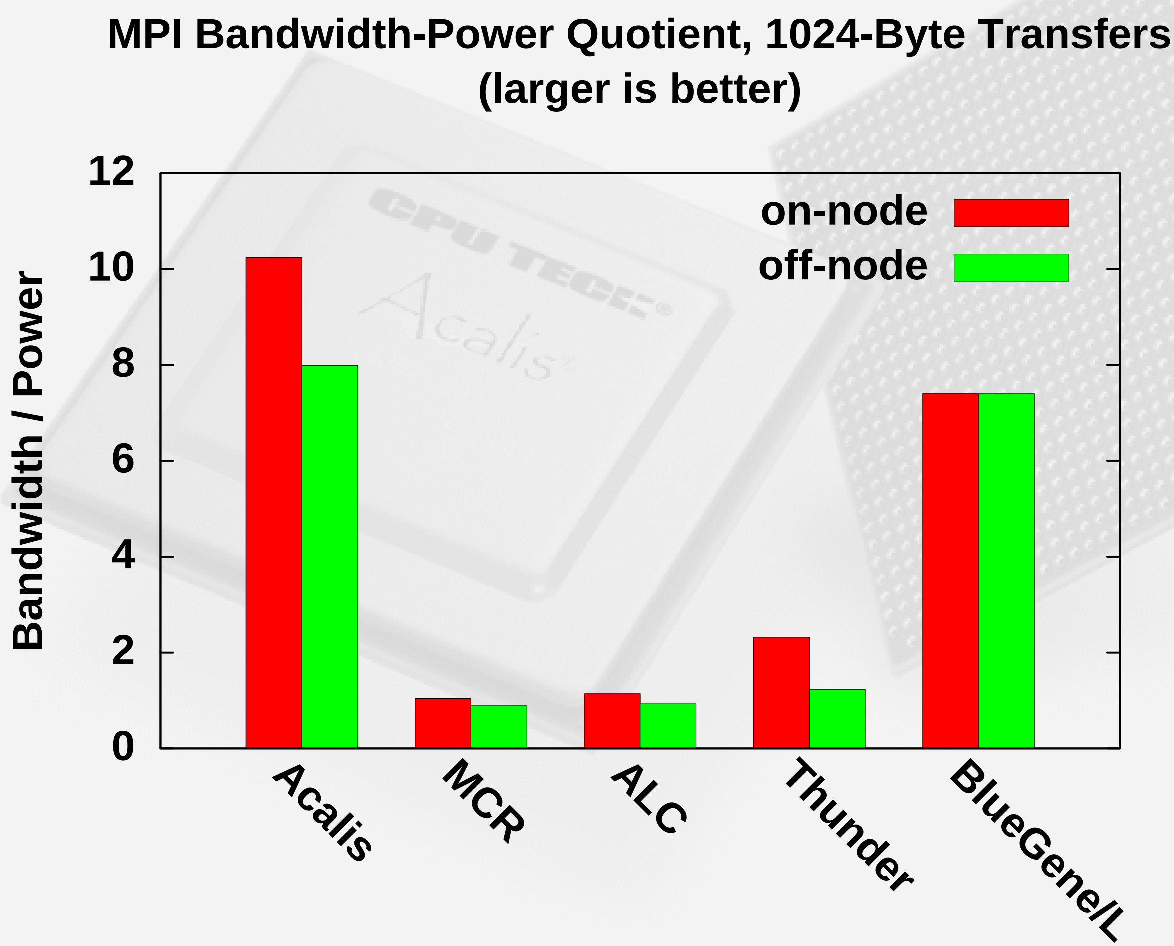

Multi-Processor Token Passing

- **Single token passed around ring of processors, *n* times**
- **Each processor reads the token, increments it, passes token to next processor in the ring**
- **Measures finest practical parallel granularity.**

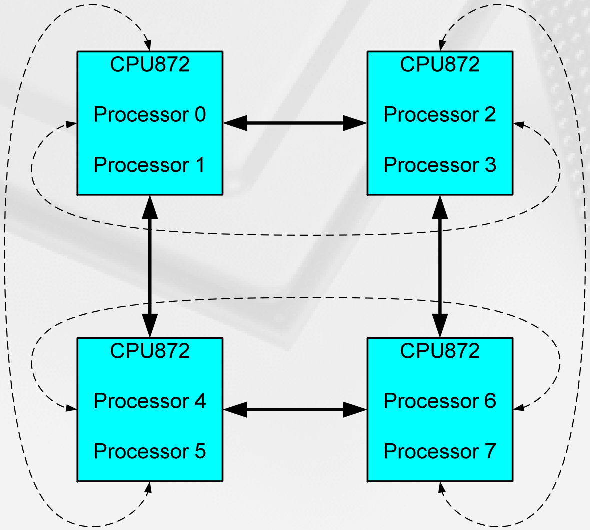

Multi-Processor Token-Passing Code

```
StreamingRegister  * nextProcToken, * localToken;
long * localTokenValid, * nextTokenValid;
if (processor == 0) {
   *localToken = 0; *localTokenValid = 1;
} else *localTokenValid = 0;
forall processors {
   for (i = 0; i < n; i++) {
#ifdef MEMORY_TOKEN
      while (!*localTokenValid)  ; *localTokenValid = 0;
       *nextProcToken = *localToken + 1;  *nextTokenValid = 1;
#else
      *nextProcToken = *localToken + 1;
#endif
   }
}
```

Token Passing Timing

- **667 MHz Clock, mix of on-chip and one-hop paths**
- **Memory-based token data and flag:**
  - **552 ns to increment, pass a token (average)**
  - **Parallel grain size is 368 cycles of work (350-700 instruction-execution times)**
- **Streaming Register-based token:**
  - **350 ns to increment, pass a token (average)**
  - **Parallel grain size is 233 cycles of work (220-450 instruction-execution times)**
  - **Communicate with other processors up to 2.8 million times/second without penalty**

Pipelined Smith-Waterman Algorithm

- **Smith-Waterman algorithm finds the best match of one string to another, assuming that characters may have been inserted, deleted, or substituted.**
- **Smith-Waterman algorithm can be applied to large classes of problems involving imperfect matching, including:**
  - **DNA similarity analysis, using an alphabet of 4 bases (Adenine, Cytosine, Guanine, Thymine)**
  - **Protein similarity analysis, using an alphabet of 20 amino acids**
  - **Plagiarism detection, using an alphabet of English words**

Pipelined Smith-Waterman Algorithm

- **Example: matching DNA string GCTGATATAGCT to DNA string GGGTGATTAGCT**
- **Similarity matrix defines scores for pairs of characters (eg. T-T = 1, C-G = -1)**
- **Penalties assigned for gaps (insertions/deletions)**

```
-  G  C  T  G  A  T  A  T  A  G  C  T
   |     |  |  |  |     |  |  |  |  |
G  G  G  T  G  A  T  -  T  A  G  C  T
```

Pipelined Smith-Waterman Algorithm

- **Smith-Waterman algorithm constructs a matrix of global similarity scores**
- **Algorithm run-time is quadratic in the lengths of the sequences**
- **Smallest human chromosome has over 1 million base-pairs**
- **Short DNA sequences (over 4000) each have 20-25 base-pairs**

Pipelined Smith–Waterman Algorithm

- **Large chromosome data distributed across all available processors**
- **Processors connected in a linear chain**
- **Short DNA sequences, similarity array data, best score data pipeline through chain of processors**
- **Program organized to allow load-time mapping of problem to available processors**

Pipelined Smith–Waterman Results

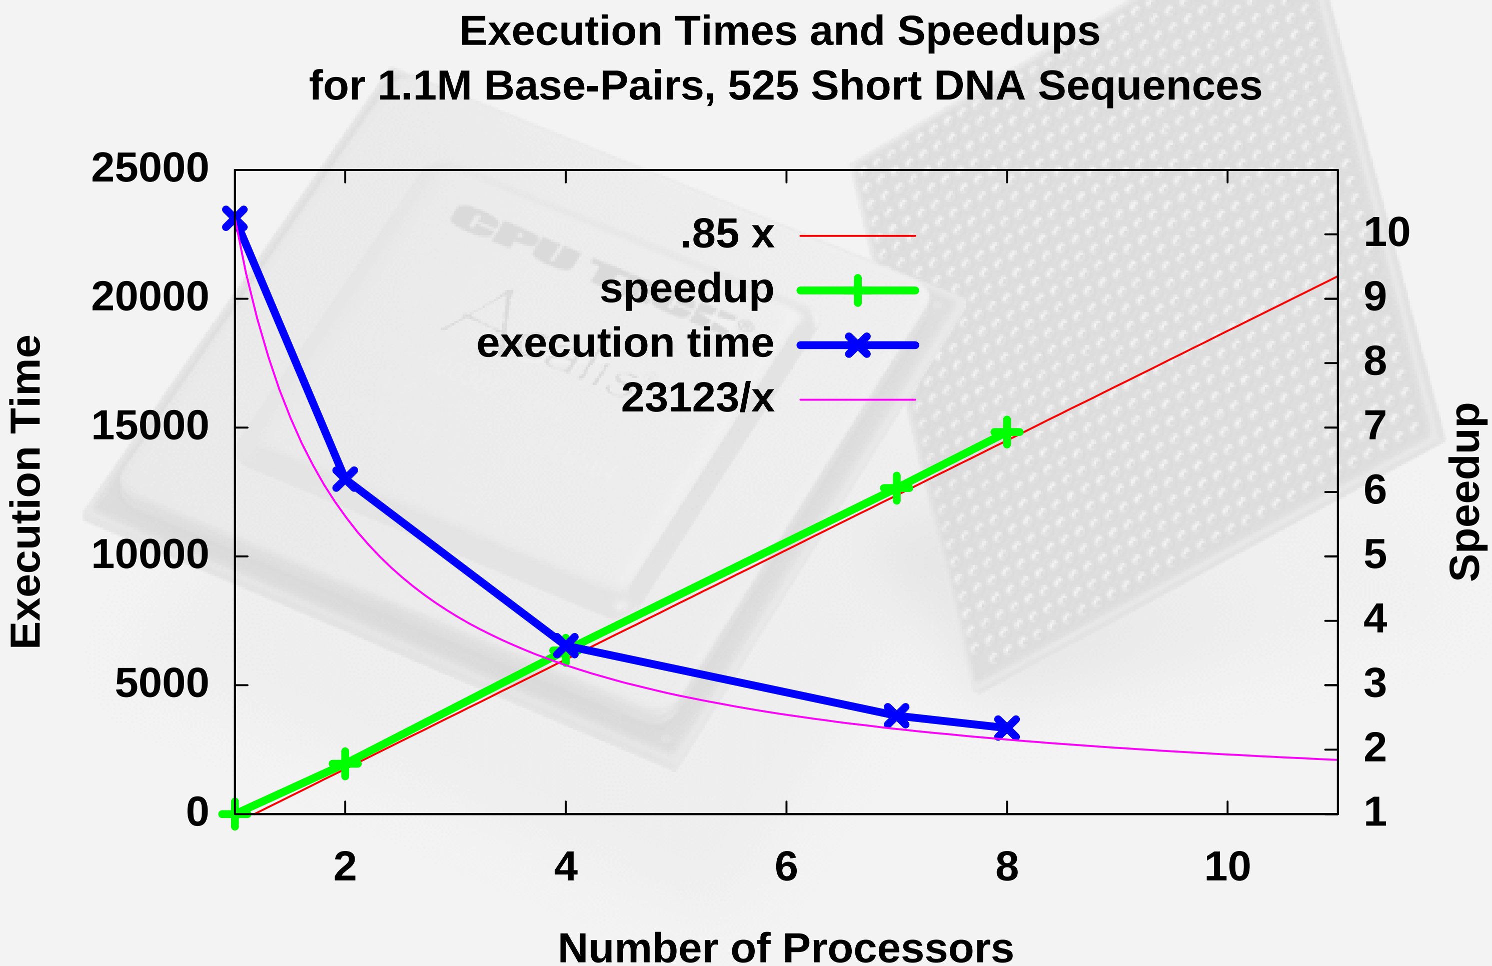

Summary

- **Acalis MPI hardware supports low latency-power product (< 120 us-W) and high bandwidth-power quotient (> 8 MB/sec/W) for inter-processor communication**
- **Acalis Streaming Registers provide low-latency communication to support medium-grained parallelism (< 500 instructions between synchronization)**
- **High-speed interconnects, on-chip and off-chip memories allow pipelined and parallel applications to run with minimal overhead (< 15%)**

Summary

- **High-performance secure processing COTS/NDI device for use in vital applications**
- **Industry-Leading Power / Performance Efficiency**
- **On-chip hardware Firewall, Anti-Tamper, and Anti-Reverse Engineering protects IP and sensitive data, without impacting performance**
- **Readily scalable from single chips to large arrays to meet any class of computing needs without additional logic**

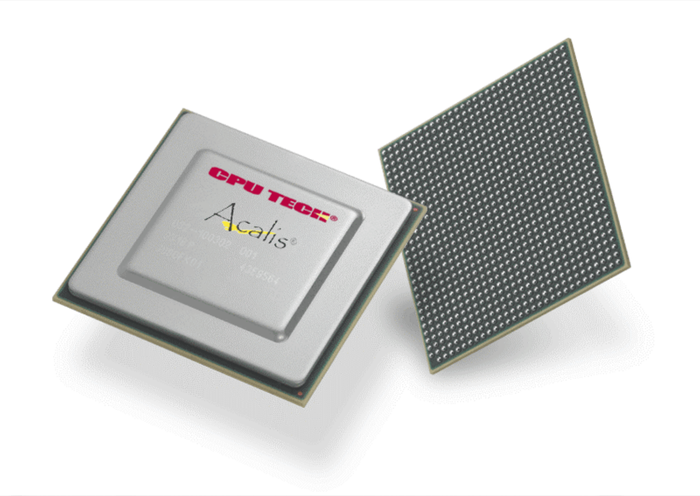

John Swensen
swensen@cputech.com

Gail Walters
g.walters@cputech.com

Acknowledgements
This work was accomplished under the sponsorship of the U.S. Navy.