PART 2

INTRODUCING THE T9000 TRANSPUTER

inmos

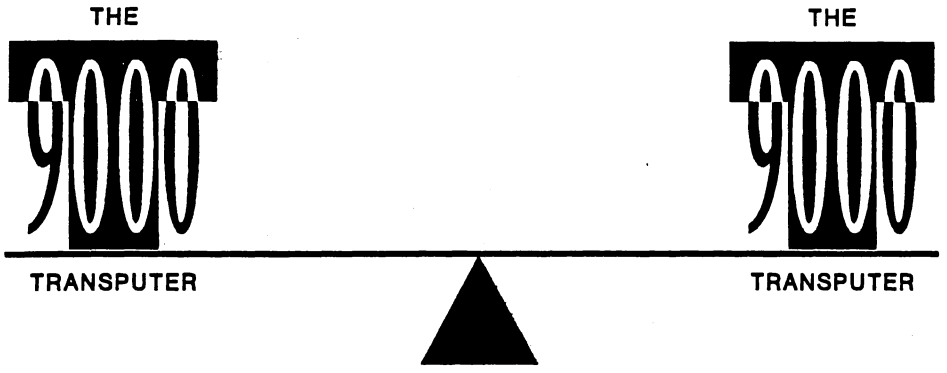

THE RIGHT BALANCE

SGS-THOMSON MICROELECTRONICS

inmos

INTRODUCING THE T9000 TRANSPUTER FAMILY

SGS-THOMSON MICROELECTRONICS

inmos

INTRODUCING THE T9000 TRANSPUTER FAMILY

Paul Strzelecki

SGS-THOMSON MICROELECTRONICS

inmos

"The features that now sell a computer product, such as the amount of memory or the speed of the processing unit, will become secondary to whether or not the product can share data and communicate with other equipment."

S.T. McClellan:

The Coming Computer Industry Shakeout: Winners, Losers and Survivors

SGS-THOMSON MICROELECTRONICS

inmos

Traditional Microprocessors (RISC & CISC)

Out of Balance

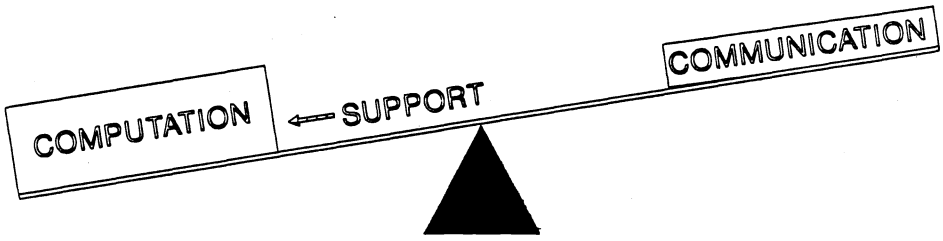

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

The Fundamental Transputer Balance

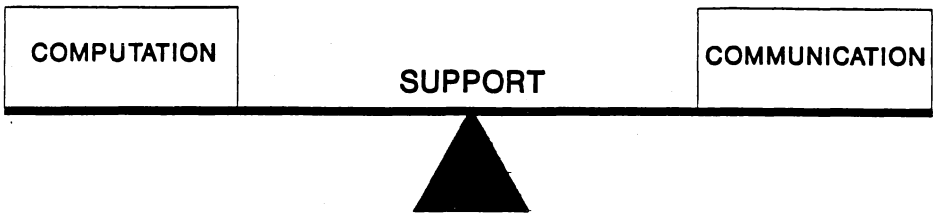

- High-performance, balanced and COMPATIBLE microprocessors, vital for embedded systems
- Balanced for MULTI-TASKING, multi-processing and parallel applications
- CONCURRENT computation, input and output

SGS-THOMSON MICROELECTRONICS

inmos

The Right Scalable Balance

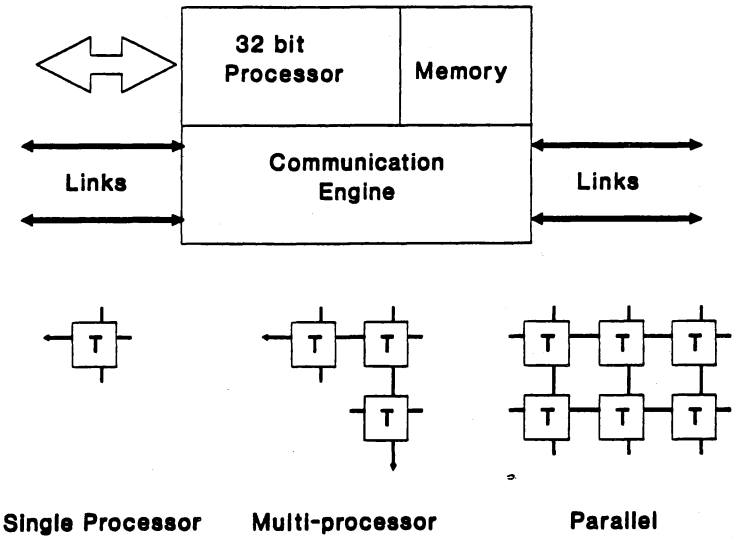

SGS-THOMSON MICROELECTRONICS

inmos

Complete Transputer Product Support for Embedded Systems

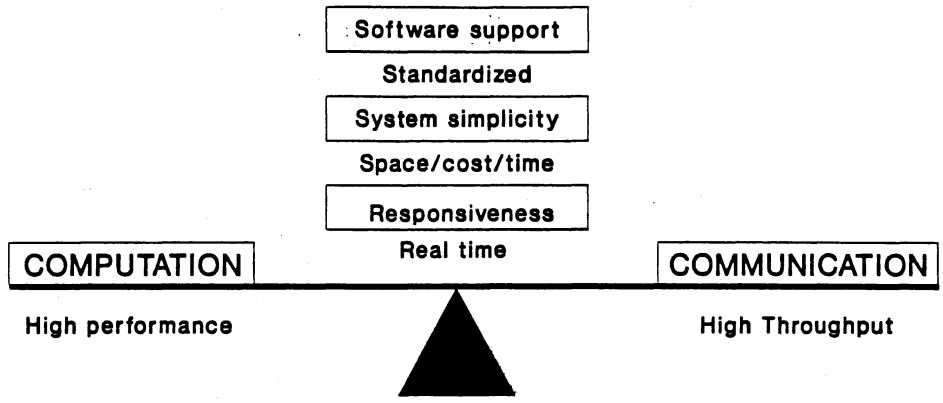

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Transputer Software Support

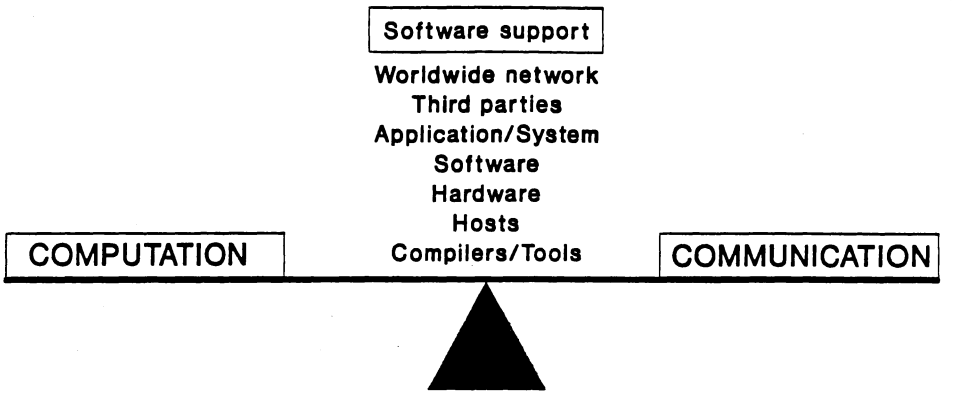

STANDARDIZATION in balance with INNOVATION

SGS-THOMSON MICROELECTRONICS

inmos

The Transputer Family

- Launched in 1985
- Range of 16/32-bit and floating point products, T2xx, T4xx, T8xx including military specification
- Low cost T400 (2$/MIP) introduced in 1989
- Variant strategy - application/customer driven
  Porting link technology to semi-custom library
- Outsold all 32-bit RISC microprocessors in 1990
- "Transputing 91" - California, April 22-25
  Worldwide user conference

SGS-THOMSON MICROELECTRONICS

inmos

32-bit Microprocessor Market

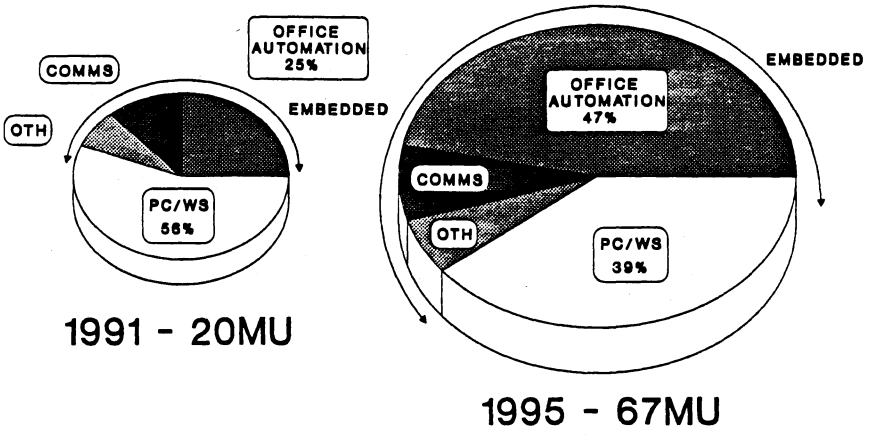

Distributed Computing + Smart Peripherals
- Fuels requirement for communications.

Short design-cycle times demand high-level languages
- Drives 32-bit microprocessor growth.

SOURCE: DATAQUEST + WSTS

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

EMBEDDED SYSTEMS

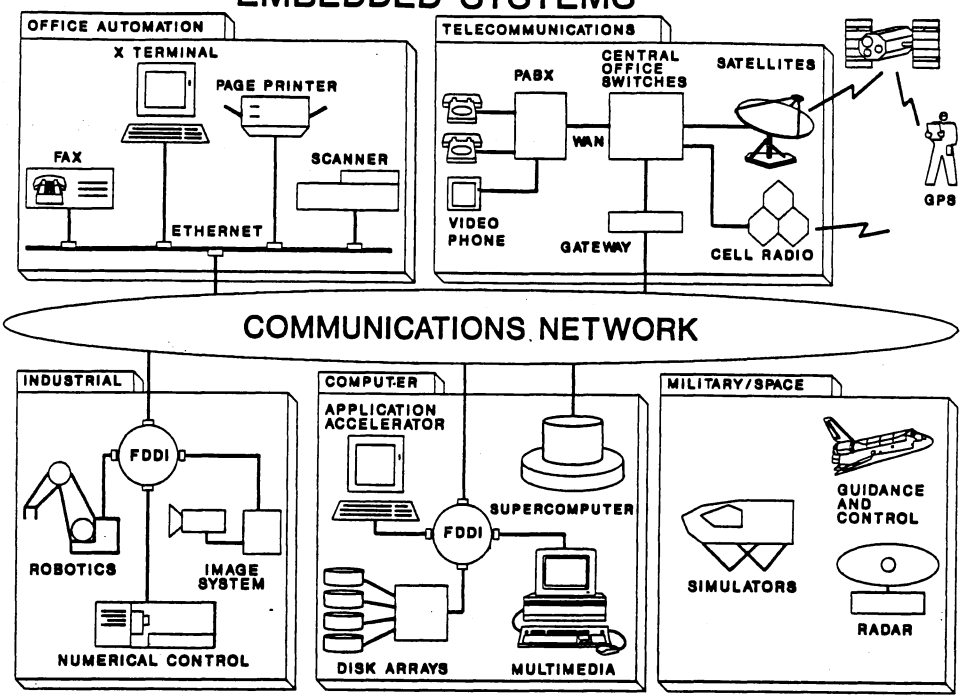

SGS-THOMSON MICROELECTRONICS

inmos

EMBEDDED SYSTEMS

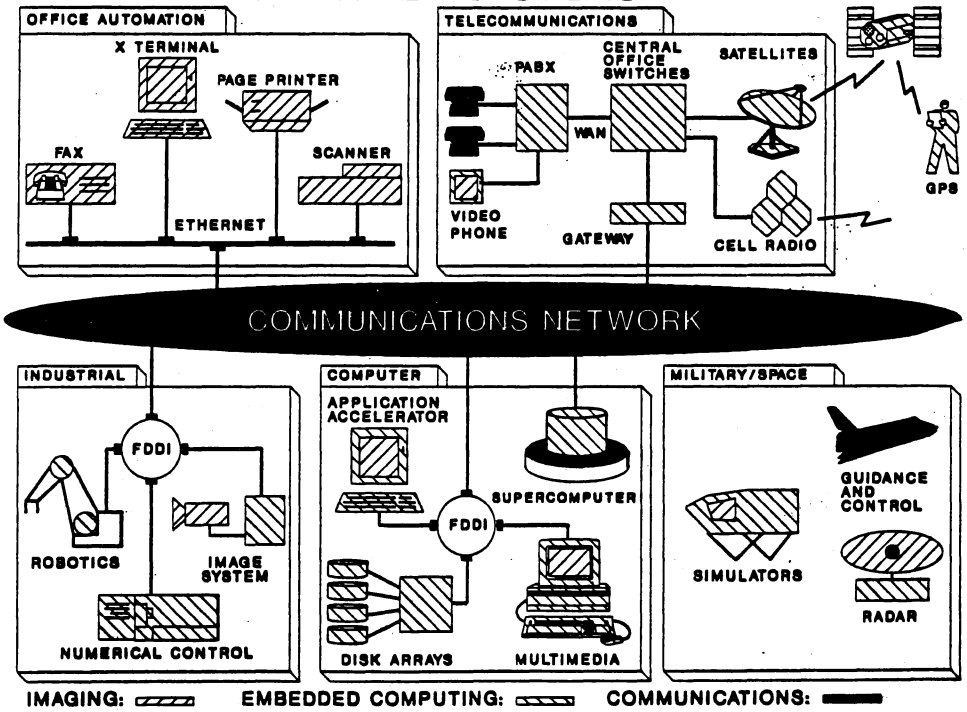

IMAGING: EMBEDDED COMPUTING: COMMUNICATIONS:

SGS-THOMSON MICROELECTRONICS

inmos

Trends in Imaging

Capture/generation, manipulation and transmission of image data

Trends:

- Higher image resolutions eg HDTV, PC graphics
- More color definition eg graphic workstations
- Real-time image processing
  eg security, special effects, defense
- Merging of video and graphics eg MULTIMEDIA

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Trends in Communications

Internetworking, Interfacing, Switching

Trends:

- Migration to digital telecommunication systems
  eg 64Kb/s → 155Mb/s ISDN
- Increase in LAN connections and hierarchy
  eg Ethernet to FDDI
- Digital radio comms requiring portability
  eg mobile cellular, PCN, Satellite and GPS
- Integrity and security of data - fault tolerance

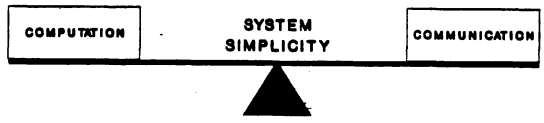

SGS-THOMSON MICROELECTRONICS

inmos

Trends in Embedded Computing

Multi-Tasking, High-Performance
Application Accelerators

Trends:

- Increased CPU performance
- Distributed computing (via LANs)
- Multiprocessing (intelligent peripherals)
- Parallel processing (databases, supercomputers)
- Distributed operating systems

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Maintaining The Balance

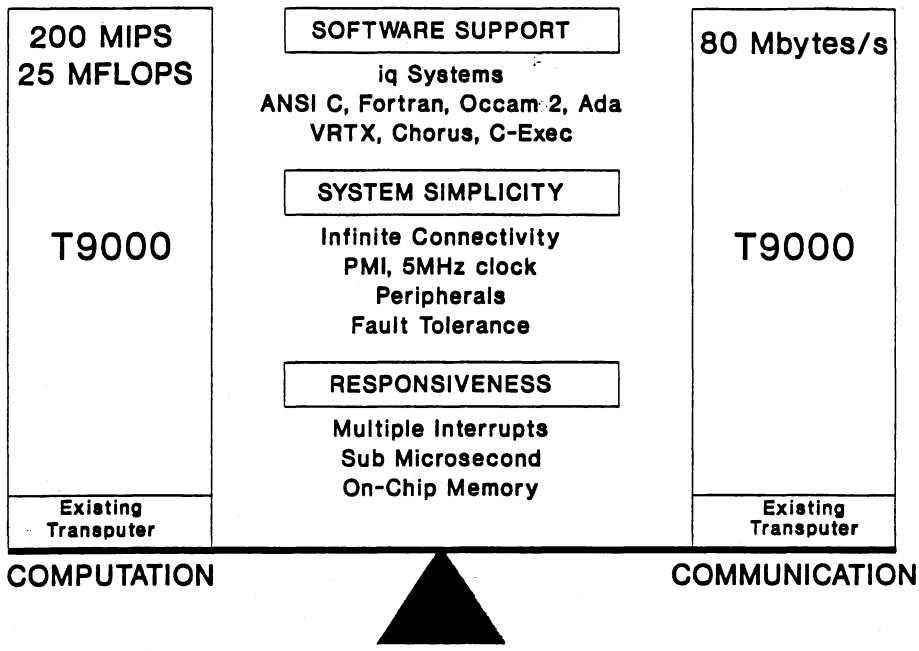

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Maintaining The Balance

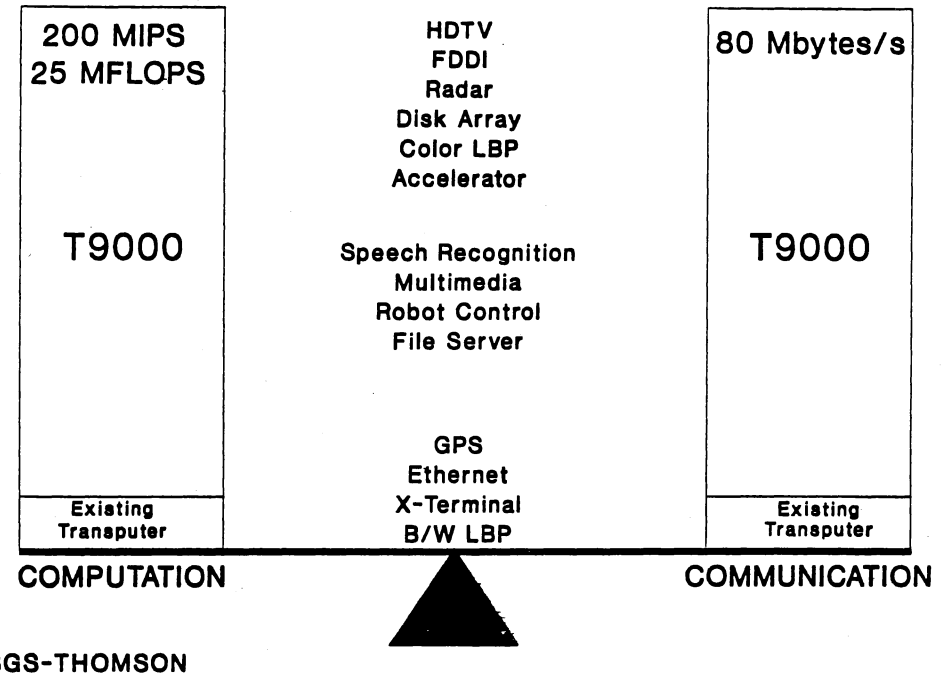

SGS-THOMSON MICROELECTRONICS

inmos

T9000 - Architecture

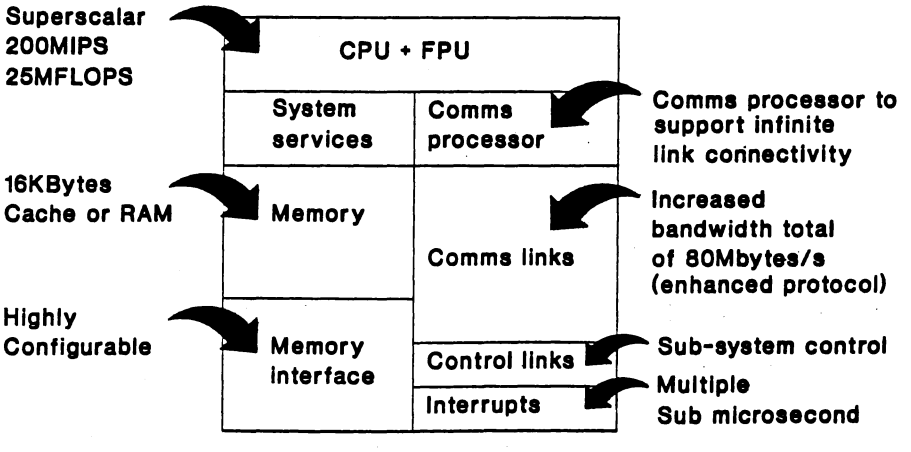

Order of magnitude increase in performance maintaining critical compute/communications balance

SGS-THOMSON MICROELECTRONICS

inmos

Single Chip to Massively Parallel Solutions

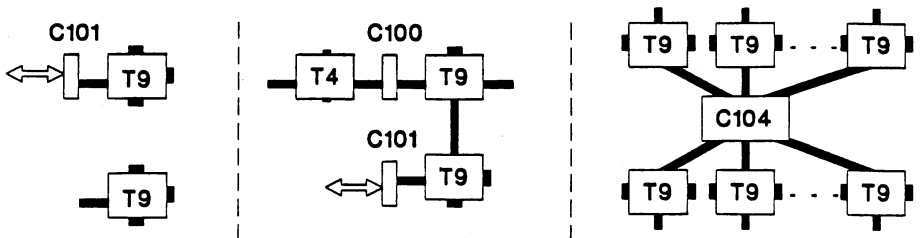

Uni-processor (multi-tasking)	Multi-processor (distributed)	Massively Parallel (data)
X-terminal	Laser printer	Supercomputer
GPS	Robotics	Radar/sonar
	FDDI i/f	Image Processing

C101 link adaptor

C100 link convertor

C104 packet routing switch

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Imaging
Real-Time Video Processing

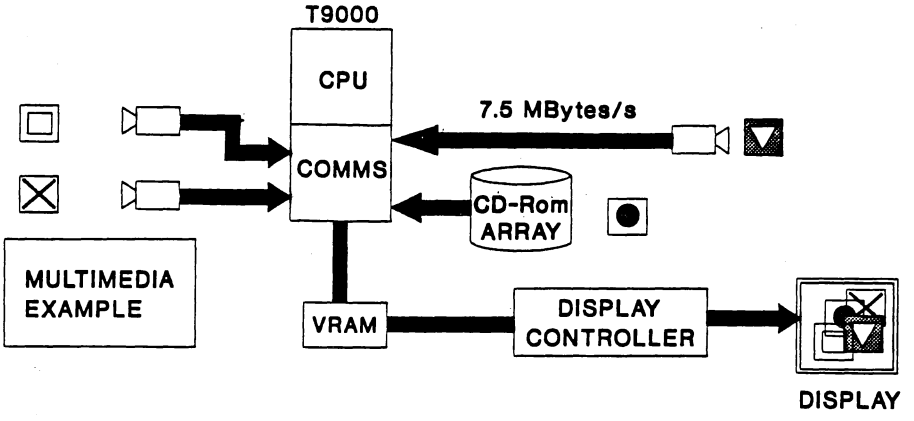

- 200 MIPS and 25 MFLOPS still available for concurrent computation eg. edge detection
- Applications: Multimedia, medical imaging, industrial inspection, security systems, video mixing units and HDTV

SGS-THOMSON MICROELECTRONICS

inmos

Communications

Internetworking eg FDDI to Ethernet Bridge

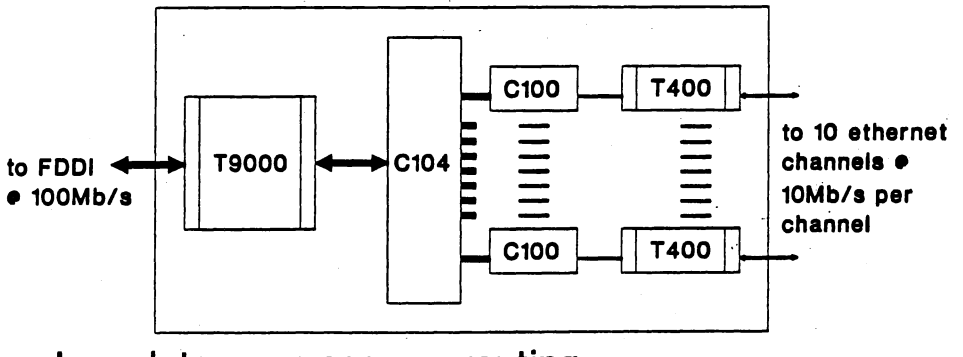

- Low-latency message routing
  - avoids multiple buffers
- Independent operation of communication links
- High compute performance
  - available for protocol conversion

SGS-THOMSON MICROELECTRONICS

inmos

T9000 - Balanced Performance

NETWORKING EXAMPLES

SUSTAINED INSTR	CYCLES (50MHz)	SERV TIME (microseconds)	INT. RESP	FRAME TIME	FRAME /SEC	• CHAN /LINK	DATARATE Mb/s	
~570	285	5.7	<1	6.7	150K	1	80/100	FDDI
~410	205	4.1	<1	5.1	140K	10	10	E'NET
~21000	10550	211	<1	212	5K	~36	2	ATM
~170	85	1.7	<1	2.7	370K	1 (2 LINKS)	155	(53 BYTE CELL)
~624900	312450	6249	<1	6250	160	>1000	.064	ISDN BASE RATE
~5800	2900	58	<1	59	17	50	2	PRIMARY RATE
COMPUTATION		RESPONSIVENESS				COMMUNICATION		

• 4 LINKS PER T9000

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Computing

Application Accelerator

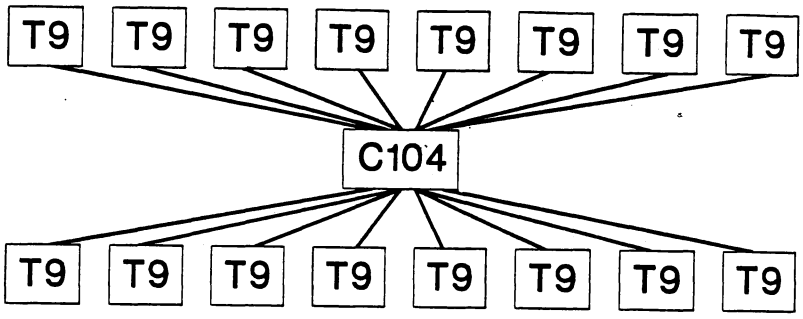

- 16 fully connected T9000s ▪ 400 MFLOPS, 3,200 MIPS
- Footprint of workstation (25cm x 25cm)
- Desktop supercomputer performance
  eg financial numerical simulation, transaction processing

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Family Timescales

	Q391	Q491	Q192
T9000	1st Silicon		Market Availability
C100/104	1st Silicon		Market Availability
C101		1st Silicon	Market Availability
Simulator		Availability	
Development Tools (Existing)	Available Today		
Optimizing ANSI C, Fortran 77, occam 2	Beta		Product Availability

SGS-THOMSON MICROELECTRONICS

inmos

Summary

- 32 bit embedded market - rapid growth
  47% CAGR 1991 - 95
- Current transputer family established
- T9000 family
  - is compatible and suited to single, multi and parallel processor systems
  - Targetted at embedded areas of compute, communications and imaging markets
  - Balanced

Computation	- 200MIPS/25MFLOPS
Communication	- 80MBytes/s
Responsiveness	- Sub microsecond context switch
System Simplicity	- On-chip support + peripherals
Software	- Standardized: Development, system and application

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

THE TECHNOLOGY & DESIGN OBJECTIVES FOR THE T9000 FAMILY

Ian R. Pearson

SGS-THOMSON MICROELECTRONICS

inmos

THE TECHNOLOGY & DESIGN OBJECTIVES FOR THE T9000 FAMILY

SGS-THOMSON MICROELECTRONICS

inmos

Current Transputers

The basic architecture overcomes communication bottlenecks associated with shared bus architectures. Communications bandwidth increase with number of processors.

For example

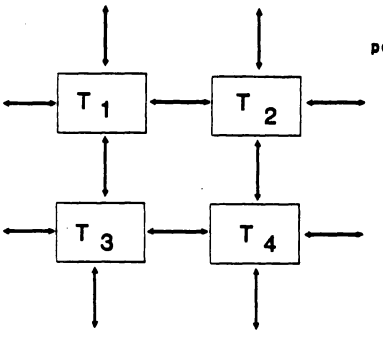

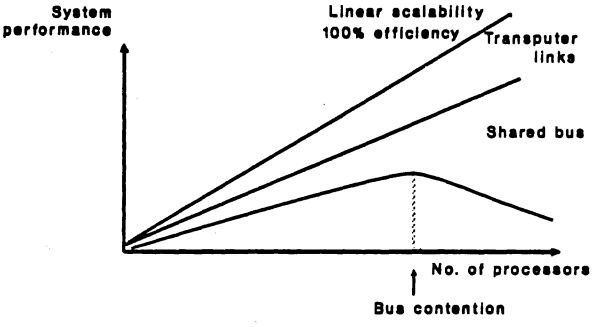

4 processors - 8 links No. of transputers	Communications bandwidth of system
1	8 Mbytes/sec
2	12 Mbytes/sec
4	16 Mbytes/sec
6	20 Mbytes/sec

SGS-THOMSON MICROELECTRONICS

inmos

Current Transputers

Key features

- Cost/ performance
- Highly integrated
- Fast context switching
- On chip RAM
- Point to point communications(links)
- Single 5MHz clock

Generic architecture

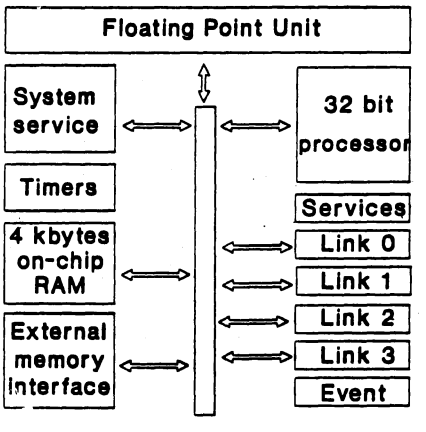

Family members

- 32 bit processors
  - Integer + FPU T805-30
  - Integer - T425-25
  - Non multiplex T801-25
- 16 bit processor
  - Integer - T225-25
- Communications
  - Link adaptor C011/12
  - 32 way crossbar C004

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Family Design Objectives

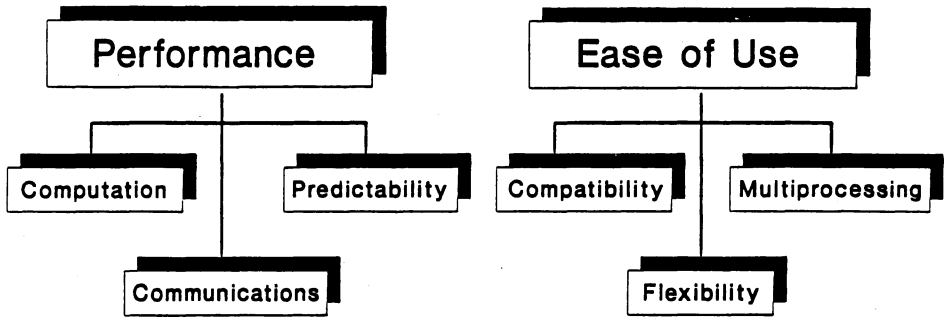

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

The T9000 Family Design Objectives

PERFORMANCE

COMPUTATION

10xT805 (200 Mips, 25 Mflops)

- 50MHz Clock (20ns cycletime)
- Multiple instructions per cycle - Superscalar Architecture. 5 stage pipeline plus instruction grouper.
- Integrated 64 bit FPU.
- Fast caching: 16 KBytes Cache - 200MW/sec.

SGS-THOMSON MICROELECTRONICS

inmos

The T9000 Family Design Objectives

PERFORMANCE

COMMUNICATIONS

Match of I/O performance to processing speed:

- Fast 64 bit memory interface - 50Mw/sec
- 100 Mbit/sec full duplex links - 20 Mbytes/sec per link.
- Dedicated Communications Processor (Virtual Channel Processor-VCP)
- Separate System Control & Status Monitoring

SGS-THOMSON MICROELECTRONICS

inmos

The T9000 Family Design Objectives

PERFORMANCE

PREDICTABILITY

Fast response to interrupts

- Hardware Scheduling & Fast Context Switching.
- Support for Industry Standard Real Time Kernels

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

The T9000 Family Design Objectives

EASE OF USE

COMPATIBILITY Code, Communications, and Operating Environments.

- Binary compatibility T805 code
- Communications compatibility (T8 -› T9) via link converter (C100) & link interfaces (C101)
- Standard Programming Languages - optimising compilers C, C++, F77, Occam
- Standard Distributed OS - VRTX, UNIX™ (Chorus)
- Standard networking support Ethernet (TCP/IP), NFS

SGS-THOMSON MICROELECTRONICS

inmos

The T9000 Family Design Objectives

EASE OF USE

FLEXIBILITY Interfacing & Configurability

- Self sufficient memory interface supports mixed memory systems
- Configurable cache - cache v internal memory
- Single 5Mhz clock - no high frequency clocks
- Small footprint packages - 208 lead (28mm) CQFP

SGS-THOMSON MICROELECTRONICS

inmos

The T9000 Family Design Objectives

EASE OF USE

MULTI-PROCESSING SUPPORT — Communications & Context Switching

- Low latency (<1 micro second) communication - in conjunction with C104
- Hardware multiplexing of messages (VCP)
- Fast context switching. <1 micro second.

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

The T9000 Family Employed In A Multiprocessing System

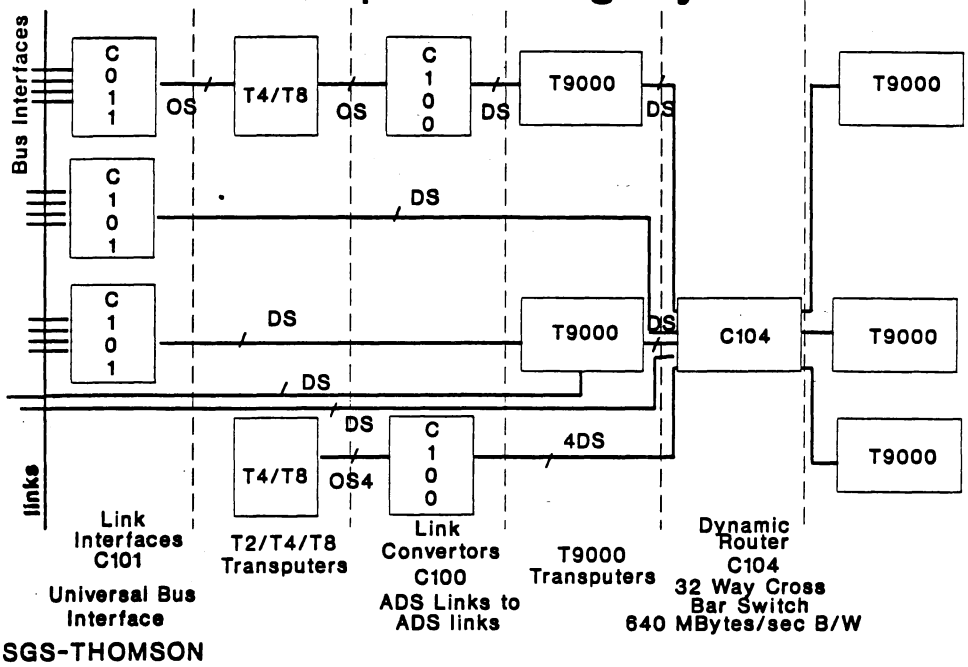

SGS-THOMSON MICROELECTRONICS

inmos

Transputer Technology

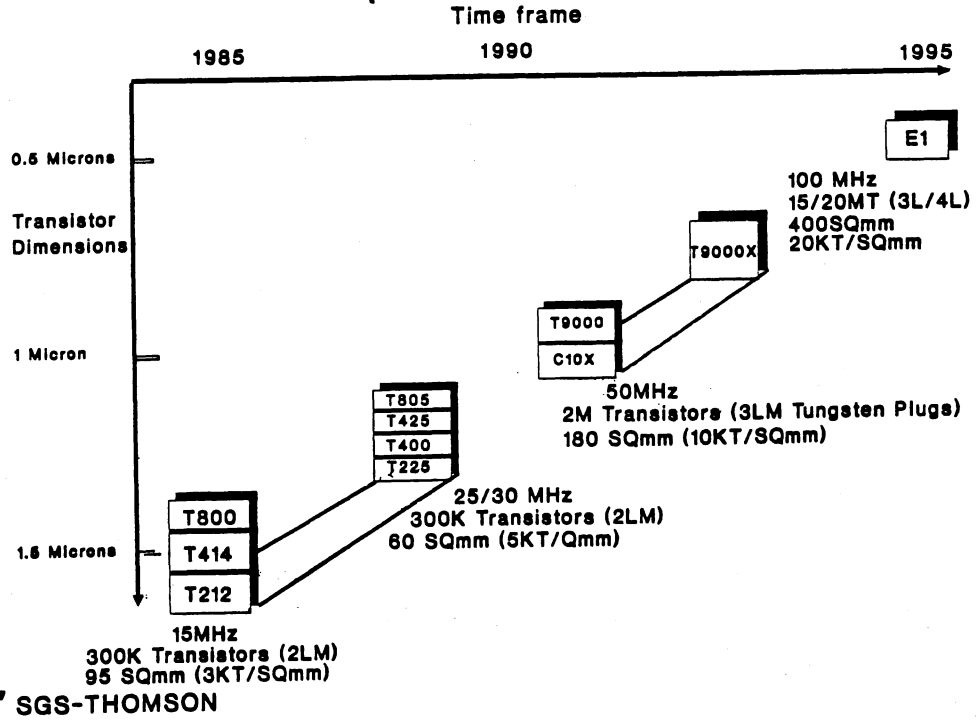

SGS-THOMSON MICROELECTRONICS

inmos

The T9000 Product Portfolio

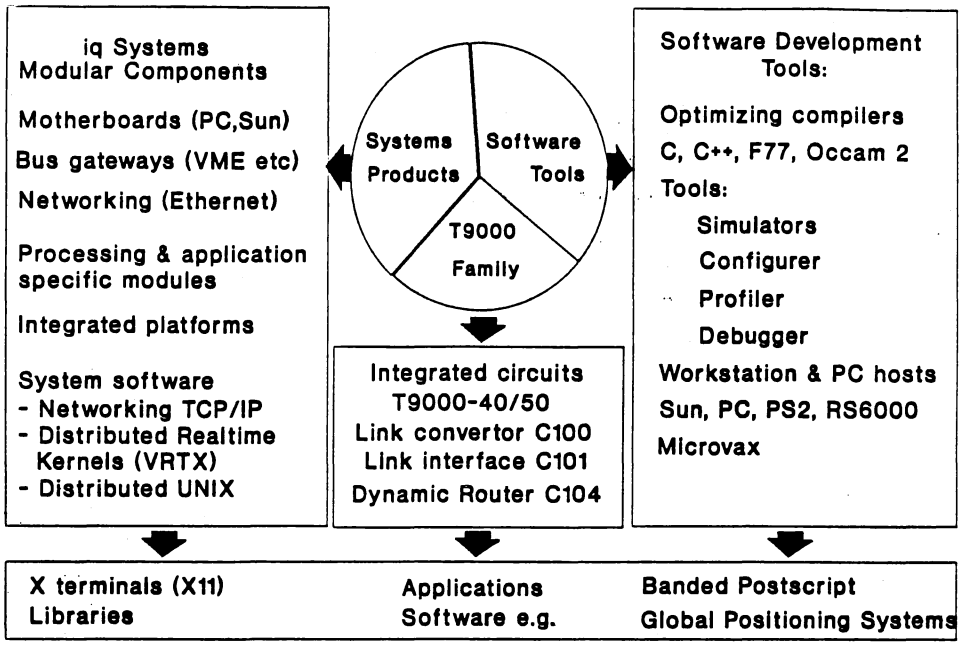

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

INMOS T9000 Software Strategy

- Evolution of Current Development Tools
- Continued Compiler Optimization
- Support for Multi-processing
- Support for Open Systems and Industry Standards
- Collaboration with Leading Third Party Suppliers
  Compilers, Kernels & Operating Systems

SGS-THOMSON MICROELECTRONICS

inmos

Full Range of Cross-Development Tools

Current INMOS Toolsets enhanced
to support T9000

● Host Machines:	PC, Sun3, Sun4, VAX
● Languages:	ANSI C, C++, Fortran, Occam 2
● Download:	via Host Motherboard or Ethernet

SGS-THOMSON MICROELECTRONICS

inmos

INMOS Compilers

- Languages: ANSI C, occam 2, FORTRAN
- First Validated ANSI C Compiler (BSI Aug 1990)
- Focus on quality of code generation
- Common backend and optimization for all compilers
- T9000 Specific Optimizations

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Common Backend Optimizations

- Global Common Subexpression Elimination
- Loop Invariant Removal
- Dead Code Elimination
- Constant Propagation
- Tail-Call and Tail-Recursion Elimination
- Improved Workspace Allocation

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Specific Improvements

- Half-word support
- New fpu, sqrt, and rem instructions used
- Removal of Floating Point Error Checking
- Various Idioms adjusted to balance CPU Pipeline
- Use of Variable Length I/O Capability
- Use of Semaphore instructions
- Full instruction set access through Assembler inse

SGS-THOMSON MICROELECTRONICS

inmos

Support for Multi-Processing

- Simplified Programming
- Automatic use of hardware message routing and multiplexing
- Software Routing used in absence of routing chips
- Details of hardware network topology and routing tables can be hidden from the applications programmer

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Network Description Language

- Extension of Current Hardware Description
- Allows declaration of devices and attributes
- T9000 Attributes: memory configuration, memory size, cache size, link speed, link protocol, boot from ROM options
- C104 Attributes: link speeds, routing table intervals, header deletion, randomized routing
- Describes construction of control network and data network
- Allows naming of specific routes through network

SGS-THOMSON MICROELECTRONICS

inmos

Network Initialization Tool

Performs checking operations on Network Description

- Connectivity Characteristics
- Construction of Control System
- Disjointness of routing table intervals
- Deadlock in routing tables
- Consistent use of header deletion
- Declared routes are routable

Produces Network Initialization File to be loaded down control link

SGS-THOMSON MICROELECTRONICS

inmos

Software Configuration Tools

- C and occam style Software Configuration Languages
- Describes Software Network and Mapping to Hardware
- No Topology Constraints
- Sets up Communication Control Blocks
- Compatible with current Configuration Languages

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Network Simulator

- Functional Simulation of T9000 and C104 networks
- Machine Level Debugging Capability
- Supports Debugging of Initalization and Bootstrapping
- Executes on a T800 Network

SGS-THOMSON MICROELECTRONICS

inmos

New Source Level Debugger

- OSF/Motif point and click interface
- Designed to support parallel programming
- Multiple Windows to support programming model
- Support for Mixed Languages
- Remote debugging over Ethernet

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Features Exploited in Debugger

- L-Process Trap Handlers
- Single Stepping
- Hardware Watchpoints
- Protection (second revision)

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Preparing Programs for Debugging

1. Compile with debug switch to generate symbol info -
   small code changes for occam only
2. Link with debug run-time - new routines for process creation and communication
3. Configuration Tools add Debugger Kernel processes

- Low Memory Overhead per Processor (about 20 KBytes)
- Efficient Code Execution

SGS-THOMSON MICROELECTRONICS

inmos

Programming Model

- Hardware Network, Processor, Connection, Edge
- Software Network, Process, Channel, Thread, Semaphore
- Program Source, Module, Function
- Dynamic State viewed through Program Browser
- Source Texts through Source Browser

SGS-THOMSON MICROELECTRONICS

inmos

Debugging Operations

- Conditional Breakpoints
- Single Stepping: Into, Over, Out of
- Watchpoints
- Function Entry and Exit
- Stack Backtrace
- Variable inspection and update

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Additional Event Mechanisms

- Thread Creation and Completion
- Channel Communication
- Semaphore Operations
- Error Traps
- User Interrupt

SGS-THOMSON MICROELECTRONICS

inmos

Third Party Compiler Support

Glockenspiel C++

- Developed by Glockenspiel
- Pre-processor to INMOS ANSI C compiler
- Source level debugging through C debugger
- Available from INMOS

Ada

- Full Ada development environment
- Available from Alsys

SGS-THOMSON MICROELECTRONICS

inmos

Third Party Operating Systems and Kernels

CHORUS/MiX: Chorus Systemes, Telmat

- Distributed UNIX System V
- Supports X/Open, POSIX, and NFS Standards

VRTX32: Ready Systems

- The Leading Multi-Tasking Real Time Executive
- Extensions to support Multi-Processing
- Support for RTScope Debugger
- Compatible with INMOS ANSI C Toolset

C-Executive: JMI

- Popular Multi-Tasking Real Time Executive
- Compatible with INMOS ANSI C Toolset

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

OCCAM 3

- Structured datatypes, Named types,
  Records, Unions
- Extended process model
  Remote procedure calls, Shared resources
- Modules and Libraries defined in language
- Occam 2 is a subset
- Simple, very efficient, parallel programming

SGS-THOMSON MICROELECTRONICS

inmos

Cross-Development Tools

- Use of Standard Development Platforms
- Current Hosts: IBM PC, NEC PC, Sun 3, Sun 4, VAX (VMS)
- Portability of Development Tools
- Fast Download and Remote Debugging via EtherNet and TCP/IP or via Host Interface Board

SGS-THOMSON MICROELECTRONICS

inmos

Major Toolset Enhancements

Common Back-end and Optimizer
- For all INMOS Compilers

New source-level, Network Debugger
- OSF/Motif based interface

T9000 Network Simulator
- Machine-level debugging

Network configuration tools
- Supports T9000 communications
- Simplifies applications programming

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Third Party Software

Compilers

- C++ , Glockenspiel
- Ada , Alsys

Real-Time Kernels & Operating Systems

- Distributed UNIX , Chorus
- VRTX32/T , Ready Systems
- C-Exec , JMI Software

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

The T9000
Transputer Product Family

A Major Advance
in Embedded Processing

SGS-THOMSON MICROELECTRONICS

inmos

The T9000
Transputer Product Family

A Major Advance
in Embedded Processing

Clive Dyson

SGS-THOMSON MICROELECTRONICS

inmos

The T9000

› 2 Million transistors

50 MHz

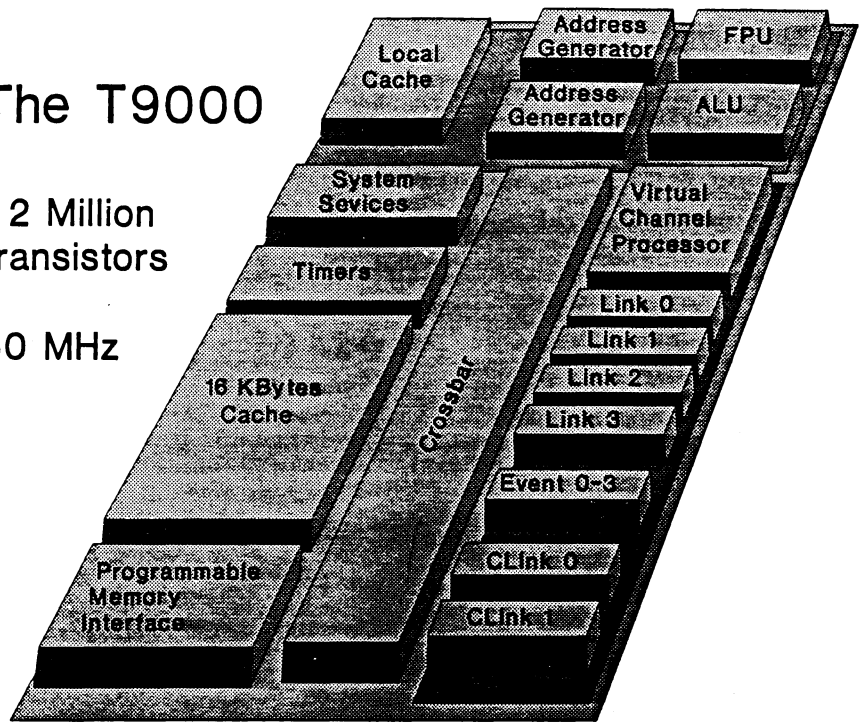

SGS-THOMSON MICROELECTRONICS

inmos

Balance of Computation and Communication

Assuming 2 FLOPS per Word

T805-20

OS-Links
1.25 MWords/s

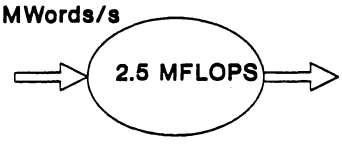

2.5 MFLOPS Peak

T9000-50

DS-Links
8.25 MWords/s

25 MFLOPS Peak

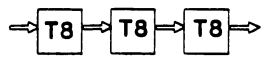

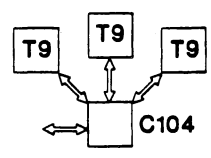

SGS-THOMSON MICROELECTRONICS

inmos

Data Bandwidths

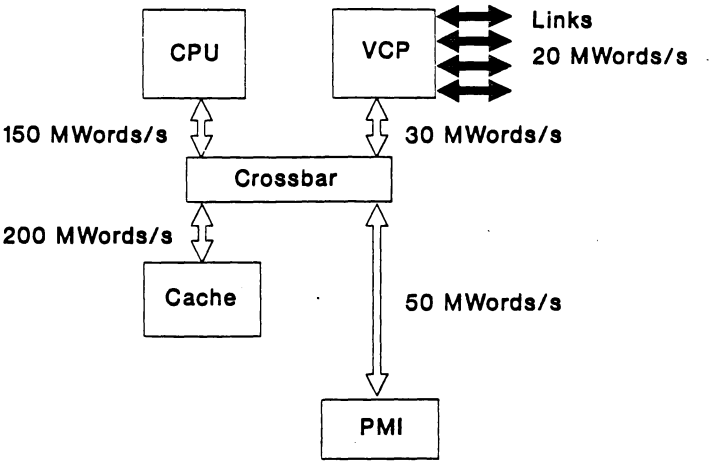

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Benchmarks

	T805-20	T9000-50 (Optimized)	T9000-50 (T805 code)
1024 point single-precision Complex FFT	33 ms 6.9 MIPS 1.3 MFLOPS		3.2 ms 71 MIPS 13.8 MFLOPS
Color Transform	8.8 MIPS	144 MIPS	
PostScript "Golfer"	‹ 17 seconds	‹ 1.5 seconds	
Single-precision Whetstone Double-precision Whetstone	4.5 MWhet/s 3.0 MWhet/s	40 MWhet/s 32 MWhet/s	30 MWhet/s 24 MWhet/s
Dhrystone 2.1	6.6 kDhry/s	64 kDhry/s	47 kDhry/s

- All compiled from source
- All fully IEEE compatible FP

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

THE

9000

TRANSPUTER

T9000 PROCESSOR

COMMUNICATIONS

ENHANCED PROCESS MODEL

SGS-THOMSON MICROELECTRONICS

inmos

Instruction Set Architecture

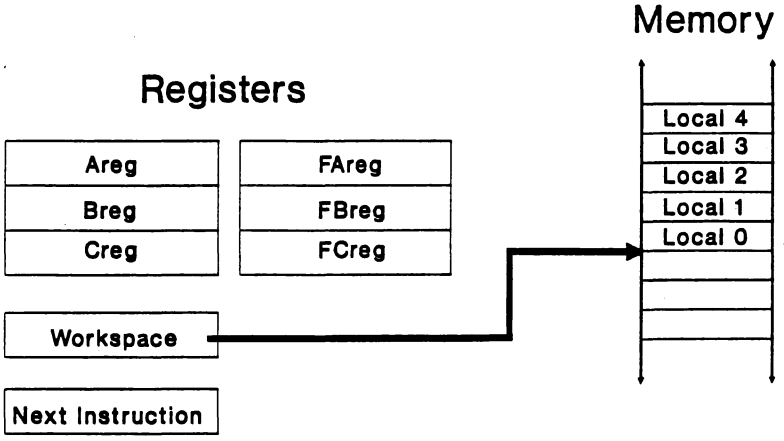

- Sub-Microsecond Context Switch
- 32 Locals in Workspace Cache

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Workspace Cache

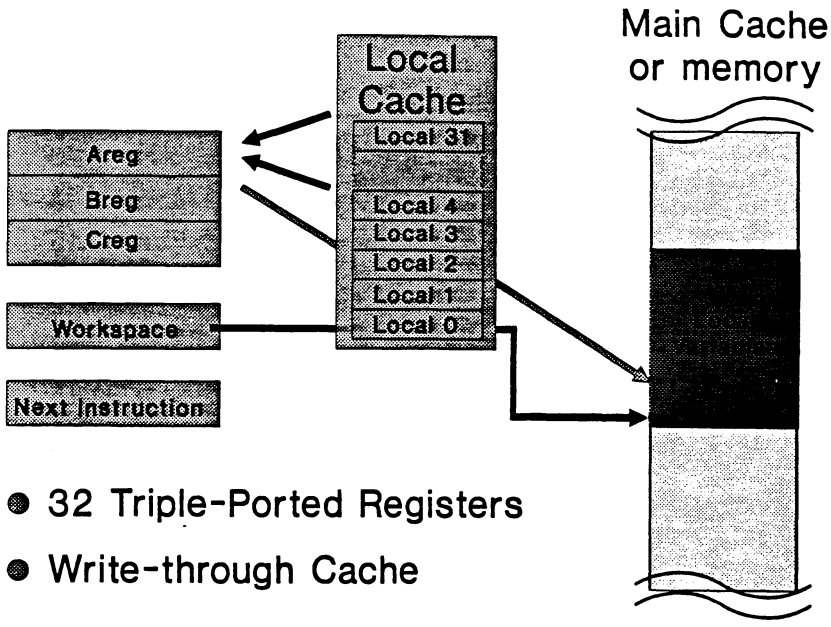

- 32 Triple-Ported Registers
- Write-through Cache

SGS-THOMSON MICROELECTRONICS

inmos

Instruction Grouping

```
a[i+1] = b[j+15] + c[k+7];
```

- 32 Instruction Fetch-Ahead Buffer
- Hardware Grouper
- T9000 issues up to 8 Instructions per cycle

1st Group

ldl j	Load value in j
ldl b	Load address of b
wsub	Address b[j]
ldnl 15	Value of b[j+15]

2nd Group

ldl k	
ldl c	
wsub	
ldnl 7	Value of c[k+7]
add	Add stack values

3rd Group

ldl i	
ldl a	
wsub	Address of a[i]
stnl 1	Write result

SGS-THOMSON MICROELECTRONICS

inmos

Pipelining

```
a[i+1] = b[j+15] + c[k+7];
```

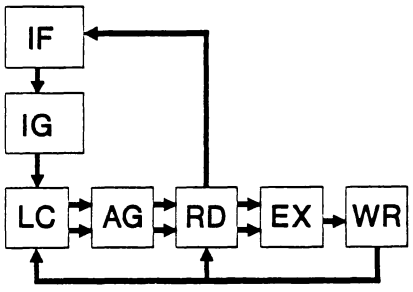

1st Group

LC	ldl j	Load value in j
LC	ldl b	Load address of b
AG	wsub	Address b[j]
AG RD	ldnl 15	Value of b[j+15]

2nd Group

LC	ldl k	
LC	ldl c	
AG	wsub	
AG RD	ldnl 7	Value of c[k+7]
EX	add	Add stack values

3rd Group

LC	ldl i	
LC	ldl a	
AG	wsub	Address of a[i]
AG WR	stnl 1	Write result

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Pipelining Instruction Groups

- 1 instruction group per clock cycle

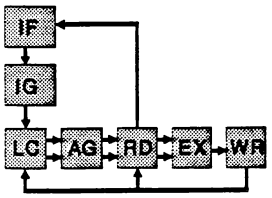

G2

LC	ldl k
LC	ldl c
AG	wsub
AG RD	ldnl 7
EX	add

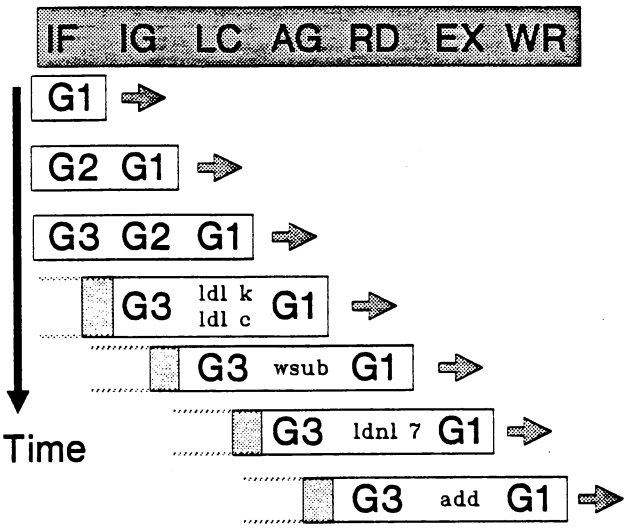

SGS-THOMSON MICROELECTRONICS

inmos

Instruction Cycle Counts

Integer

	T805-20		T9000-50	
shift left/right	n + 2	< 1,800 ns	1	20 ns
multiply	38	1,900 ns	2-5	< 100 ns
divide	39	1,950 ns	5-12	< 240 ns

Floating Point

	T805-20		T9000-50	
fp add (32 & 64)	6-9	< 450 ns	2	40 ns
fp multiply (32)	11-18	< 900 ns	2	40 ns
fp multiply (64)	18-27	< 1,350 ns	3	60 ns
fp divide (32)	16-28	< 1,400 ns	8	160 ns
fp square root (32)	119-122	< 6,100 ns	8	160 ns

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Datapath Crossbar

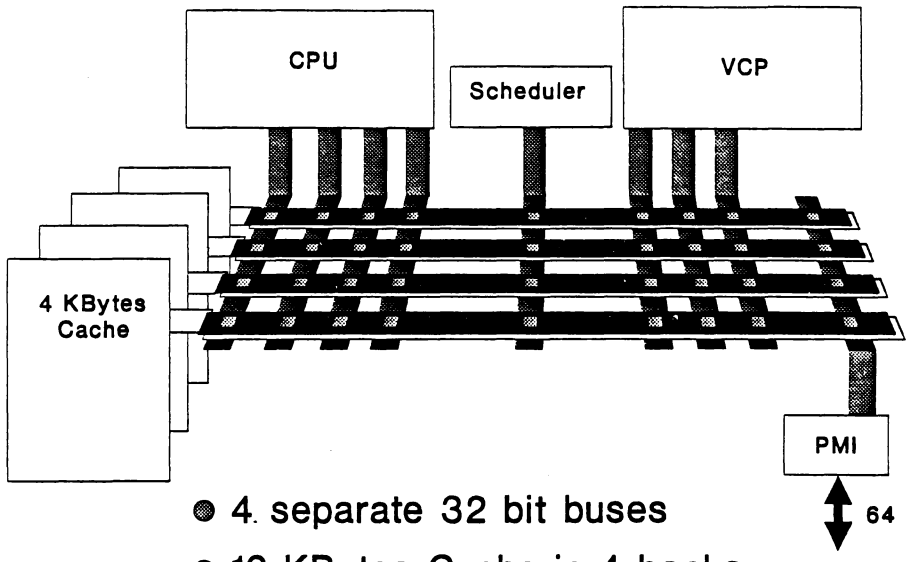

- 4 separate 32 bit buses
- 16 KBytes Cache in 4 banks

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Main Cache

- 16 KBytes shared by Data & Instructions
  Write-back
- 4 banks, 256 lines of 4 words
  Fully Associative
- Four accesses per cycle, 800 MBytes/s
- Random Replacement Strategy

SGS-THOMSON MICROELECTRONICS

inmos

Minimal System

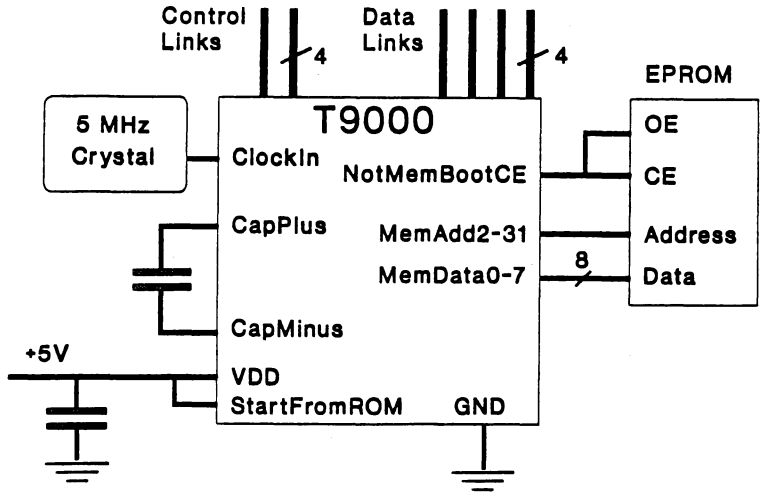

- Phase Lock Loop
- 16 KBytes On-Chip RAM

SGS-THOMSON MICROELECTRONICS

inmos

Programmable Memory Interface

- System Simplicity & Compactness
  Usually zero Glue Logic
- 64 bit Data Bus, 32 bit Address Bus
  Bandwidth up to 200 MBytes/s
- 4 Independent Banks
  Support for DRAM,SRAM,EPROM,VRAM,FROM
- Optimized for Cache Refill
  Fast Page-Mode DRAM support
- Programmable Bus Widths
  Configurable Timing for each Bank

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

T9000 PMI

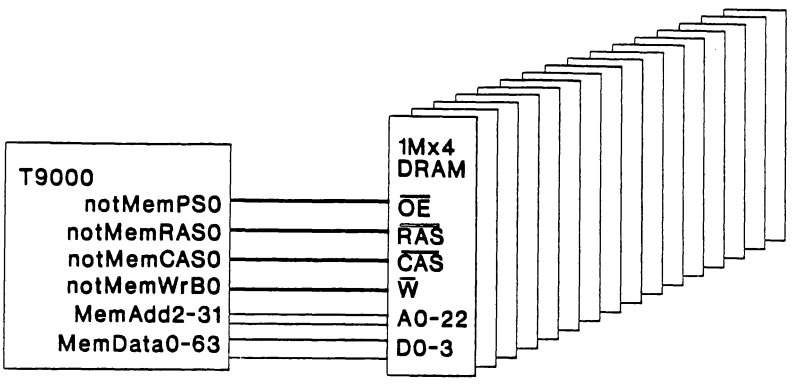

- 64 bit Data Bus
- 8 MBytes DRAM with zero Glue Logic

SGS-THOMSON MICROELECTRONICS

inmos

Page-mode DRAM

32-bit interface cache refill from bank 2

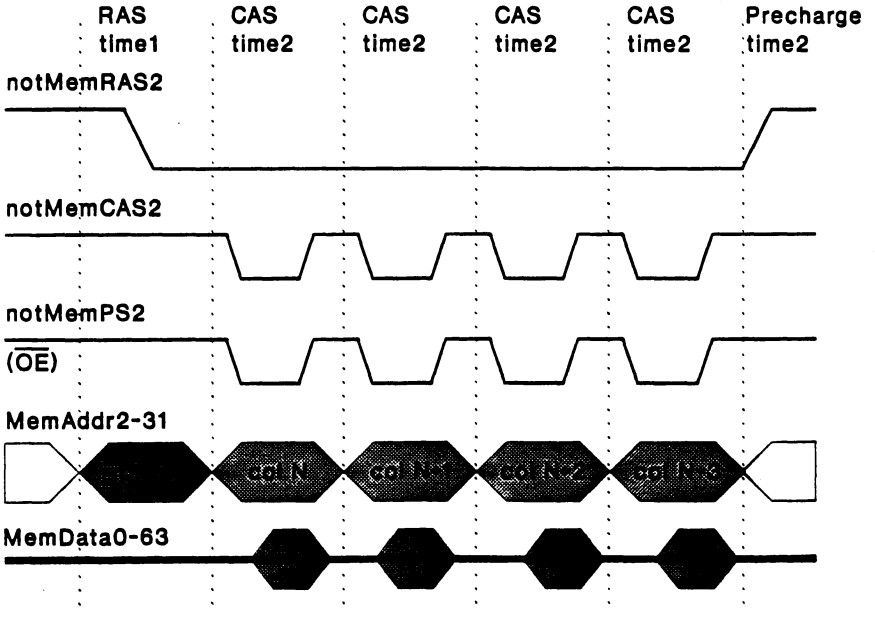

SGS-THOMSON MICROELECTRONICS

inmos

Half-word and byte support

- New T9000 instructions (except *)
- Enhanced C performance

Description	8 bit word		16 bit word
load	*lb*	*	*ls*
load and sign extend	*lbx*		*lsx*
store	*sb*	*	*ss*
sign extend	*xbword*		*xsword*
check signed	*cb*		*cs*
check unsigned	*cbu*		*csu*
subscript	*bsub*	*	*ssub*

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Processor - summary

- T805 code compatibility
- Superscalar, pipelined implementation
- Workspace cache
- Large on-chip cache
- Instruction cycle count reductions
- Byte and half-word support
- Faster clock speeds

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

THE

9000

TRANSPUTER

T9000 PROCESSOR

COMMUNICATIONS

ENHANCED PROCESS MODEL

SGS-THOMSON MICROELECTRONICS

inmos

Control Links

- Resetting
- Configuring
- Booting
- Error Reporting
- Debugging

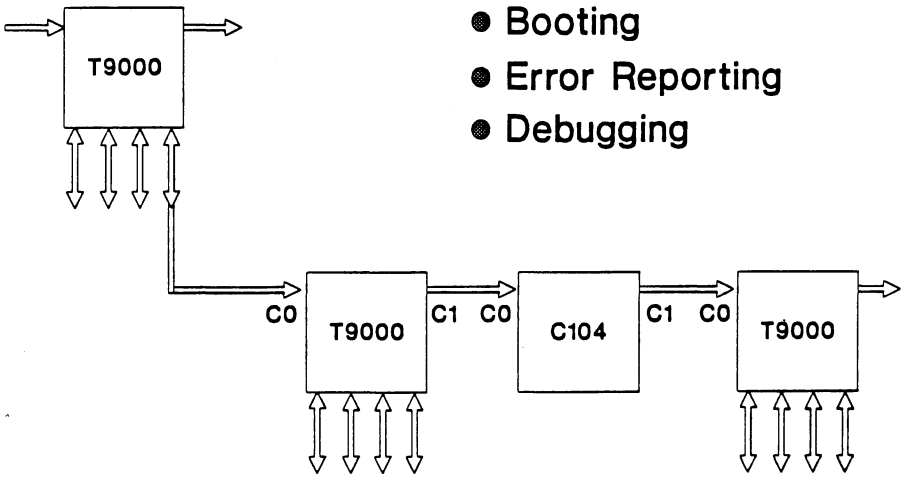

SGS-THOMSON MICROELECTRONICS

inmos

Packet Synchronization

EOP = End Of Packet Token
EOM = End Of Message Token

Example: 80 byte message

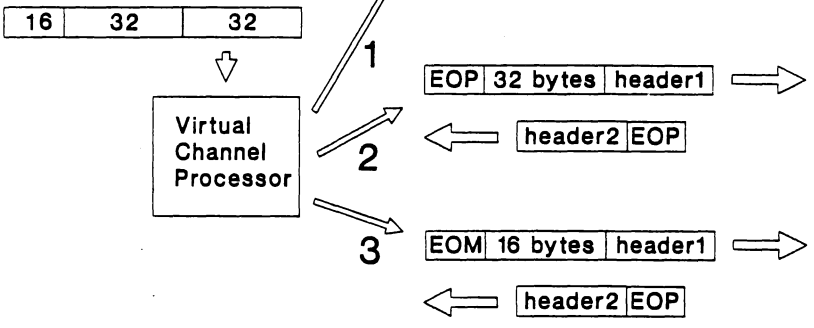

- Acknowledges typically overlap messages

SGS-THOMSON MICROELECTRONICS

inmos

Message Multiplexing

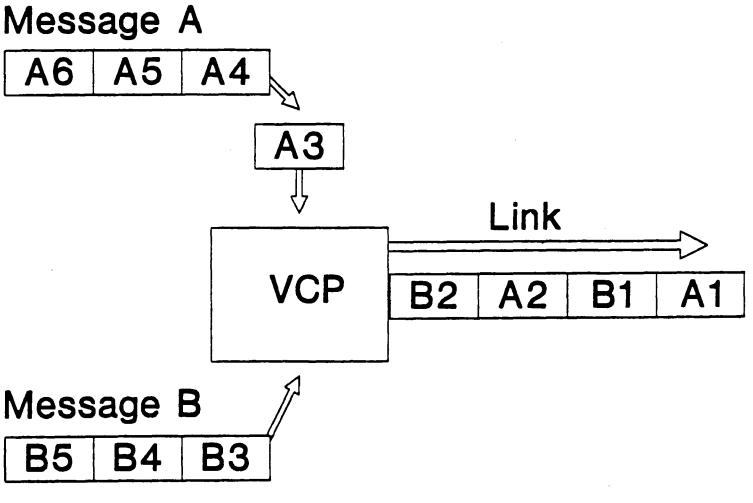

SGS-THOMSON MICROELECTRONICS

inmos

Virtual Link Control Blocks

- 8 word block used by VCP
- Two virtual channels
- 64K virtual links require 4MBytes memory

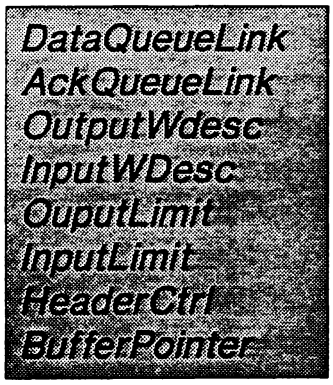

Packets to send
Acknowledges to send
Outputting process
Inputting Process
Sender's data upper bound
Receiver's upper bound
Packet header format
Address of 32 byte buffer

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Packet Header Decode

- Virtual Link Control Blocks
- Buffers for process not ready

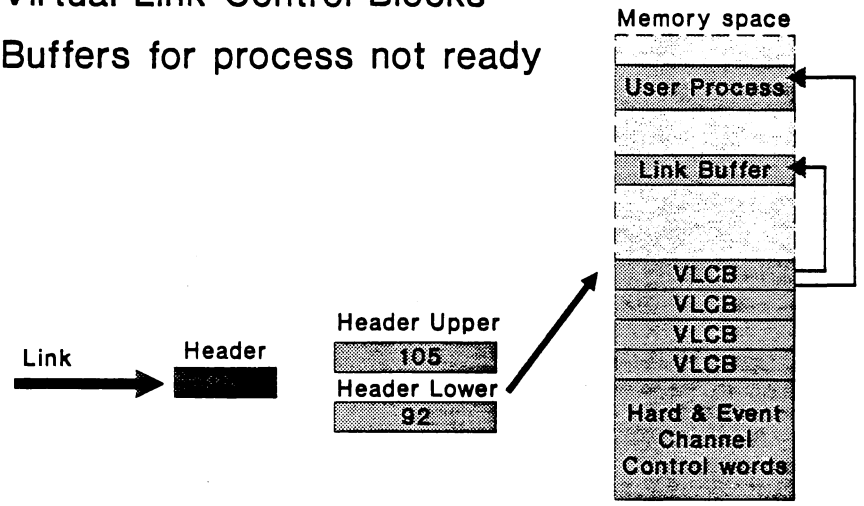

address = v.base + ((95 - 92) « v.shift)

SGS-THOMSON MICROELECTRONICS

inmos

DS-Links

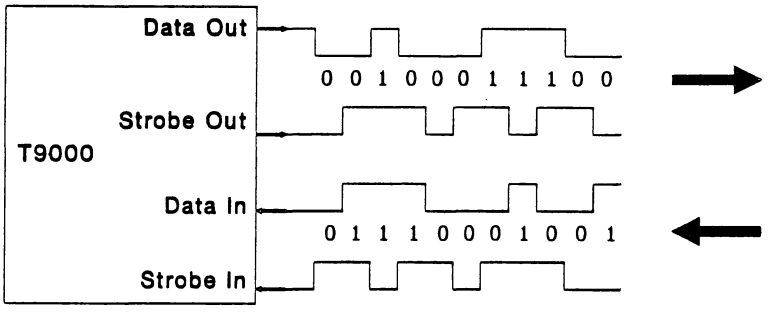

- 100 MBits/s, TTL, 100 Ohm
- Parity & Disconnection trapped
- Chip to chip Flow Control

SGS-THOMSON MICROELECTRONICS

inmos

DS-Link Format

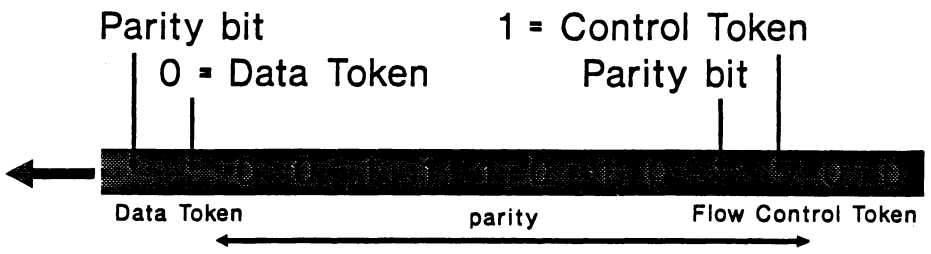

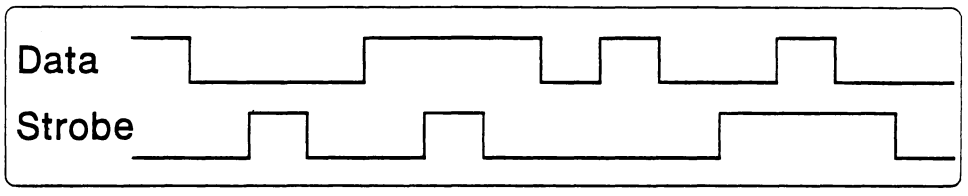

Data Token	P 0 D D D D D D D D
Flow Control Token	P 1 0 0
End Of Packet	P 1 0 1
End Of Message	P 1 1 0
Escape Token	P 1 1 1
Null Token	P 1 1 1 P 1 0 0

P = parity
D = data

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Chip-to-Chip Flow Control

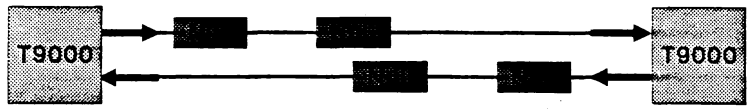

When idle, links send Null tokens

FCT indicates 8 available token buffers

FCT is sent as 8 token buffers clear

SGS-THOMSON MICROELECTRONICS

inmos

Layered Message Protocol

- Applications user only sees messages

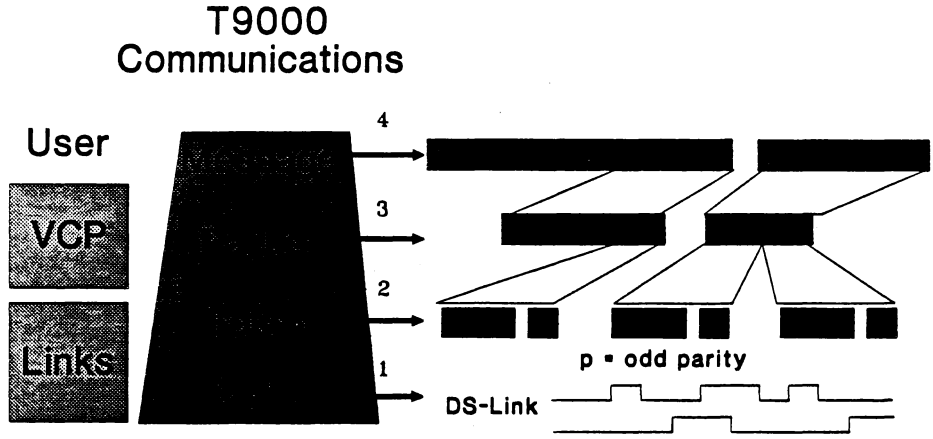

SGS-THOMSON MICROELECTRONICS

inmos

Low Latency Forwarding

- Only Header is read
- No External Memory
- Wormhole Routing

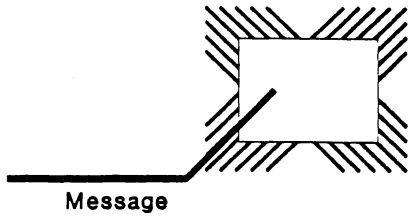

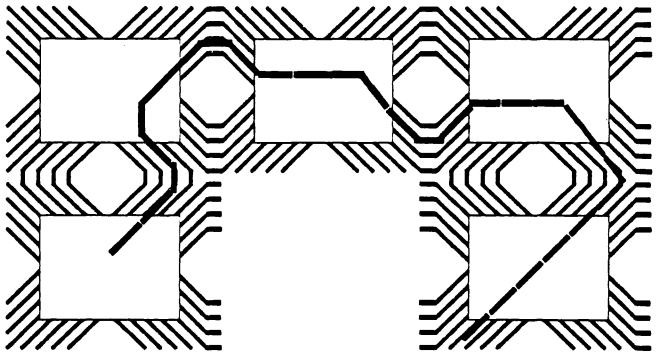

SGS-THOMSON MICROELECTRONICS

inmos

C104, Interval Labelling

- Renumber Virtual Links into Intervals
- Concurrent Interval Comparisons

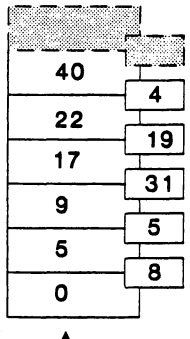

Route to Link 31

Header = 16

SGS-THOMSON MICROELECTRONICS

inmos

Computer Communications Networks

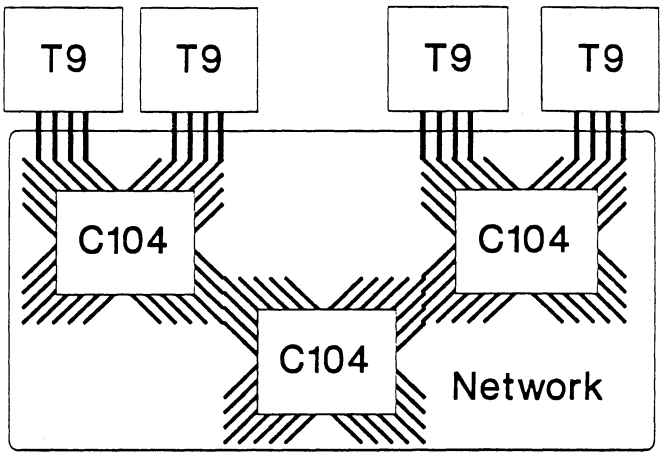

- Messages routed by Packet Switching
- Low Latency, 700ns per C104
- Reliable, efficient, Routing Algorithm

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Why Interval Labeling?

- Complete
  all packets reach destination
- Deadlock free

- Efficient
  low overhead; near shortest route
- Simple
  minimize latency; implementable
- Network independent
  large, small, different topologies

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

Labeling with 2-byte headers

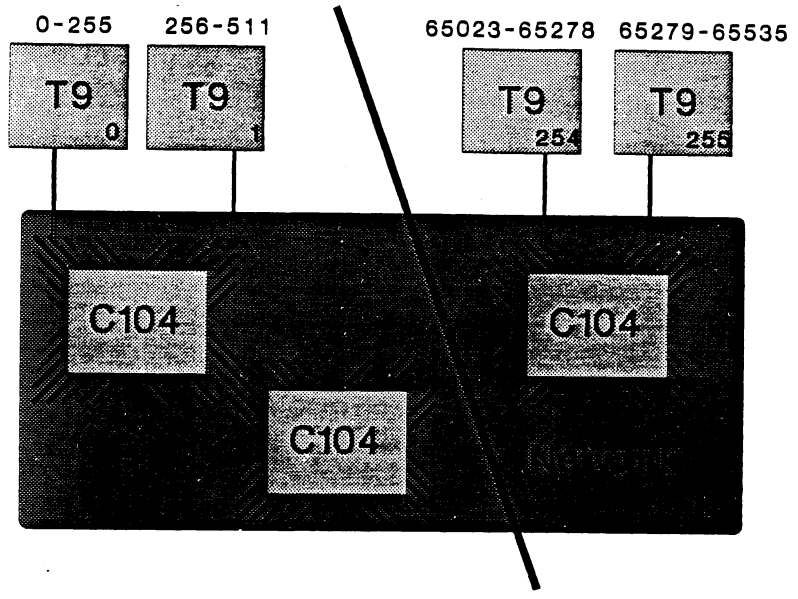

- The T9000 can recognize 1 or 2 byte headers

SGS-THOMSON MICROELECTRONICS

Concatenated sub-headers

- The T9000 can transmit multiple sub-headers

eg. System Board Processor Virtual channel

- Every sub-header selects a route
- C104 header deletion exposes a new header
- Hierarchical routing decisions

SGS-THOMSON MICROELECTRONICS

inmos

Header Deletion

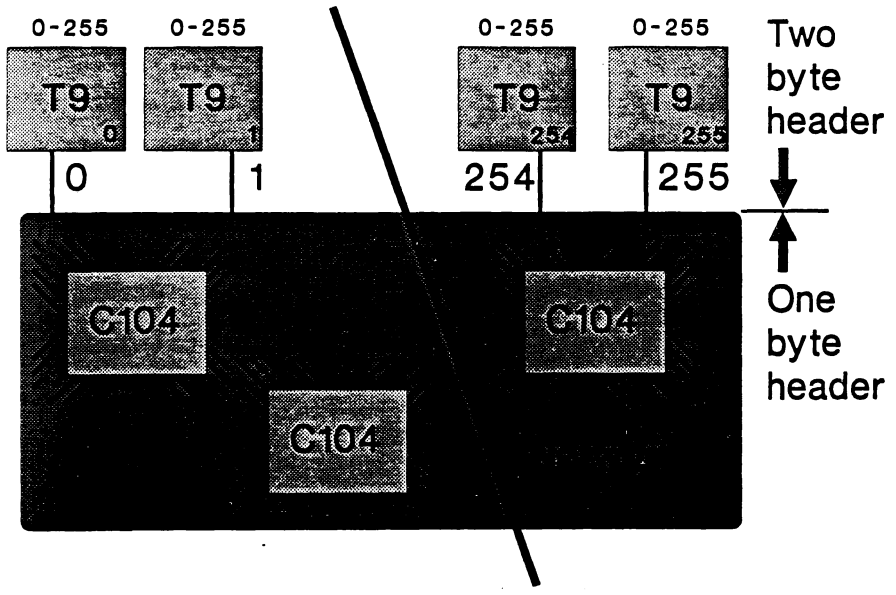

- C104 links can delete first header on output

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Universal Routing

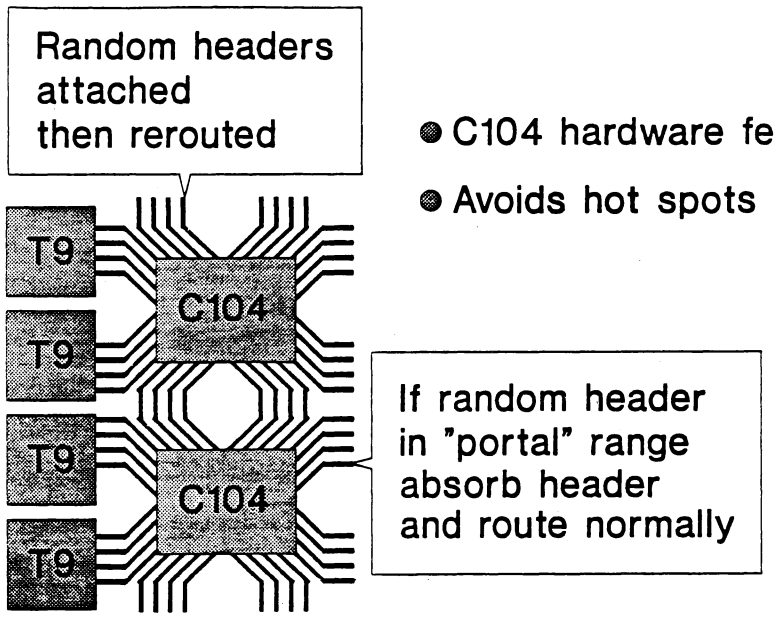

- C104 hardware feature
- Avoids hot spots

SGS-THOMSON MICROELECTRONICS

inmos

Universal Routing

or hot-spot avoidance

Under high communication loads:
- increases capacity
- reduces latency

Tradeoff between
- best-case latency
- average latency

Two-phase routing algorithm
- route to a random node
- then route normally

SGS-THOMSON MICROELECTRONICS

inmos

Single C104 Network

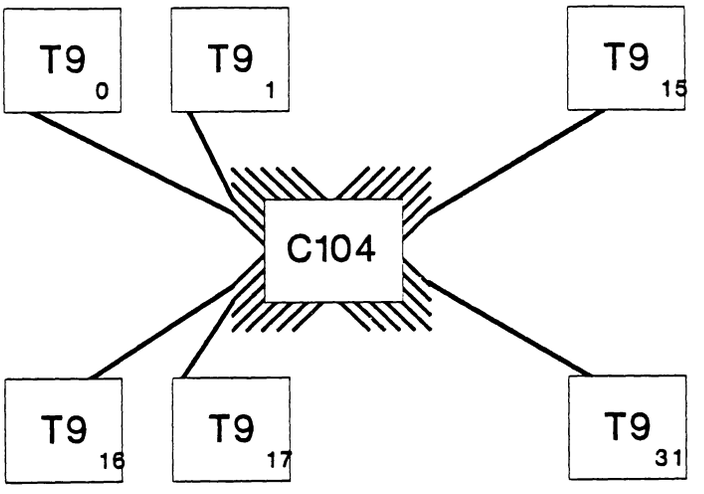

- Bisectional Bandwidth 320 MBytes/s

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Triple C104 Network

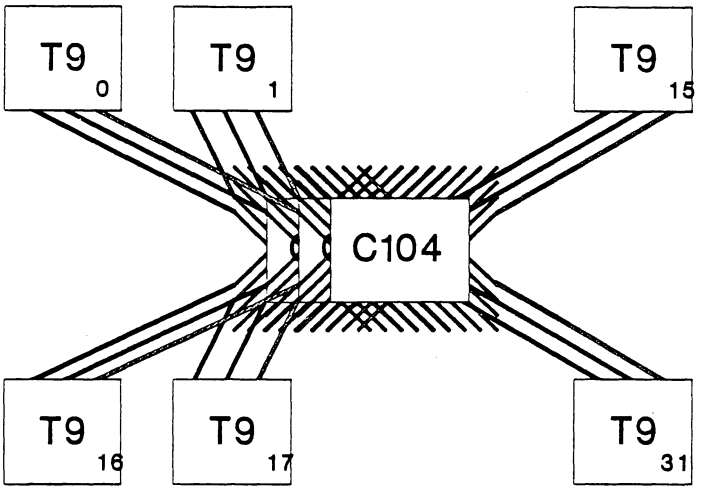

- Bisectional Bandwidth 960 MBytes/s
  Using 4 C104s, 1.28 GBytes/s

SGS-THOMSON MICROELECTRONICS

inmos

Real-Time Video Generation

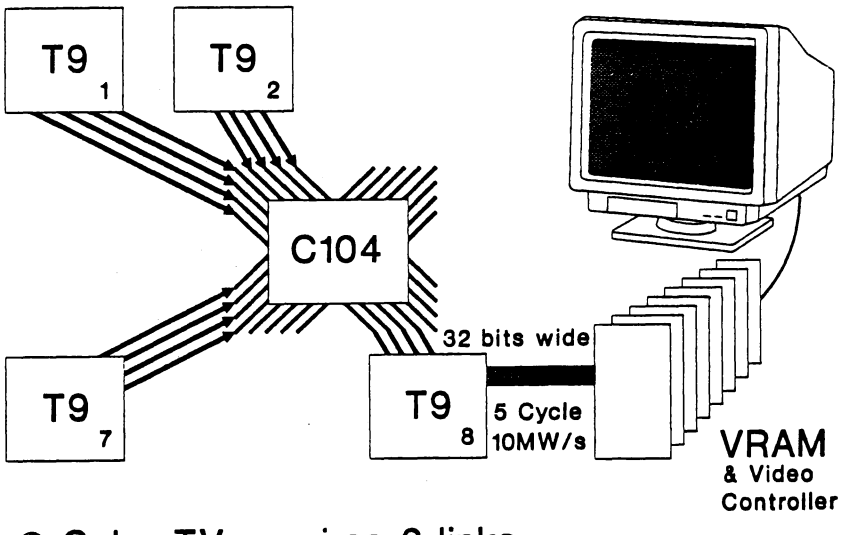

- Color TV requires 2 links
- Scalable Computation

SGS-THOMSON MICROELECTRONICS

inmos

Link Conversion

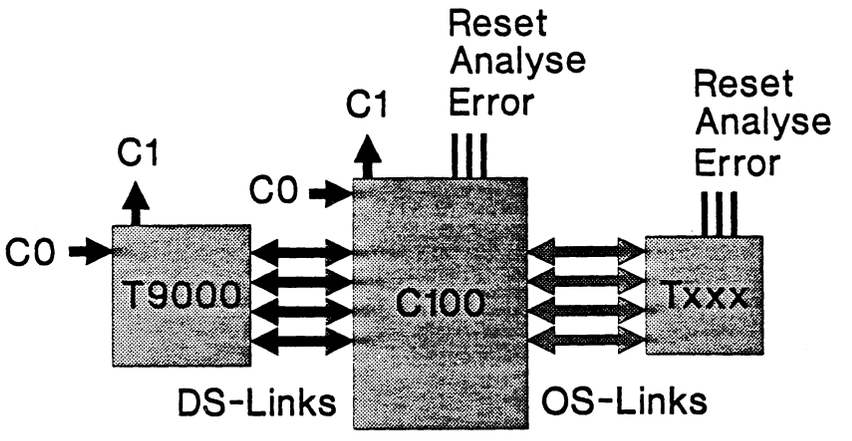

- T9000 or Txxx can be master
- Links can be byte stream
  or packetized (Txxx Library)

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Communications - summary

Virtual Channel Processor
- Message-based communications in hardware

Virtual links
- Programs independent of network topology

C104
- Low latency
- Hierarchical networks
- Universal routing

DS-Links
- New high-bandwidth link protocol

C100
- Compatibility

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

THE

9000

TRANSPUTER

T9000 PROCESSOR

COMMUNICATIONS

ENHANCED PROCESS MODEL

SGS-THOMSON MICROELECTRONICS

inmos

Local Processes

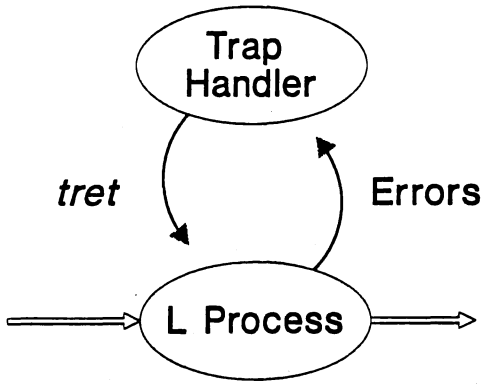

SGS-THOMSON MICROELECTRONICS

mos

G and L processes

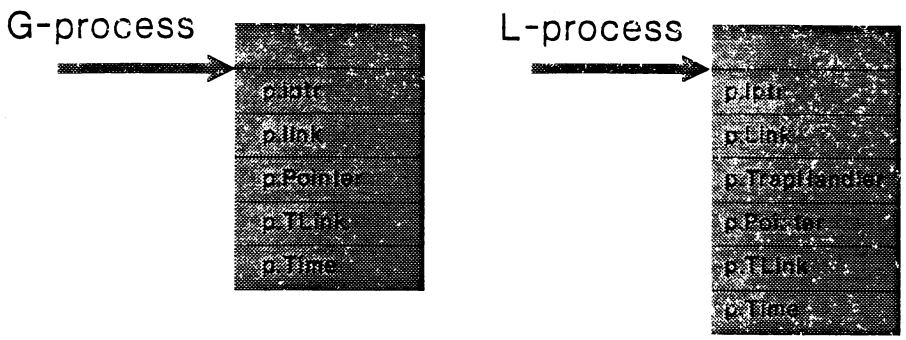

- G process is T805 compatible
- L process has local error handling
- WdescReg bit 1 determines process type

mos

Trap handler data structures

ThInUse bit in control word
allows shared trap handlers

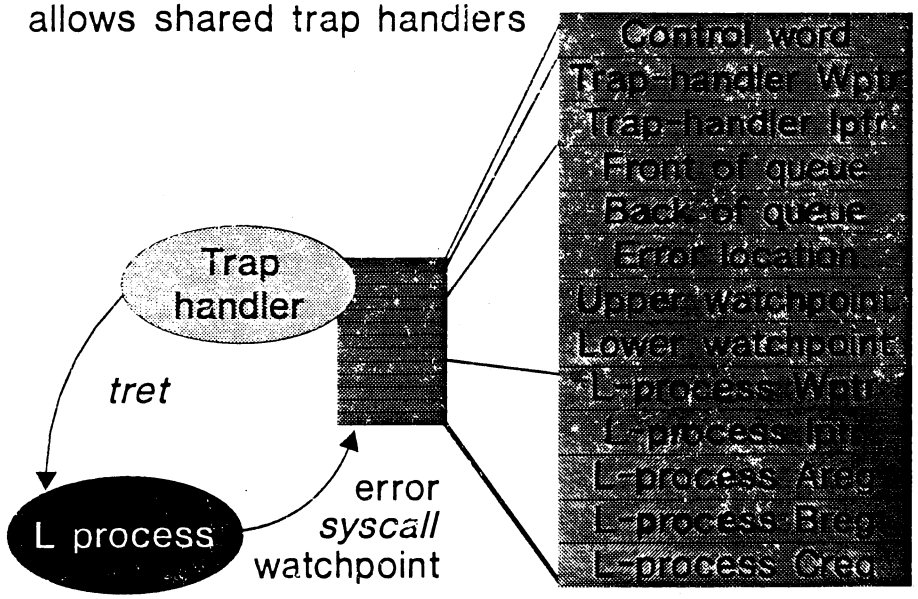

inmos

Memory Protection

- P Processes run at full speed

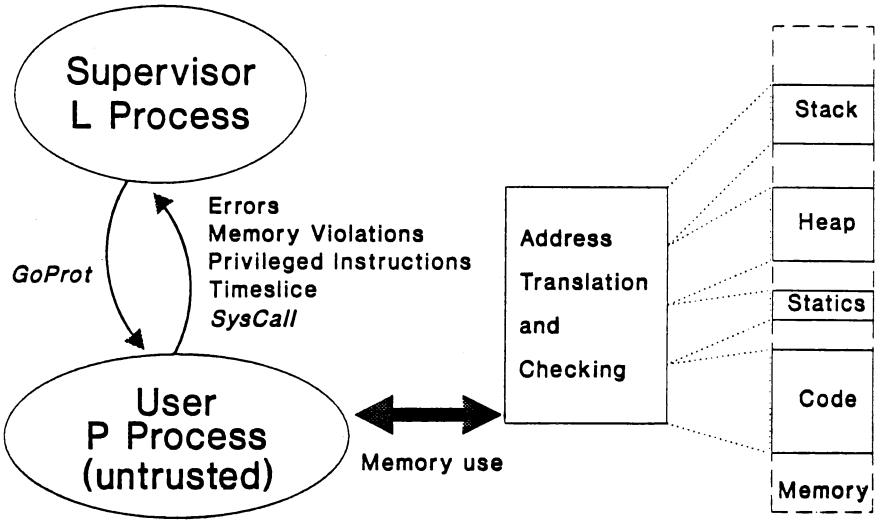

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Protected processes

L-process acts as Supervisor

P-process is started by Supervisor
- Supervisor initializes region descriptors

P-process uses region description registers for memory translation and relocation

On error, privilege violations, timeslice or system call
- P-process state is written into memory
- P-process returns reason code in Areg

For privileged instruction violation Supervisor can check then execute

SGS-THOMSON MICROELECTRONICS

inmos

Region descriptors

- P-processes use 4 memory regions
- A region is specified by 1 word
- Address translation hardware adds no overhead

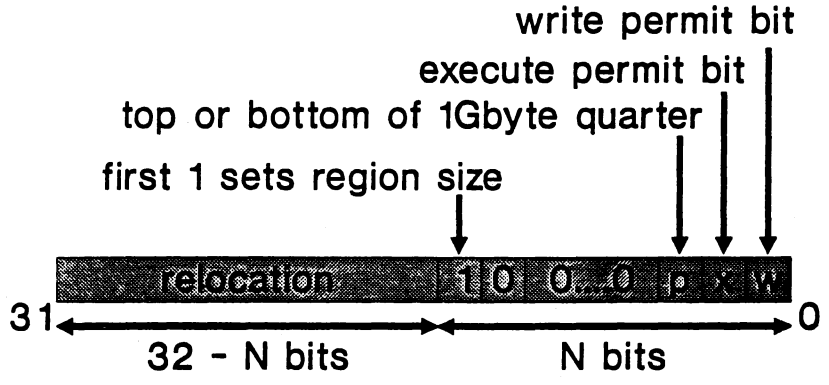

SGS-THOMSON MICROELECTRONICS

The translation mechanism

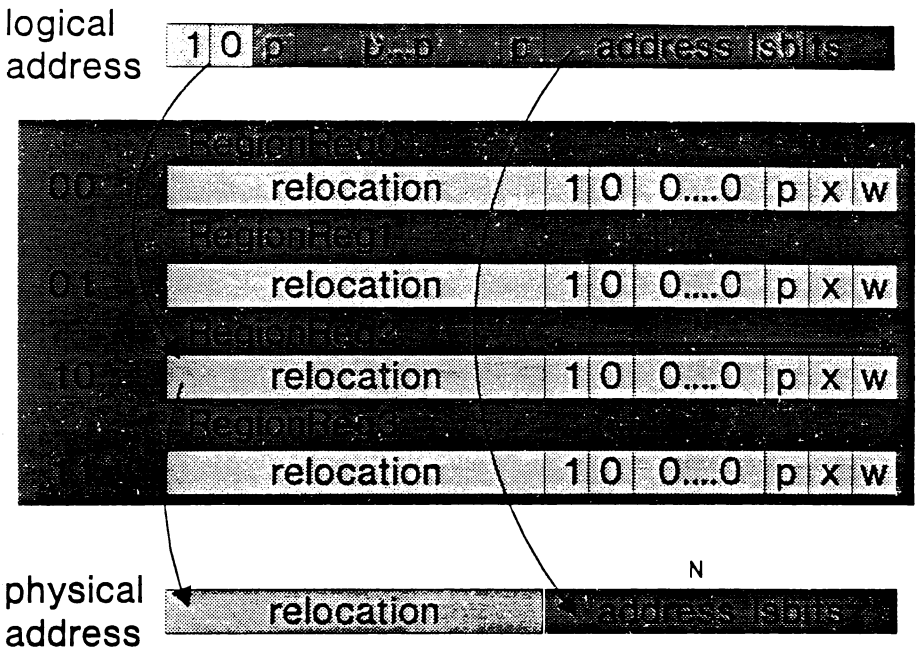

inmos

SGS-THOMSON MICROELECTRONICS

Logical address space

Region sizes
2^m 2^n 2^p 2^q

on any
2^N boundary

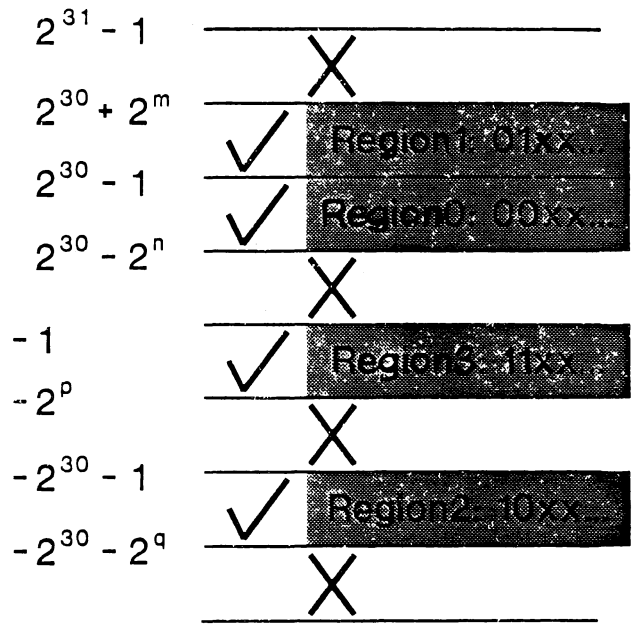

Classic N-valued semaphores

- Processes on the same processor

Semaphore
data
structure

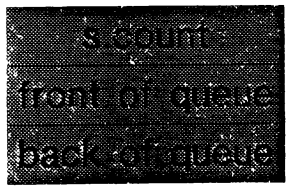

Wait

If s.count = 0
  deschedule
otherwise
  decrement s.count

Signal

If process on queue
  reschedule it
otherwise
  increment s.count

Resource Channels

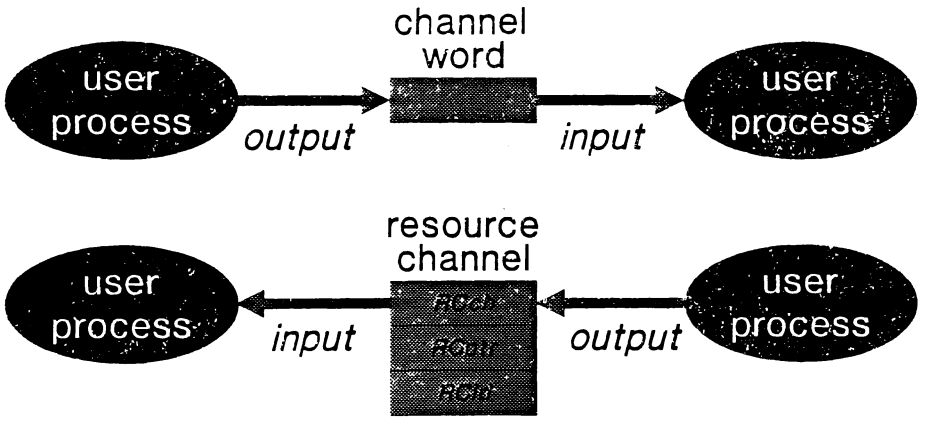

- Like other channels until used for resources

RCch	Descheduled process pointer
RCptr	For RC queues or return channel
RCid	User programmable identifier

inmos

SGS-THOMSON MICROELECTRONICS

Resource Data Structures

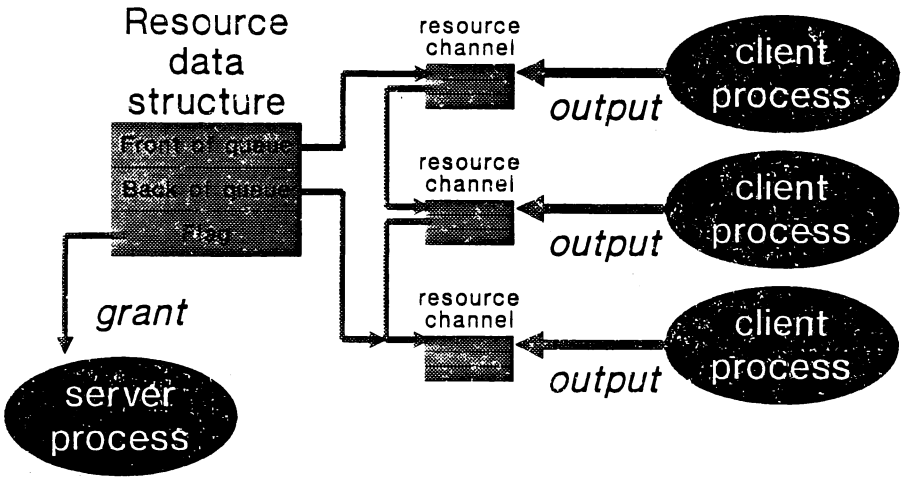

- Solves 1000-way ALT problem
- *grant* enables RClink as return channel

SGS-THOMSON MICROELECTRONICS

Resources across networks

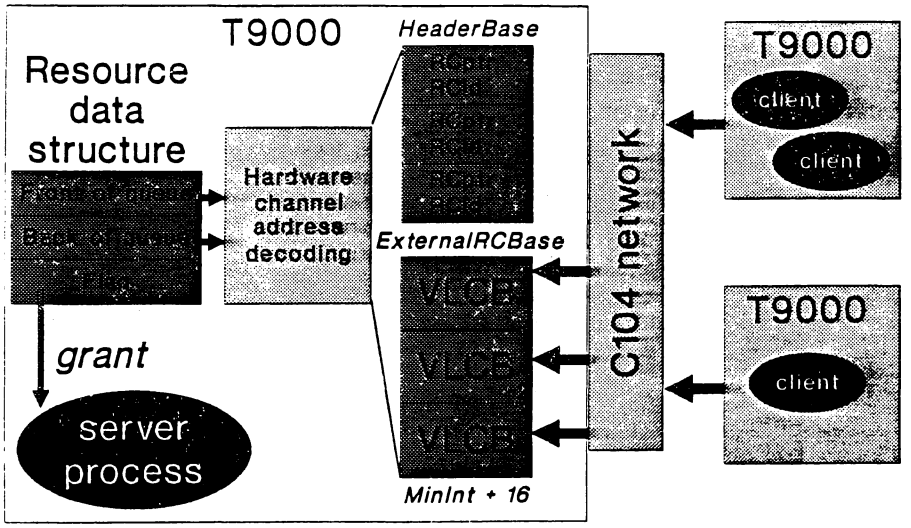

- Hardware handles virtual link connections
- Same instruction for internal/external clients

SGS-THOMSON MICROELECTRONICS

inmos

Enhanced Connectivity

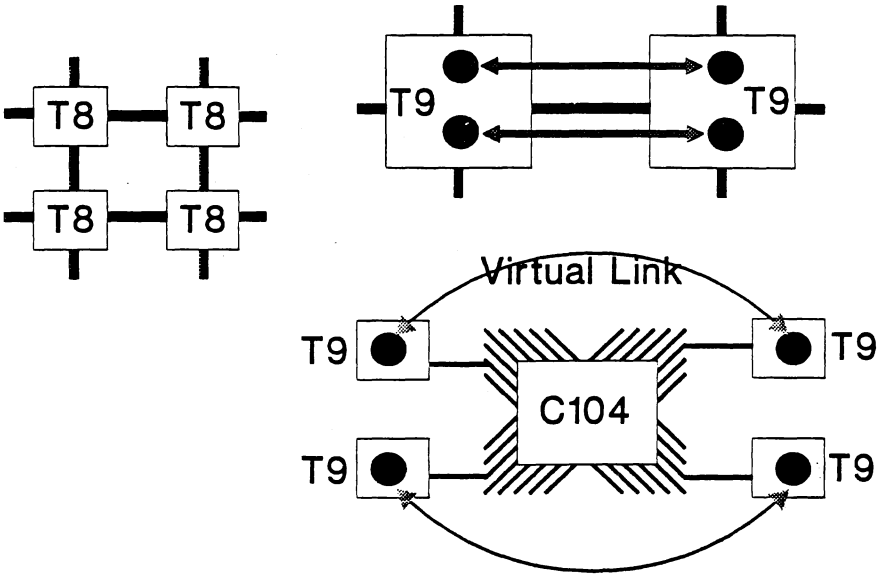

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

Control Link Commands

Resetting	*start*	Sets ID & channel header
	reset	Sets reset level
Configuring	*cpeek/cpoke*	Config. space setup
Booting	*identify*	Device revision number
	boot/bootdata	Load boot code
	run	Begin execution
	reboot	Boot from ROM
Reporting	*error*	Slave reports to master
Debugging	*stop*	Halt and analyze
	peek/poke	Read/write memory word

- All command messages are handshaken

SGS-THOMSON MICROELECTRONICS

Debugging support

- Breakpoints
  - j0 for G-process is T805 compatible
  - j0 for L-process invokes trap handler
  - j0 for P-process returns to supervisor
- Watch regions
  - Can be enabled for L- and P-processes
- Single stepping
  - Available for L- and P-processes

SGS-THOMSON MICROELECTRONICS

inmos

Debugging networks

Control links
- Independent debugging network
- Link speed programmable
- Ultimate post-mortem capability

Virtual channels
- Logically separate debugging networks

Protection and relocation
- Debugging monitors
- Operating systems

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

Enhanced kernel writing support

Interrupted process state access
- *ldshadow* load shadow registers from memory
- *stshadow* store shadow registers in memory

Scheduling queue manipulation
- *swapqueue* swap complete scheduler queue
- *insertqueue* insert at front of scheduler queue

Timer queue manipulation
- *swaptimer* swap complete timer queue

Timeslice control
- *settimeslice* enable/disable timeslicing
- *timeslice* deschedule to back of queue

SGS-THOMSON MICROELECTRONICS

inmos

Events

4 Event inputs & 4 Event outputs

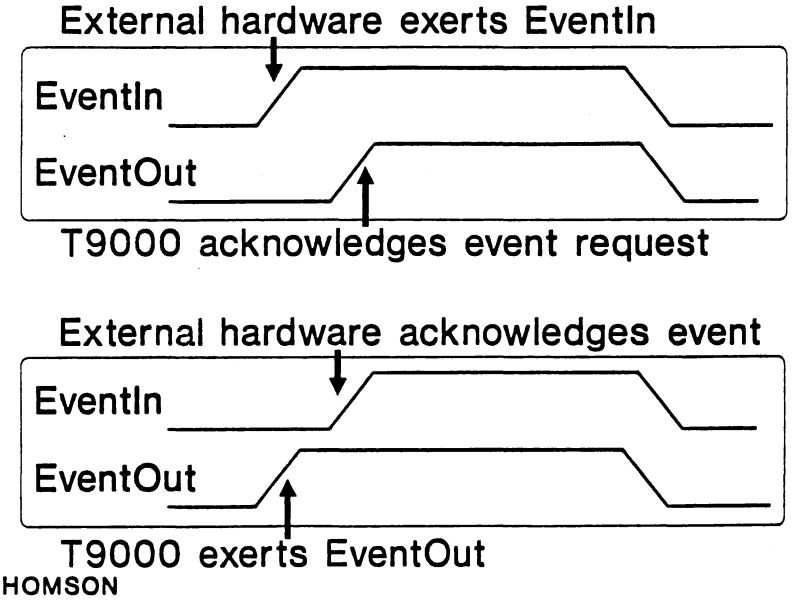

SGS-THOMSON MICROELECTRONICS

inmos

Enhanced process model - summary

- Virtual channels
- Resources
- Semaphores
- Per-process error handling
- Memory protection and relocation
- Enhanced support for debugging
- Enhanced real-time kernel support

SGS-THOMSON MICROELECTRONICS

inmos

SGS-THOMSON MICROELECTRONICS

inmos

T9000 Product Family

- Performance
- Responsiveness
- Connectivity
- System simplicity
- Software

SGS-THOMSON MICROELECTRONICS

inmos

Transputing '91

T9000 Software Presentation

Tony Debling

SGS-THOMSON MICROELECTRONICS