Chapter 1 : Transputer and OCCAM Engineering Series

Transputer Systems Ongoing Research edited by Alastair Allen (University of Aberdeen) WoTUG Proceedings of the 15th World Occam and Transputer User Group Technical Meeting.*

The T shared most features with the T, but moved several pieces of the design into hardware and added several features for superscalar support. Unlike the earlier models, the T had a true 16 kB high-speed cache using random replacement instead of RAM, but also allowed it to be used as memory and included MMU-like functionality to handle all of this termed the PMI. For more speed the T cached the top 32 locations of the stack, instead of three as in earlier versions. The T used a five-stage pipeline for even more speed. An interesting addition was the grouper [11] which would collect instructions out of the cache and group them into larger packages of 4 bytes to feed the pipeline faster. Groups then completed in one cycle, as if they were single larger instructions working on a faster CPU. The T also added link routing hardware called the VCP Virtual Channel Processor which changed the links from point-to-point to a true network, allowing for the creation of any number of virtual channels on the links. This meant programs no longer had to be aware of the physical layout of the connections. A range of DS-Link support chips were also developed, including the C way crossbar switch, and the C link adapter. It consistently failed to reach its own performance goal of beating the T by a factor of ten. The production delays gave rise to the quip that the best host architecture for a T was an overhead projector. This was too much for Inmos, which did not have the funding needed to continue development. By this time, the company had been sold to SGS-Thomson now STMicroelectronics , whose focus was the embedded systems market, and eventually the T project was abandoned. However, a comprehensively redesigned bit transputer intended for embedded applications, the ST20 series, was later produced, using some technology developed for the T ST20[edit] Although not strictly a transputer, the ST20 was heavily influenced by the T4 and T9 and formed the basis of the T, which was arguably the last of the transputers. The mission of the ST20 was to be a reusable core in the then emerging SoC market. The architecture was loosely based on the original T4 architecture with a microcode-controlled data path. However, it was a full redesign, using VHDL as the design language and with an optimized and rewritten microcode compiler. The project was conceived as early as when it was realized that the T9 would be too big for many applications. Actual design work started in mid Several trial designs were done, ranging from a very simple RISC-style CPU with complex instructions implemented in software via traps to a rather complex superscalar design similar in concept to the Tomasulo algorithm. The final design looked very similar to the original T4 core although some simple instruction grouping and a workspace cache were added to help with performance. Adoption[edit] While the transputer was simple but powerful compared to many contemporary designs, it never came close to meeting its goal of being used universally in both CPU and microcontroller roles. In the microcontroller market, the market was dominated by 8-bit machines where cost was the most serious consideration. Here, even the T2s were too powerful and costly for most users. This was excellent performance for the early s, but by the time the floating-point unit FPU equipped T was shipping, other RISC designs had surpassed it. Few transputer-based workstation systems were designed; the most notable likely being the Atari Transputer Workstation. The transputer was more successful in the field of massively parallel computing, where several vendors produced transputer-based systems in the late s. These controlled both the readout of the custom detector electronics and ran reconstruction algorithms for physics event selection. The ability to quickly transform digital images in preparation for print gave the firm a significant edge over their competitors. Within a few years, the processing ability of even desktop computers ended the need for custom multi-processing systems for the firm. Please help improve it by rewriting it in an encyclopedic style. Learn how and when to remove this template message Growing internal parallelism has been one driving force behind improvements in conventional CPU designs. Instead of explicit thread-level parallelism as is used in the transputer , CPU designs exploited implicit parallelism at the instruction-level, inspecting code sequences for data dependencies and issuing multiple independent instructions to different execution units. This is termed superscalar processing. Superscalar processors are suited for optimising the execution of sequentially

constructed fragments of code. Given these substantial and regular performance improvements to existing code there was little incentive to rewrite software in languages or coding styles which expose more task-level parallelism. Nevertheless, the model of cooperating concurrent processors can still be found in cluster computing systems that dominate supercomputer design in the 21st century. Unlike the transputer architecture, the processing units in these systems typically use superscalar CPUs with access to substantial amounts of memory and disk storage, running conventional operating systems and network interfaces. Resulting from the more complex nodes, the software architecture used for coordinating the parallelism in such systems is typically far more heavyweight than in the transputer architecture. The fundamental transputer motive remains, yet was masked for over 20 years by the repeated doubling of transistor counts. Inevitably, microprocessor designers finally ran out of uses for the greater physical resources, almost at the same time when technology scaling began to hit its limits. Power consumption, and thus heat dissipation needs, render further clock rate increases unfeasible. These factors led the industry towards solutions little different in essence from those proposed by Inmos. The most powerful supercomputers in the world, based on designs from Columbia University and built as IBM Blue Gene , are real-world incarnations of the transputer dream. They are vast assemblies of identical, relatively low-performance SoCs. Recent trends have also tried to solve the transistor dilemma in ways that would have been too futuristic even for Inmos. On top of adding components to the CPU die and placing multiple dies in one system, modern processors increasingly place multiple cores in one die. The transputer designers struggled to fit even one core into its transistor budget. Today designers, working with a fold increase in transistor densities, can now typically place many. One of the most recent commercial developments has emerged from the firm XMOS , which has developed a family of embedded multi-core multi-threaded processors which resonate strongly with the transputer and Inmos. The transputer and Inmos helped establish Bristol , UK, as a hub for microelectronic design and innovation.

Chapter 2 : Research - Bank of Canada

Get this from a library! Transputer systems, ongoing research: WoTUG, proceedings of the 15th World Occam and Transputer User Group Technical Meeting, th April , Aberdeen, Scotland, UK.

Chapter 3 : Research Programs & Contacts | NIDDK

Buy Transputer Systems Ongoing Research (WoTUG), (Transputer and Occam Engineering Systems, 25) on www.nxgvision.com FREE SHIPPING on qualified orders.

Chapter 4 : Statistics & Research | Children's Bureau | ACF

Transputer Systems - Ongoing Research by World Occam and Transputer User Group starting at $ Transputer Systems - Ongoing Research has 1 available editions to buy at Alibris.

Chapter 5 : Artificial Intelligence research at Microsoft

Transputer Systems Ongoing Research (Wotug Proceedings of the 15th World Occam and Transputer User Group Technical Meeting: Alastair Allen, World Occam and Transputer User Group Technical Meeting aberdeen: Books - www.nxgvision.com

Chapter 6 : Research | University of Pittsburgh

We work out the trending price by crunching the data on the product's sale price over the last 90 days. New refers to a brand-new, unused, unopened, undamaged item, while Used refers to an item that has been used previously.

DOWNLOAD PDF TRANSPUTER SYSTEMS, ONGOING RESEARCH

Chapter 7 : Transputer - Wikipedia

Transputer Systems Ongoing Research (WoTUG) Volume 25 Transputer and OCCAM Engineering Series Published in Transputer and OCCAM Engineering Series.

Chapter 8 : Clinical Trials - Clinical Trials - Mayo Clinic Research

Transputer Systems Ongoing Research, Aberdeen University, Aberdeen, Scotland, April 19, Occam and the Transputer - Research and Applications, University of.

Chapter 9 : Research & Innovation | Computer Science & Engineering

The transputer is a series of pioneering microprocessors from the s, featuring integrated memory and serial communication links, intended for parallel www.nxgvision.com were designed and produced by Inmos, a semiconductor company based in Bristol, United Kingdom.