Intel® Multi-Core™ Technology provides higher performance in PC-based automation technology

Powered up PC-based Control with Multi-Core™ Technology

A trend towards multi-processor systems is emerging for PC-based control technology. The operating system, PLC, Motion Control systems and HMI can now be distributed over different processor cores. How multi-processor technologies can affect new control concepts is explained by Ramon Barth, Manager Software Development for Beckhoff, and Sebastian Richter, Field Application Engineer for Intel GmbH in Germany.

Due to the widespread distribution and steadily growing processing power – with relatively low costs – PC technology offers an ideal platform for automation technology. It enables automation tasks to be solved by software, which previously could only be carried out at substantially higher costs with dedicated hardware. In addition, the software solution for automation is highly scalable: for example, the limitation on the number of controllable NC axes is only dependent on the available CPU processing power. Today, x86 systems with integrated hardware floating point units and applications with 50 to 100 servo axes are not uncommon. In this connection, the number of axes is only a parameter which can be configured by the user.

The basis of PC-based automation technology

The core principle of Beckhoff "New Automation Technology" is PC-based control technology. As a result, all control system and visualization tasks are executed by a powerful central CPU and decentralized (non-intelligent) I/Os. Because as a rule in somewhat more complex applications, the PC is used as a visualization "front-end", which is not utilized to full capacity by this task, the idea arises to let the same PC calculate automation tasks such as PLC and Motion Control. For this purpose, the Windows operating system is extended by an integrated, lean real-time system which looks like a normal driver to the operating system. Automation

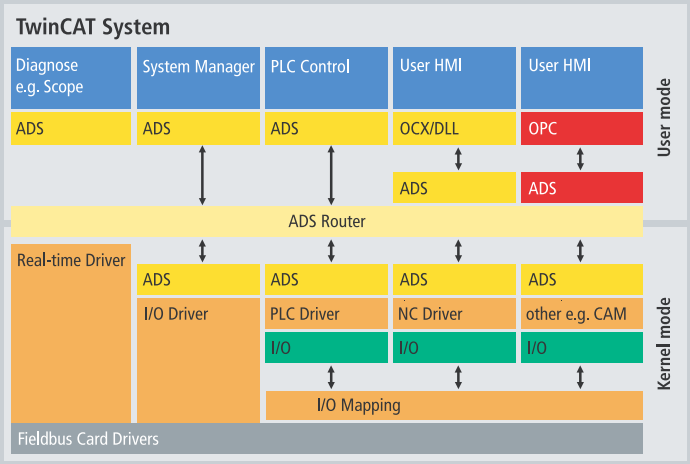

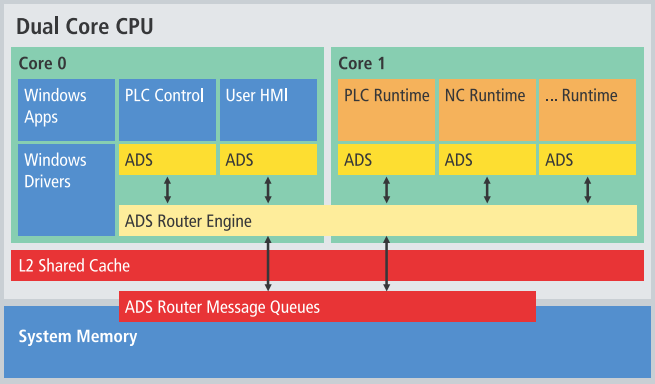

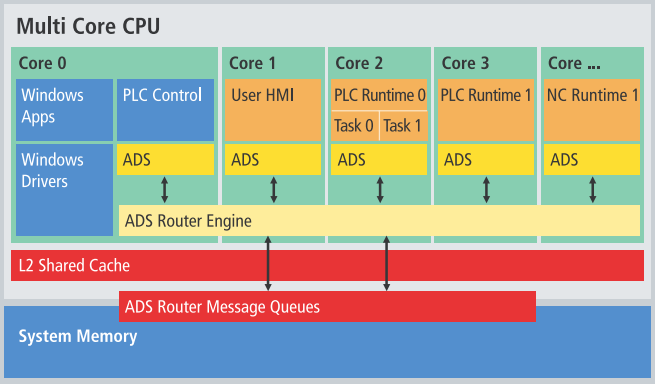

processes run within the real-time system as tasks, e.g. PLC, Motion Control, cam controller and linear path control. Access to the required periphery takes place via familiar fieldbus systems as well as via Ethernet. The Ethernet-based fieldbus system EtherCAT eliminates the "bottleneck" previously caused by traditional fieldbuses. The modular architecture of PC-based control technology also satisfies the requirements for the most effective utilization of multi-core processors.

Intel® Multi-Core Technology

In recent years, there has been a neck-and-neck race between Intel® and their competitors as to who could bring the fastest processor onto the market. This vigorous competition over the desktop market resulted in an increase in power consumption for industrial applications. New processor generations are going down the route of parallelism, in which several processor cores are integrated into one package with reduction of the operation frequency in some cases.

The introduction of multi-core systems led at first to a reduction in the clock frequency, which superficially had a negative effect on the total power of the system because programs running in series are not executed as quickly due to the lower clock frequency. In order to extract the most benefit from the new computer architecture, a new way of thinking is called for from the application development engineers: What had been executed as fast as possible in series up until now, needs to run simultaneously on both cores in the future.

Using skilled functional or data-based division of tasks, an automation application can be obtained which supports a noticeable acceleration above the 50 % mark. In order to achieve only an approximate growth in performance through an increase in frequency, the IPC manufacturer has to accept high power consumption and invest in considerably higher-priced products.

In 2006 with the second generation of Dual-Core™ products, Intel® introduced the "Core™ Microarchitecture", known under the label "Core™2 Duo". Due to the reduction of the internal feature size to 65 nm and fundamental changes to the architecture – compared to Netburst or Yonah – the Dual-Core™ processors underwent a further power boost alongside greatly reduced power consumption.

New power potential was therefore available to the Industrial PC, which until then had been reserved for the server segment and completely new application areas opened up.

Practical application of dual-core technology

A further innovation enabled by the dual-core system is the physically parallel execution of different functions which previously could only be carried out quasi in parallel on one system – with the disadvantage of mutually influencing the processing power.

A good example is supplied by the allocation of resources on the control system and the HMI. The former is executed on a core under an embedded real-time system in order to preserve the real-time capability and to exclude any influence exerted by a graphic-based output. Conversely, the non-real-time application also benefits from its own CPU core, to which it has exclusive access. The second core provides a general operating system platform for the HMI, in which the graphical options rather than the deterministic features are the primary focus. This means that e.g. the image switch-over times on the user interface are shorter and the compiling times for the PLC project, when compiled on the target system, are significantly reduced. Where a 2.8 GHz P4 system at 50 % utilization by the real-time application needs approx. 6 minutes for a completely new compilation of a 4.5 MB PLC project, a 2.16 GHz Core™2 Duo system only needs 40 s for the same project. For the user, the TwinCAT system makes clear use of the improved performance as described.

On single-core systems the TwinCAT real-time functions are executed deterministically with the maximum priority quasi in parallel. Due to the patented switch-over process between Windows and real-time applications, time-critical Windows

functions are processed at the right time in the time base of the real-time system. However, an increased processor power consumption from the real-time applications is at the expense of the performance of the Windows applications. Communication between the two worlds takes place via the message-based TwinCAT ADS communication system; in this case, the TwinCAT ADS router synchronizes the process communication.

Dual-core systems are automatically recognized by TwinCAT: One of the two cores (see Fig.: TwinCAT dual-core support) is occupied by the real-time functionality, whereas the core computing time not used by the real-time core is available for Windows. The ADS messaging system in turn is responsible for the synchronization between Windows applications and the real-time applications.

For optimum use of the dual-core CPUs, the TwinCAT real-time system and ADS router were upgraded – while ensuring compatibility with previous versions. Because both CPU cores access the same memory and a common L2 cache, the ADS router only has to be supplemented by the multi-processor synchronization in order to distribute the ADS messages safely to the registered processes. The result for the user consists of a higher processing power for the entire application with minimum migration expense.

In the future, over and above dual-core CPUs, there will also be four-core or eight-core systems available at a reasonable cost. Software-based solutions benefit here again, because these are able to allocate tasks depending on the number of CPU cores available. Moreover, it requires a considerably lower effort in order to carry out a functional division. Functional units can be allocated to dedicated cores. The TwinCAT system from Beckhoff will considerably simplify the use of multi-core systems for the user via appropriate configuration and diagnostic tools. In the TwinCAT System Manager, for example, run-times for real-time tasks can be monitored and priorities or task processing sequences can be configured manually.

Using "Load Balancing" in the standard operating mode, the task to be activated is either assigned a free core or – depending on the priority – the optimum core. Tasks can also be statically allocated to a core using configurable core affinities. In this way, a traditional division into PLC or NC run-time systems can again be produced by making use of ready-made profiles. As the TwinCAT user today already thinks in PLC tasks and run-time systems, the transition from single or dual-core systems to multi-core systems will also take place smoothly.

Integrated IPC series using Core™ Duo/Core™2 Duo

As the company has already done with Pentium® M technology, Beckhoff has transferred Intel® Core™ Duo/Core™2 Duo processors onto its own motherboards. This ensures that a chosen chipset and type of processor will be gradually introduced into all Beckhoff IPC series. The user can select the most appropriate IPC for their application from up to 13 different device families, equipped with Core™ Duo performance, irrespective of whether the requirement is for a flexible Panel PC C3640 with numerous free PCI and PCIe slots, or for an ultra compact control cabinet Industrial PC such as the C6920.

In addition to the increase in performance, Core™ Duo/Core™2 Duo technology – and the 945 chipset associated with it – has further benefits for the user: The low power consumption of the Core™ Duo/Core™2 Duo processors results in extremely stable systems, thermally. In other words, the maximum ambient temperature for the IPC can be up to 55 °C (131 °F). In addition, the Intel® chipset 945 with its ICH-7R used by Beckhoff offers an on-board SATA RAID 1 controller, which enables reliable and rapid data backup. All that is required here is an IPC that integrates two hard disks. Here again, the ideally suited Beckhoff product range offers benefits for its users.

In addition, with ADD2, Beckhoff is offering a PCI express plug-in card with two DVI monitor connectors. This means that two displays can be connected as standard, enabling clone, extended and twin modes.

As users are going for the newest technology with the Core™ Duo/Core™2 Duo processors, a high degree of long-term availability is guaranteed in addition to all the high-end technical features.

→ www.beckhoff.com/IPC

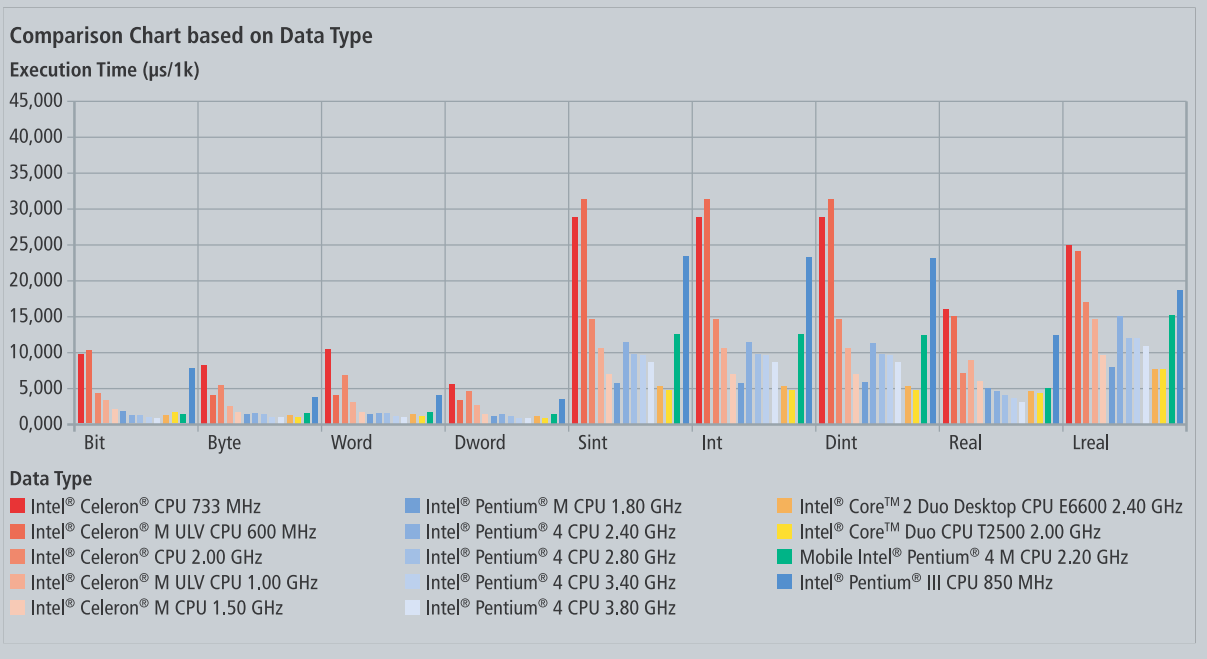

CPU benchmark comparison between special operation types:

Operation types (data type)	Command
Bit operation (Bit, Byte, Word, Dword)	AND, OR, XOR, NOT
Integer operation (Sint, Int, Dint)	ADD, SUB, MUL, DIV
Real operation (Real, Lreal)	ADD, SUB, MUL, DIV

Demands on the application development engineers

For the developers of real-time or PLC applications in the TwinCAT system environment, the change-over from single-core systems to dual-core systems is seamless. The real-time run-time environment continues to use only one CPU so that existing PLC projects can be taken over one-to-one.

As TwinCAT makes unused CPU time available for Windows applications, the Windows operating system sees two CPUs, of which one is running at part capacity. Windows applications built from several program threads can profit from this. The Windows operating system distributes the application threads to the available CPUs. These threads run physically in parallel and the CPU hardware is used optimally. However, synchronization gaps present in the application occur more readily in physical parallel processing than in the quasi-parallel execution of threads.

In order to use multi-core systems optimally in the future, all applications need to be divided in a modular way into threads or tasks as far as possible. This gives both Windows and TwinCAT the possibility of dividing the execution of program parts as best as possible on the available CPU cores (i.e. load balancing). Monolithic programs will still continue to work, but they will only be able to use the existing computing capacity to an ever-decreasing degree. For the development of applications with several threads, attention should be paid to a clean synchronization of the different program parts when accessing common system resources. For this purpose, each operating system provides appropriate mechanisms such as semaphores, critical section, etc. Even for experienced development engineers, the physically parallel execution of program parts calls for a partial shift in thinking. In real-time systems with strict priority control, the highest priority thread runs to its end without being interrupted. On a single-core system, this feature justifies the assumption that during the run-time of the high priority threads no lower priority thread is active. However, this no longer applies to multi-core systems. While the highest priority thread in a core is being executed, other lower priority threads in other cores can be processed at the same time. The disabling of interrupts is likewise no longer an appropriate synchronization method in multi-core CPUs, because blocking only applies to the particular core being used.

More processing power thanks to new CPU architecture

But even considering only one core, each new CPU architecture displays a steeply increasing processing power graph. Beckhoff carries out a PLC benchmark (cf. PLC benchmark) for every CPU installed, which tests the execution times of important IEC 61131 PLC operations. What is noticeable here is the difference in performance between the Netburst generation and CPUs which are based on "Core™ Microarchitecture" (this applies to a lesser degree for Pentium® M products). The effect of the processor architecture illustrates the fact that Core™ Microarchitecture CPUs, despite significantly lower clock frequencies, achieve equal or better performance values in tests. The performance of a CPU cannot be based only on the designated clock frequency.

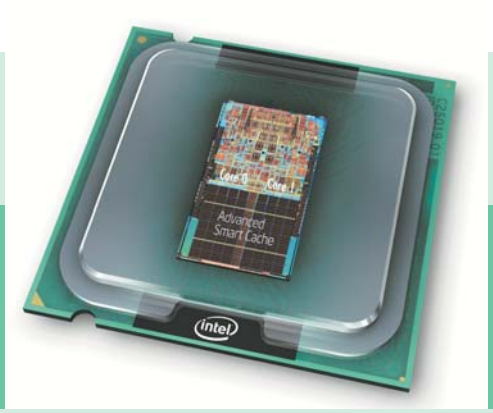

Intel® Core™ Microarchitecture

With the launch of the Woodcrest, Conroe and Merom products, Intel® has introduced the new Core™ Microarchitecture to all segments of the embedded roadmap, from the performance line via the scalable line to the low-power line.

Since the Netburst architecture is out of the question for a new CPU generation for the reasons given above, the Core™ Microarchitecture combines the benefits of Netburst and the mobile Yonah architecture and goes far beyond this with additional improvements. The most important innovations and their effects on system performance are:

Wide Dynamic Execution

Since the introduction of "Dynamic Execution" in the P6 architecture (Pentium® II/Pentium® III) this has been expanded in the course of Netburst architecture to "Advanced Dynamic Execution". Among these are technologies such as data flow analysis, the speculative implementation of instructions as well as what is known as out-of-order execution. This means that commands are not necessarily executed one after another, but in a sequence that makes best use of the processor.

By means of an extended "branch prediction," the prediction precision is increased, thus preventing reloading of instructions due to incorrect predictions.

The Core™ Microarchitecture increases these capabilities, in that it loads, decodes and executes more instructions per cycle (four in comparison to three instructions in previous architectures).

In order to push energy efficiency further (performance/watt) these CPUs are capable, using what is known as "macrofusion", of transferring very frequently occurring x86 instructions into a micro operation and so saving cycles during execution.

In this way, the decoder creates e.g. a single micro operation (CMPJNE) from a comparison (CMP) and the ensuing conditional jump (JNE). The ALUs were expanded accordingly and can now deal with these combined commands.

Advanced Smart Cache

An important innovation in the processor architecture concerns the design of the L2 cache. While in other dual/multi-core designs, each core has a separate L2 cache, both cores of a Core™2 Duo CPU, for example, have a common L2 cache.

This structure is ideal for supporting software designs where the application threads are running on both cores. Each part of this application needs access to the data for processing – ideally, those which are already available in the L2 cache. As a result of joint utilization, each core has complete access to the cache and can see the data. If the caches are separate, then time-consuming transfers between the memories are necessary.

Smart Memory Access

The improvements already described do not take effect if the information is not in the right place at the right time. In this context, a good prediction of the processing units naturally plays a part: but the effect of the memory and bandwidth management is far greater.

Smart Memory combines technologies which should prevent a delay in access to the memory buses through optimal utilization of the bandwidth. A new feature here is what is known as "memory disambiguation," which enables the CPU to make an intelligent decision as to whether a Load can take place before a Store.

In old processor architectures this was strictly forbidden in order to guarantee data coherency, as otherwise data could be loaded that were already out-of-date and would be overwritten with the next Store. This means that an out-of-order processing of commands is possible to an increasing degree, which in turn saves time and therefore energy. Even a miscalculation of this algorithm does not lead to an invalidation of the data, but to a reloading with the correct data.

A further feature of the new processor architecture deals with the problem of memory bandwidth. The Intel® Core™ Microarchitecture introduces a total of eight so-called "Prefetchers," which ensure that the data are loaded as early as possible and at a time when the bus is not too heavily loaded. Each core of a dual-core CPU possesses three of these units (two for data, one for instructions): the latter two fill the divided L2 cache with data.

Taken together, all these changes in the CPU architecture cause an increase in performance of up to 90 %, with up to 40 % reduced power consumption. Due in particular to the intelligent memory subsystem and the more efficient utilization of the FSB, no latency problems whatsoever occur with critical real-time applications in automation technology.

Software which is designed for this processor architecture has a guaranteed future, as it is compatible with the forthcoming generation of multi-core processors and, due to its modular configuration, offers more functions or can provide functions with more resources.

→ **www.intel.com**