Methods of optimizing power consumption.

The dynamic power consumption of the processor can be reduced by reducing the clock frequency or the CPU supply voltage. This method is called "dynamic scaling of voltage and frequency". The supply voltage and clock frequency of the processor are reduced in a software way so that the SHIELD operates at the minimum speed sufficient for the execution of the application or task. When performing tasks that are not demanding on speed, the work slows down.

You can reduce the power consumed by the device in several ways:

- by reducing the supply voltage;

- clock frequency reduction;

- changes to the cache structure;

- reducing the length of interconnections;

- using conveyor mechanisms.

To avoid errors or delays, a scheduler is used that determines when each application request (task) should be executed by the processor, and at what intervals the work can be slowed down to reduce consumption.

Task priorities can be set statically (fixed priorities) or dynamically (priorities change from request to request). Dynamic priorities are used, for example, in EDF algorithms, when the CPU first executes the task with the earliest deadline.

Static priorities are good for recurring tasks and even scheduling.

Blocking.

Dynamic power consumption depends on the number of switching elements. Accordingly, it can be reduced either by reducing the number of logic elements in the circuit or by reducing the clock frequency. At the same time, the speed of calculations increases, which also reduces energy consumption.

Usually, 20-40% of the total consumption is accounted for by the charging scheme. Gating is used to turn off unused elements.

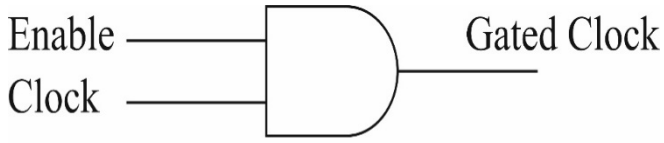

Gating of a logical element.

Gating algorithms are divided into three groups: system, sequential, and combinational. System gating blocks clocking for the entire unit and disables it. The other two types stop the timing signal only for some parts of the block. The unit itself continues to work.

Power supply gate.

This is one of the most effective methods of reducing both subthreshold leakage and gate leakage at the same time, as the supply to the element is completely blocked. Gain in dynamic power occurs due to the use of elements with a low ultimate voltage. At the same time, the static power consumption may increase.

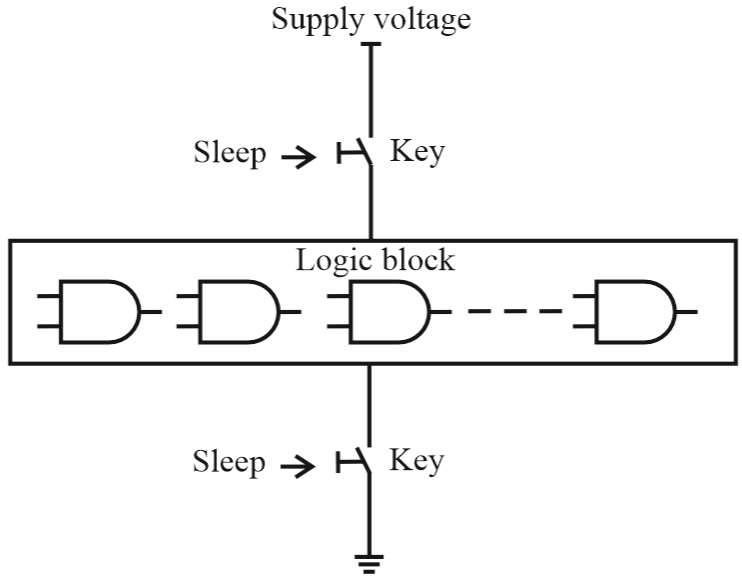

A simplified diagram of a logic block with power gating.

Power reduction methods at the component level.

Cache.

On-chip cache modules dissipate 25-50% of power in various processors. This value can be reduced in several ways:

- improvement of the technological process;

- changing the structure of the memory cell;

- voltage reduction;

- structure optimization.

Adding buffers.

The most common approach to reducing cache power consumption is to add small cache or buffer to fetch data from it without accessing the general cache.

Tires.

As the size of the elements decreases, the power consumption of the conductors and the delays introduced by them become increasingly important. This problem is most acutely manifested in bus connections, since they can dissipate up to a third of all consumed power.

There are several approaches to optimizing bus power consumption:

- reducing the amplitude of the signal;

- use of codes that ensure the minimum number of switche;

- localization of connections.

To reduce the amplitude of the signal, you should use a level converter, which reduces the amplitude of the input signal (from the transmitter), and then does the reverse in the receiver.

The longer the conductor, the greater its parasitic capacitance. Accordingly, since the length of the links cannot always be reduced, it is necessary to reduce the number of switches in such sections by means of signal coding.

There is another approach - dividing the bus into segments using buffers. Energy is consumed only by active sections of the bus. In addition, due to the reduction of the length of the connections, the parasitic effects are weakened, which also allows to reduce the consumption.

Self-synchronous processors.

Self-synchronous circuits are among the best in terms of power control. Thanks to the absence of a charging scheme (this is one of the main sources of power dissipation), it is possible to significantly reduce energy consumption. In addition, the issue of dephasing of clock pulses is removed and noise is reduced.

One of the first examples of a self-synchronous processor is AMULET. This is the first development of APM in this region (1993). The processor consists of several functional blocks that work independently and in parallel. Other processors work according to the same scheme, for example, SNAP, ARM996HS, HT80C51.

Another approach is GALS (globally asynchronous and locally synchronous) - asynchronous circuits with local synchronization. It combines the advantages of both asynchronous and synchronous circuits, in addition to their disadvantages.

The GALS device is divided into three types:

- clock generators with the possibility of stopping (pausible-сиоскХ

- FIFO buffers

- limit synchronization schemes.

Architectures of self-synchronous processors.

RISC processors with a reconfigurable instruction set.

RISC processors have the task of reducing one-time costs for design and implementation in production. Processors are economical in terms of consumption and have high speed.

ASIP processors with a specialized set of commands.

ASIP processors are a cross between general-purpose processors and specialized ultra-large ICs. They have a set of built-in operations. Building an optimal set of extended instructions is the main task of the developer, the performance of the processor depends on its solution.

Commercial ASIP processors include Extensa Processors (Tensillica), NIOS (Altera), Microblaze (Xilinx).

MISC-processor without a set of instructions.

Another approach is implemented in processors that do not have a set of instructions. The data tract is created for a specific application, and the code is compiled not in instructions, but in control words.

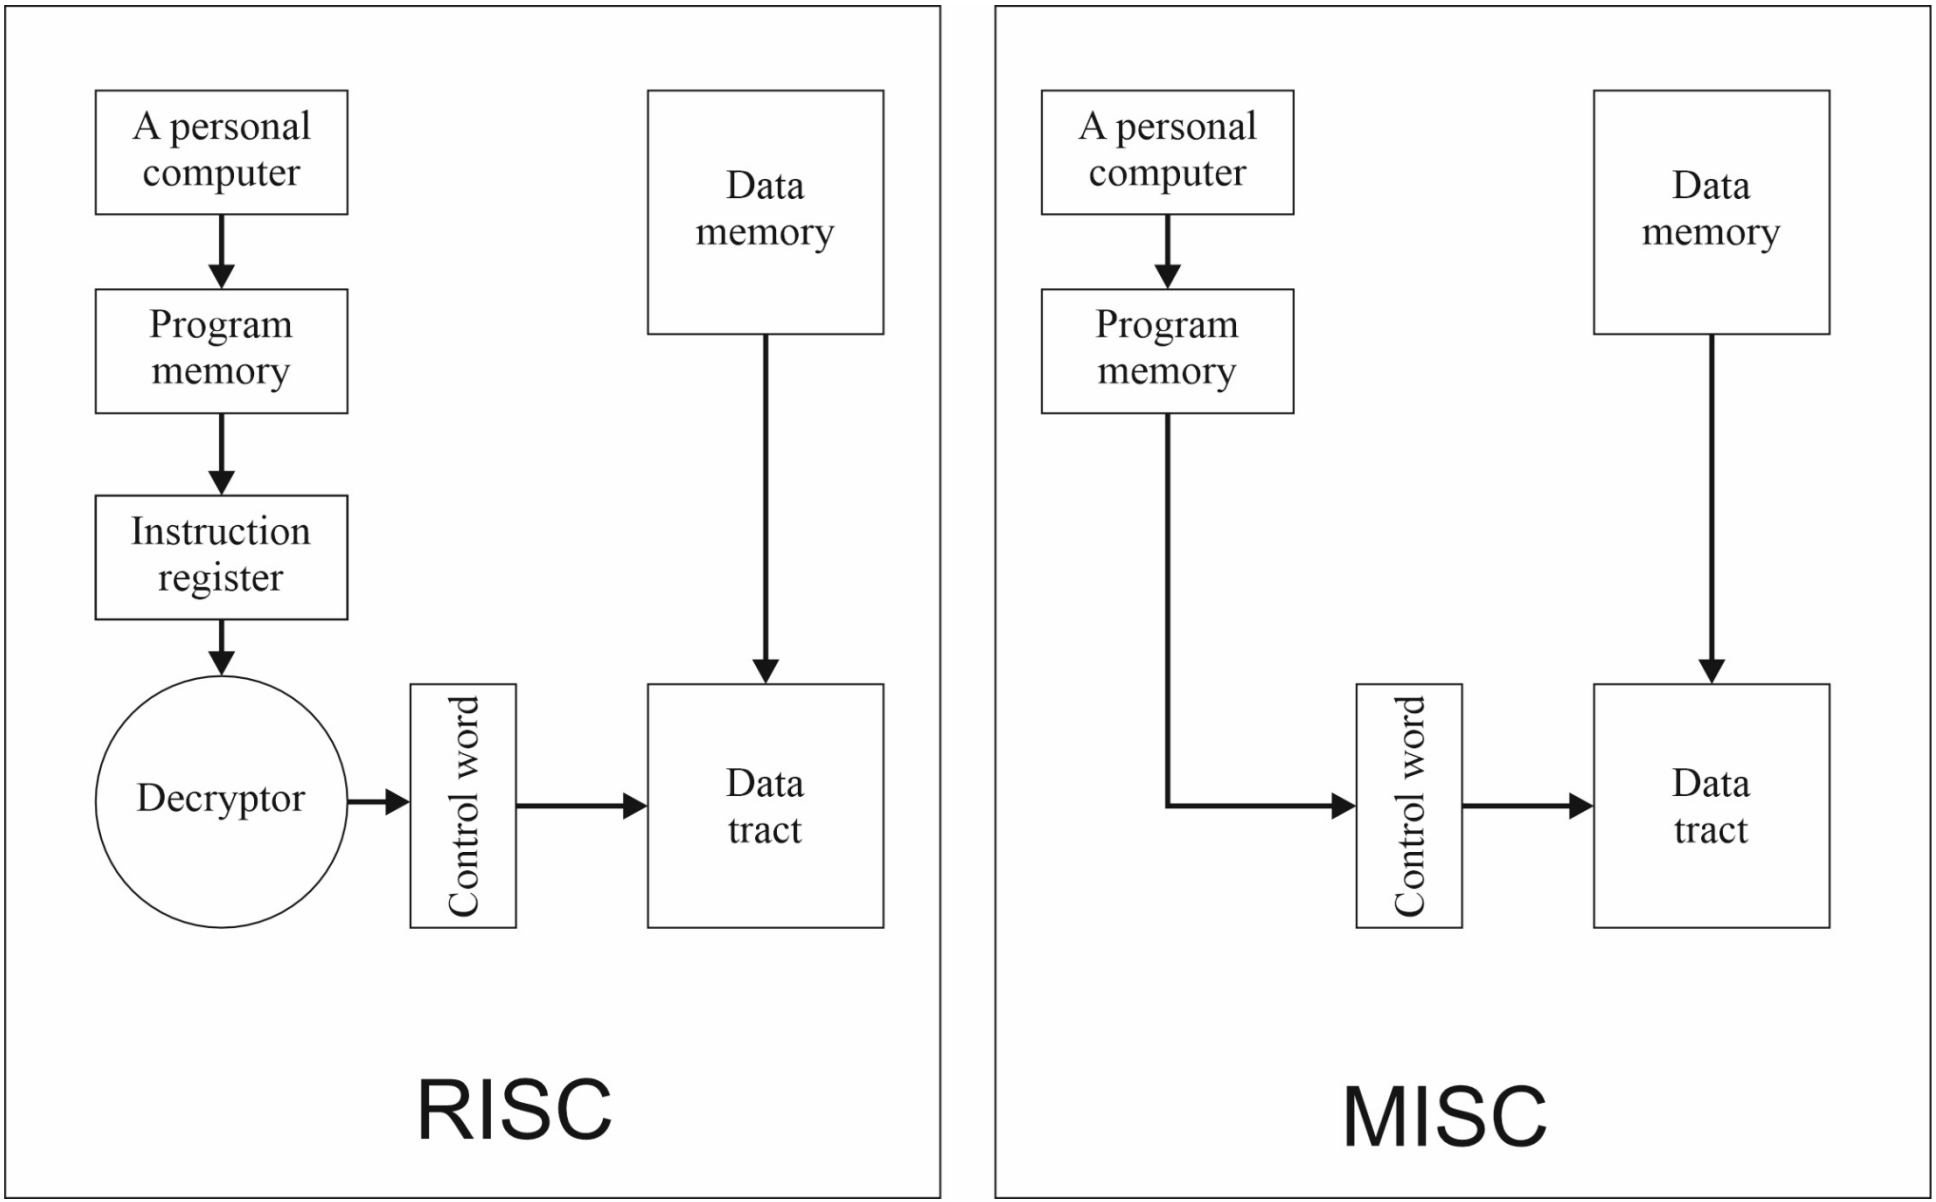
A personal computer
Program memory
Instruction register
Decryptor
Control word
Data tract
Data memory
RISC
A personal computer
Program memory
Control word
Data tract
Data memory
MISC

Most often, when designing embedded systems, the developer faces a difficult choice between single-board computers (SBC - single board computer) and processor modules (COM - computer on module).

Today, computing systems have penetrated all spheres of human life: high-performance embedded systems that are used in completely different areas, starting from the control of production lines and ending with medical equipment. In order to make an informed choice in favor of one or another system, it is necessary to take into account the performance and architecture of the processor, implemented interfaces, consumed energy, used software, cost and development time. Developers must navigate the technical and organizational and commercial issues affecting design in order to choose the optimal solution.

Single-board computers are a ready-made solution that allows you to eliminate the stage of development and production in the case of using processor modules to create a carrier board, therefore, the developed systems concentrate only on software issues.

The RS/104 standard allows you to avoid the need for your own development of a carrier board and to configure the computer more optimally, using modules from different manufacturers. It should be noted that during the long existence of the RS/104 standard, many manufacturers have developed a huge number of modules with different purposes and functions. Additional freedom from a specific computer board is provided by the fact that in the PC/104, the interface connectors on the case are connected to the computer board or expansion module with cables. A limitation in the use of PC/104 single-board computers is the use of low-power processors. The power consumption of a PC/104 single-board computer should be no more than 25 W.

Thanks to the standardization of the modules, there is a possibility of simple modification of the system by replacing the modules on the carrier board. Thus, it is possible to fine-tune the system according to the customer's requirements and release versions of the system with more or less productive processors, or even change the processor architecture by simply replacing the module. Interchangeability of modules allows for easy system recovery when new processors are introduced, which also extends the product life cycle, makes it free from module obsolescence and phase-out, and eliminates dependency on a single manufacturer.

Types of processor modules.

Reducing the power consumption of computer hardware

Ensuring processor energy efficiency

At the beginning of the microprocessor era, most processors used 5V (later 3.3V, 3.5V) for the processor and I/O circuits. However, the 8080A microprocessor had +5, -5.-12V.

Currently, PC equipment is considered within the framework of the chipset installed on the motherboard, which is made according to the standard that specifies the overall dimensions of the product and describes the technical parameters (shape, types of elements, their position and orientation).

Ensuring energy efficiency of memory

To ensure the energy efficiency of the dynamic RAM, a reduction in the supply voltage is used.

Supply voltage DDR = 2.5 V.

Supply voltage DDR2 (double-data-rate four synchronous dynamic random access memory) = 1.8 V.

DDR3 supply voltage = 1.5V.

In 2010 p. ASUS introduced a new Dual Intelligent Processors technology, which, according to the announcement, can instantly speed up the PC by 37% or reduce its power consumption by 80%. This technology at the hardware level is implemented in two ways: Turbo Processing Unit (TPU) and EPU, which are responsible for overclocking and energy efficiency, respectively.

Ultra Durable technology (from the Gigabyte company) has been developed - to improve the temperature regime and reliability of the motherboard, which provides:

- an increased thickness of copper layers with a thickness of 70 microns for both the power supply layer and the grounding layer of the system board, which reduces the total resistance of the board by 50%, and ensures a decrease in the operating temperature of the computer, an increase in energy efficiency, and an improvement in the stability of the system during overclocking;

- the use of field-effect transistors with reduced resistance in the open state (RDS(on)), because the transistors of the +12 V power converters emit a lot of heat;

- use of chokes with a ferrite core - they provide lower energy losses and a lower level of electromagnetic radiation;

- use of lead-free solder; reuse of cardboard and plastic packaging (not electrical efficiency, but ecological).

ACPI (Advanced Configuration and Power Interface) is an open standard (released in 1996). defines a common interface for hardware discovery, power management, and motherboard and device configuration.

ACPI distinguishes the following basic system states:

- G0 (S0) (Working) - normal work;

- G1 (Suspend, Sleeping, Sleeping Legacy) - the machine is turned off, but the current system context (system context) is saved, work can be continued without rebooting. For each device, the degree of information loss is determined, as well as the places where the information should be stored, where it will be read upon awakening, and the time of awakening;

- G2 (S5) (soft-off) - soft (program) shutdown; the system is completely stopped, but energized, ready to turn on at any moment, the system context is lost;

- G3 (mechanical off) - mechanical shutdown of the system, the power supply unit is disconnected.

There are four processor operating states:

- C0 - operational (working) mode:

- C1 (Halt) - a state in which the processor does not execute instructions, but can immediately return to the working state:

- C2 (Stop-Clock) - a state in which the processor is detected by applications, but it takes time to switch to working mode;

- S3 (Sleep) - a state in which the processor disables its own cache, but is ready to transition to other states.

There are 4 sleep states:

- S1 - a state in which all processor caches are reset and processors have stopped executing instructions. However, CPU and RAM power is supported; devices that are not needed can be disabled

- S2 - an additional deeper state of sleep than S1, when the central processor is disabled while not in use;

- S3 "Suspend to RAM" (STR) in BIOS, "standby mode" ("Standby") - in this state, power continues to be supplied to the random access memory (RAM), and it remains almost the only component that consumes energy. Since the state of the operating system and all applications, open documents, etc. is stored in the RAM, the user can resume work exactly where he left it - the RAM state when returning from S3 is the same as before entering this mode.

The S3 has two advantages over the S4: the computer returns to working condition faster, and. if a running program (open documents, etc.) contains sensitive information, this information will not be forcibly written to disk. However, disk caches can be flushed to disk to prevent data integrity violations in the event that the system does not wake up, such as due to a power failure;

- S4 ("Sleep mode" (Hibernation) in Windows, "Safe Sleep" in Mac OS X, also known as "Suspend to disk") - in this state, the entire contents of the RAM are stored in the non-volatile memory yati (on the hard disk): the state of the operating system, all applications, open documents, etc. This means that after returning from S4, the user can pick up where they left off, similar to S3 mode.

States of operation of other devices (monitor, modem, buses, network cards, video card, disks, etc.):

- D0 - fully working (operational) state, the device is turned on;

- D1 and D2 - intermediate states, the activity is determined by the device;

- D3 - the device is turned off.

Windows 7 supports "Hybrid sleep mode", which combines the advantages of S1/S3 (wake-up speed) and S4 (protection against power failures). It is also implemented in Linux and Mac OS X.

There is an international initiative - the Green500 rating, which evaluates supercomputers by the MFLOPS/W indicator, based on the amount of electricity required to perform a fixed set of tasks.

The Microsoft Sleep Proxy system will be able to reduce the energy consumption of computers by 60-80%.

The Sleep Proxy system currently being developed by Microsoft, as it comes out, could become a greater boon for firms and companies that use a large number of computer equipment and allow them to save millions of dollars a year. In addition to saving electricity and finances, it should be noted that the use of the new system will have a favorable effect on the environment as well, reducing the amount of harmful substances emitted into the atmosphere by power plants due to the reduction of consumed energy.

Embertec power management system - a new level of energy saving.

The Embertec company has developed a new, unique and innovative technology for controlling the power supply of household appliances and other electronic devices, thanks to which it is possible to significantly reduce electricity consumption. Embertec devices can control the power supply of devices such as DVD players, televisions, computers and computer peripherals, most of which consume some, however small, amount of electricity even when turned off. The developed technology is intelligent and self-learning, it accumulates and takes into account during its work data related to typical situations of interaction between people and controlled electronic devices.