CLAIMS

What is claimed is:

1. A scheduling method, wherein the method is performed by a computing device, the computing device comprises a plurality of processors, a type of at least one instruction set of instruction sets supported by at least one of the plurality of processors is different from a type of an instruction set of instruction sets supported by another processor, and the method comprises:

obtaining a type of an instruction set of an application;

selecting a target processor from the plurality of processors, wherein the type of the instruction set of the application is a subset of types of a plurality of instruction sets of instruction sets supported by the target processor;

allocating the application to the target processor for execution;

wherein obtaining a type of an instruction set of an application comprises:

when the application is compilable, directly obtaining, by using the compilation option, the type, specified by a user, of the instruction set for executing the application; or first obtaining, by using the compilation option, a type, specified by a user, of a processor for executing the application, and then determining the type of the instruction set of the application based on a relationship between the type of the processor and the type of the instruction set; and

when the application is not compilable, obtaining the type of the instruction set of the application by using an abnormal instruction.

2. The method according to claim 1, wherein the target processor belongs to a target scheduling domain, instruction sets supported by the target scheduling domain is the same as the instruction sets supported by the target processor, and before the selecting a target processor, the method further comprises:

selecting the target scheduling domain from a plurality of scheduling domains, wherein the type of the instruction set of the application is a subset of types of a plurality of instruction sets in the instruction sets supported by the target scheduling domain.

3. The method according to claim 2, wherein

the target scheduling domain comprises only the target processor; or the target scheduling domain comprises the target processor and another processor, and instruction sets supported by the another processor is the same as the instruction sets supported by the target processor.

4. The method according to claim 1, wherein the application comprises at least one task, and

the method further comprises:

writing a type of an instruction set of the task into a task control block of the task, wherein the type of the instruction set of the task is the same as the type of the instruction set of the application; and

the obtaining a type of an instruction set of an application specifically comprises: obtaining the type of the instruction set of the task from the task control block.

5. The method according to claim 4, wherein the allocating the application to the target processor for execution specifically comprises:

allocating the task to the target processor for execution.

6. The method according to claim 5, wherein the target processor comprises at least two processing units, and the allocating the task to the target processor for execution specifically comprises:

when the target processor supports simultaneous multithreading processing, allocating the task to one processing unit of the target processor for execution;

when the target processor does not support simultaneous multithreading processing, allocating the task to one processor core of the target processor.

7. The method according to claim 5, wherein the target processor belongs to the target scheduling domain, the instruction sets supported by the target scheduling domain is the same as the instruction sets supported by the target processor, and the method further comprises:

when the target processor fails to execute the task, storing an identifier of the target scheduling domain in a first storage unit of the task control block of the task, wherein the identifier stored in the first storage unit indicates that the processor comprised in the target scheduling domain is not allocated to execute the application.

8. The method according to claim 5, wherein the target processor belongs to the target scheduling domain, the instruction sets supported by the target scheduling domain is the same as the instruction sets supported by the target processor, and the method further comprises:

when the target processor successfully executes the task, storing the identifier of the target scheduling domain in a second storage unit of the task control block of the task, wherein the identifier stored in the second storage unit indicates that the processor comprised in the target scheduling domain can be allocated to execute the application.

9. The method according to any one of claims 3 to 8, wherein the task is a process or a thread generated by the application.

10. A computing device, wherein the computing device comprises a plurality of processors, a type of at least one instruction set of instruction sets supported by at least one of the plurality of

processors is different from a type of an instruction set of instruction sets supported by another processor, wherein

a main processor in the plurality of processors is configured to:

obtain a type of an instruction set of an application;

select a target processor from the plurality of processors, wherein the type of the instruction set of the application is a subset of types of a plurality of instruction sets of instruction sets supported by the target processor;

allocate the application to the target processor for execution;

wherein the main processor is configured to:

when the application is compilable, directly obtain, by using the compilation option, the type, specified by a user, of the instruction set for executing the application; or first obtain, by using the compilation option, a type, specified by a user, of a processor for executing the application, and then determine the type of the instruction set of the application based on a relationship between the type of the processor and the type of the instruction set; and

when the application is not compilable, obtain the type of the instruction set of the application by using an abnormal instruction.

11. The computing device according to claim 10, wherein the target processor belongs to a target scheduling domain, instruction sets supported by the target scheduling domain is the same as the instruction sets supported by the target processor, and before the selecting a target processor, the main processor is further configured to:

select the target scheduling domain from a plurality of scheduling domains, wherein the type of the instruction set of the application is a subset of types of a plurality of instruction sets in the instruction sets supported by the target scheduling domain.

12. The computing device according to claim 11, wherein

the target scheduling domain comprises only the target processor; or the target scheduling domain comprises the target processor and another processor, and instruction sets supported by the another processor is the same as the instruction sets supported by the target processor.

13. The computing device according to claim 10, wherein the application comprises at least one task, and the main processor is further configured to:

write a type of an instruction set of the task into a task control block of the task, wherein the type of the instruction set of the task is the same as the type of the instruction set of the application;

obtain the type of the instruction set of the task from the task control block; and

allocate the task to the target processor for execution.

14. The computing device according to claim 13, wherein the target processor comprises at

least two processing units, and the main processor is further configured to:

when the target processor supports simultaneous multithreading processing, allocate the task to one processing unit of the target processor for execution.

15. The computing device according to claim 13, wherein the target processor comprises at least one processor core, and the main processor is further configured to:

when the target processor does not support simultaneous multithreading processing, allocate the task to one processor core of the target processor.

16. The computing device according to claim 13, wherein the target processor belongs to the target scheduling domain, the instruction sets supported by the target scheduling domain is the same as the instruction sets supported by the target processor, and the main processor is further configured to:

when the processor fails to execute the task, store an identifier of the target scheduling domain in a first storage unit of the task control block of the task, wherein the identifier stored in the first storage unit indicates that the processor comprised in the target scheduling domain is not allocated to execute the application.

17. The computing device according to claim 13, wherein the target processor belongs to the target scheduling domain, the instruction sets supported by the target scheduling domain is the same as the instruction sets supported by the target processor, and the main processor is further configured to:

when the processor successfully executes the task, store the identifier of the target scheduling domain in a second storage unit of the task control block of the task, wherein the identifier stored in the second storage unit indicates that the processor comprised in the target scheduling domain can be allocated to execute the application.

18. The computing device according to claim 13, wherein the task is a process or a thread generated by the application.