BL O/292/04

23rd September 2004

PATENTS ACT 1977

APPLICANT ARM Limited

ISSUE Whether patent application Nº GB 0025696.6 is excluded by section 1(2)(c)

HEARING OFFICER S J Probert

DECISION

Introduction

1 Patent Application Nº GB 0025696.6 was filed on 19th October 2000 in the name of ARM Limited; it claimed priority from a US application which was filed on 20 January 2000. The application is entitled "Simulation of data processing apparatus", and it was published on 12 November 2001 (serial Nº GB 2363222 A).

2 During the course of substantive examination, the examiner raised an objection concerning the patentability of the invention, in addition to other defects in the application — eg. lack of novelty, clarity, consistency and support. Most of these objections were overcome by suitable amendments and/or observations made on behalf of the applicant during the correspondence that has taken place. However, the examiner maintained that the invention, as described and as claimed, in this application is excluded from patentability by section 1(2) of the Act. More specifically, he reported that the invention relates to a program for a computer as such.

3 In his third examination report, the examiner noted that it was unlikely that further correspondence would resolve the issue and he invited the applicant to request a hearing. The applicant agreed, and requested that a hearing be appointed. The matter therefore came before me at a hearing on 14th July 2004, at which the applicant was represented by Mr Nigel Robinson and Dr Susan Keston of D. Young & Co.. The inventor, Dr Mark Burton, also attended the hearing, and gave a very helpful presentation of his invention.

The invention

4 The invention is a simulator for simulating the operation of a data processing apparatus, such as a microprocessor. It is particularly suited to simulating microprocessors that include pipelined circuit elements. There are a number of reasons why someone might want to simulate the operation of a microprocessor. For example, during the design process, the manufacturer will generally want to validate the design before committing

to the costs associated with production of a new device. More importantly in the context of this invention, it is often in the manufacturer's interest to supply potential customers with a simulated model of a microprocessor in advance of its release onto the market so that they can evaluate the performance of their own products by incorporating the model of the projected new microprocessor in a simulation of the overall circuit.

Pipelined architecture

A 'pipeline' in a microprocessor is a sequence of functional units ('stages') which perform a task in several steps, like an assembly line in a factory. Each functional unit takes inputs and produces outputs which are stored in its output buffer. One

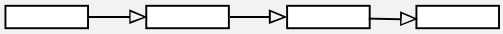

stage's output buffer is the next stage's input buffer. This arrangement allows all the stages to work in parallel thus giving greater throughput than if each input had to pass through the whole pipeline before the next input could enter.

(InstantWeb Online Computing Dictionary)

5 A simulated model of a microprocessor must provide an accurate representation of how the actual, physical device will function, but it is also important that the simulation can operate rapidly. As Mr Robinson explained, it would be possible in theory to evaluate a microprocessor design using pen and paper, but given the speed at which modern microprocessors operate today, it could easily take a thousand years to evaluate all aspects of the design fully. Using a computerised simulation typically means that a simulation may be run over a weekend. To illustrate one of the advantages of the present invention, Mr Robinson suggested that by comparison with the above examples, a simulator incorporating the features claimed in this application may run overnight.

6 In a simulator for simulating a pipelined microprocessor, each stage of the pipeline is modelled by a piece of software. That is to say, a computer program takes data from an input register, operates on it in such a way as to replicate the function of the hardware stage that is being modelled, and then loads the resulting output into an output register. Just as there may be many different stages in a pipeline processor (eg. fetch instruction, decode instruction, fetch arguments, perform arithmetic operations, store results), there must also be a corresponding number of computer programs (suitably linked together) to simulate every stage in the pipeline. In a typical, physical microprocessor, data is transferred between the stages simultaneously by clocking latches that are situated between the hardware stages; but in a conventional computer-simulated model, the data from each stage has to be transferred individually from each output register to the input register of the next stage. Furthermore, this is an operation that has to be performed at the end of each cycle of the simulated clock signal.

7 The invention in this application specifically concerns the means by which data is transferred from the output register of one stage to the input register of the next stage.

8 According to the prior art described in the application, this data transfer is usually performed by a series of COPY instructions. This means that for every stage in the pipeline, the simulator has to execute a COPY instruction in order to transfer data from each stage to the next, thereby simulating the operation of the latches in the actual pipelined microprocessor. If there were only two or three stages, this might not

necessarily be a serious overhead. But most modern microprocessors are "highly pipelined". Moreover, it is not uncommon for a microprocessor to have two or more pipelines that can run in parallel. In a simulated model of such a device, the delays generated by COPYing data between the output and input registers of adjacent stages of the model can be significant.

9 What the inventor has realised, is that by locating all of the input registers in one data storage area (eg. Memory A), and all of the output registers in another data storage area (eg. Memory B), the simulator only has to swap the two storage areas around in order to effectively COPY all of the output registers to the corresponding input registers of the next stage in a single operation.

10 The application describes several ways of swapping the two data storage areas, but the person skilled in the art will appreciate that there are also other ways that clearly fall within the scope of the claims. According to one example described in the application, the two data storage areas are located in different parts of a single memory device. The simulator saves the address of the first entry of each storage area in a special variable (called a 'pointer') and is then able to refer to the other entries in the storage area in relation to the first entry. In practice this is achieved by means of an 'offset' which, when added to the address stored in the pointer, provides the required address. At the end of each simulated clock cycle, the pointers to the two data storage areas are swapped such that the output values of the preceding cycle automatically become the input values for the next stage in the following cycle.

11 Alternatively, the simulator may be constructed from two separate pieces (or fragments) of program code — one compiled to take inputs from Memory A and store outputs in Memory B; the other compiled to take inputs from Memory B and store outputs in Memory A. In this way, rather than swapping pointers, it is possible to execute the two code fragments alternately. Although the overall size of the simulator program is increased in this embodiment, the simulation time is reduced since there is now no need to swap the pointer values between simulator clock cycles. This principle is conveniently illustrated in figure 2 of the application, part of which is reproduced below.

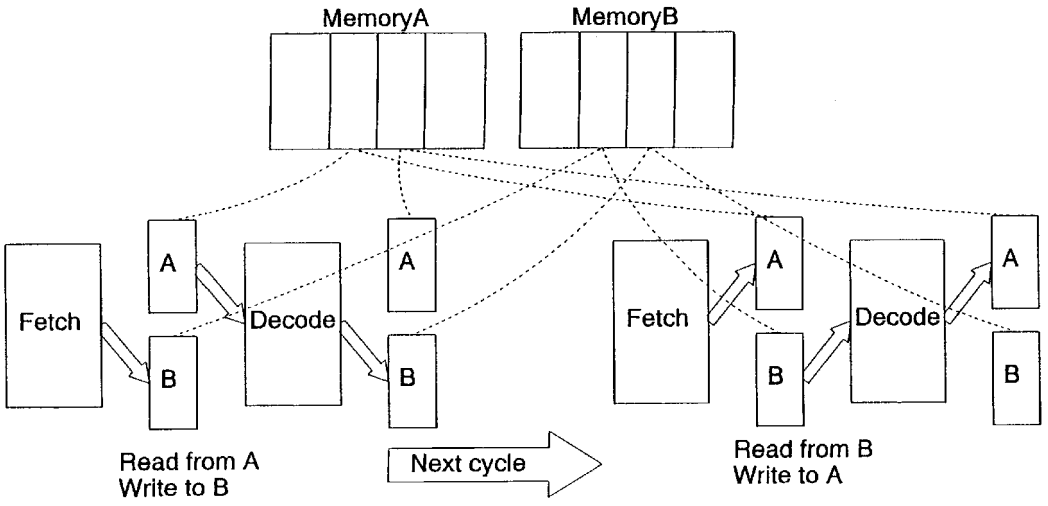

12 In the above paragraphs, I have deliberately referred to some of the simpler embodiments described in the application in order to illustrate the substance of the invention as clearly as possible. The claims (below) are somewhat more complex since they have been drafted (and amended during examination) so as to protect the invention as effectively as possible, including several more complex embodiments of the same fundamental invention.

The Claims

13 When the hearing took place, there were three independent claims in the application. Although each of the three claims is in a different form (ie. method, apparatus, program carrier), the substance of each claim is the same. For convenience I need only reproduce claim 1 in this decision:

> 1. A method of testing data processing operation including software-simulating operation of data processing apparatus including a plurality of pipelined circuit elements driven by a common clock signal, said method using a plurality of pipelined circuit element models linked by one-deep message queues and a set of data storage areas and comprising the steps of:
>
> within a first data storage area storing input data values representing respective input signals passed to each pipelined circuit element at commencement of a simulated common clock signal cycle of said common clock signal;
>
> during simulation of said simulated common clock signal cycle reading said input data values for each pipelined circuit element from said first data storage area and using one of said plurality of pipelined circuit element models for each pipelined circuit element to generate output data values representing output signals generated by a corresponding one of said plurality of pipelined circuit elements by termination of said simulated common clock signal cycle;
>
> within a second data storage area storing said output data values; and
>
> changing said second data storage area to serve as said first data storage area such that said stored output data values may serve as input data values during a following simulated clock signal cycle and changing to use a different storage area as said second storage area.

The Law

14 The examiner has reported that the application relates to a scheme, rule or method for doing business and/or a program for a computer as such. This objection is based on section 1(2) of the Act, the relevant parts of which are shown in bold below:

> **1(2) It is hereby declared that the following (among other things) are not inventions for the purposes of this Act, that is to say, anything which consists of-**
>
> (a) a discovery, scientific theory or mathematical method;
>
> (b) a literary, dramatic, musical or artistic work or any other aesthetic creation whatsoever;
>
> (c) a scheme, rule or method for performing a mental act, playing a game or doing business, or **a program for a computer;**
>
> (d) the presentation of information;

but the foregoing provision shall prevent anything from being treated as an invention for the purposes of this Act only to the extent that a patent or application for a patent relates to that thing as such.

15 It has been established by the Courts that an invention will not be excluded from patentability by the above subsection if it makes a technical contribution.[1] That is to say, if an invention makes a technical contribution, then it is ***more than*** one of the above excluded items, and cannot be regarded as ..... "that thing ***as such***".

16 Having read the application carefully, and in the light of submissions from Mr Robinson and Dr Burton at the hearing, I have come to the conclusion that the invention that I have described above ***does*** involve a technical contribution. The invention is a simulator. More importantly, it is a technically improved simulator. The invention may be implemented using a program for a computer, but the invention is more than a program for a computer as such. More specifically, the presence of a technical contribution indicates that it is more than a program for a computer ***as such***.

17 I am aware that in the majority of patent applications that have followed a similar path to this one, the Hearing Officer has usually been unable to identify any technical contribution and accordingly has refused the application. Therefore I shall try to highlight the difference in this application that has led to a different outcome.

Increase in processing speed?

18 First, in the particular circumstances of this case, it is worth stating clearly that the technical contribution does not arise simply because the simulator runs faster than conventional simulators. It is well established that an increase in processing speed, of itself, is not enough. One has to look more closely at what is going on, and examine why the simulator runs faster. On this occasion, the simulator runs faster because it operates differently at a technical level. Instead of merely increasing the speed at which the data is moved in turn from the output register of each stage to the input register of the next, the fundamental construction of the simulator has been modified such that this series of time-consuming move operations is no longer required. This strikes me as a technical contribution.

Hardware or software?

19 Those embodiments described in the greatest detail in the application explain how the invention may be implemented by swapping the two pointers that point to the two storage areas in memory. Although it is not specifically stated in the application, it seems reasonably obvious to me that these particular embodiments would not require any modification to the hardware of the simulator. These embodiments of the invention could, and most probably would, be implemented entirely within the computer program element of the simulator.

20 However, it was equally obvious to me from a first reading of the application, that the invention could also be implemented by using other means to switch the two storage

---

[1] *Fujitsu Limited's Application* [1997] RPC 14 at page 614.

areas. For example, discrete logic circuitry could be used to enable/disable different memory devices, or to map the input and output registers used by the various simulator stages to different areas of a common memory device, eg. by inverting address lines. These are all techniques that are well known in the art. Such embodiments are clearly within the scope of the claims. More importantly, there is no doubt that they would require modifications to the hardware of the simulator, as well as changes to the computer program.

21 Here too, it is worth stating that a technical contribution does not exist merely because an invention may be embodied in "hardware". This was pointed out by the EPO Technical Board of Appeal in *IBM/Document abstracting and retrieving*[2]. In that case, the appellant had argued that when a program was run on a conventional computer, it would bring about a change in the physical environment inasmuch as physical electric signals within the computer would be changed. However, in the Board's view this did not satisfy the requirement for a technical contribution. In paragraph 14 of its reasons, the Board said:

> "The foregoing considerations have been made mainly on the basis that the claimed systems and methods would involve a conventional computer controlled by a software program. The application states that this is the preferred embodiment of the invention and no other embodiment is specifically disclosed. Analogous considerations, however, apply in the case where the control of the computer would be effected by hardware ..., an option also falling within the scope of the claims, as the choice between the two possibilities is not of an essential nature but is based on technical and economic considerations which bear no relationship to the inventive concept as such."

22 I am also aware that in *Gale's Application*[3], Lord Justice Nicholls said:

> "I approach the substantial issue in this case, therefore, on the footing that it is convenient and right to strip away, as a confusing irrelevance, the fact that the claim is for 'hardware'."

23 Both of these last two authorities confirm that it is the substance of the invention that matters and not, in this instance, whether the invention may be implemented in hardware or software. Moreover, there is potentially a huge difference between replicating the function of a computer program using discrete logic on the one hand, and recognising the existence of an invention that may be implemented in hardware or software on the other hand. The former could be said of practically every computer program that has ever been written, and would (if accepted) drive a coach and horses through the provisions of section 1(2); whereas the latter is more applicable to inventions in a wide range of technical fields (eg. telecommunications, engine management, process control systems), the patentability of which would not be questioned by anyone who is familiar with patent law.

---

[2] *IBM/Document abstracting and retrieving* [1990] 1-2 OJEPO 12 (T22/85)

[3] *Gale's Application* [1991] RPC 13 at page 326

24 Having carefully considered the description of this invention, especially with a view to identifying the ***substance*** of the invention, it appears to me that the invention in this application is a technical improvement in the construction and operation of a simulator. It is a 'neat' technical solution to a technical problem. The invention is made at a much deeper level than the computer program that implements the simulator in the described embodiments. In other words, the invention itself is not directly associated with any particular computer program or sequence of instructions. For all these reasons also, I am satisfied that the invention provides a technical contribution.

Conclusion

25 I have decided that the invention described and claimed in this application involves a technical contribution, and that consequently it is not excluded from patentability by section 1(2)(c). I am therefore returning the application to the examiner in order for him to send the application forward to grant.

S J PROBERT
Deputy Director acting for the Comptroller