BESCHWERDEKAMMERN DES EUROPÄISCHEN PATENTAMTS

BOARDS OF APPEAL OF THE EUROPEAN PATENT OFFICE

CHAMBRES DE RECOURS DE L'OFFICE EUROPÉEN DES BREVETS

Internal distribution code:

(A) [-] Publication in OJ
(B) [-] To Chairmen and Members
(C) [-] To Chairmen
(D) [X] No distribution

Datasheet for the decision of 25 July 2022

Case Number:	T 1224/17 - 3.5.06
Application Number:	12161064.6
Publication Number:	2506147
IPC:	G06F9/54
Language of the proceedings:	EN

Title of invention:
Epoll optimisations

Applicant:
Xilinx, Inc.

Headword:
Epoll optimisation/XILINX

Relevant legal provisions:
EPC Art. 56

Keyword:
Inventive step - (yes) auxiliary request 1

Decisions cited:

This datasheet is not part of the Decisior
It can be changed at any time and without notic

Catchword:

Beschwerdekammern

Boards of Appeal

Chambres de recours

Boards of Appeal of the
European Patent Office
Richard-Reitzner-Allee 8
85540 Haar
GERMANY
Tel. +49 (0)89 2399-0
Fax +49 (0)89 2399-4465

Case Number: T 1224/17 - 3.5.06

DECISION
of Technical Board of Appeal 3.5.06
of 25 July 2022

Appellant: (Applicant)	Xilinx, Inc. 2100 Logic Drive San Jose, California 95124 (US)
Representative:	Driver, Virginia Rozanne Page White & Farrer Bedford House John Street London WC1N 2BF (GB)
Decision under appeal:	**Decision of the Examining Division of the European Patent Office posted on 3 January 2017 refusing European patent application No. 12161064.6 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman	M. Müller
Members:	S. Krischer
	A. Jimenez

Summary of Facts and Submissions

I. The appeal is directed against the decision of the examining division, dated 3 January 2017, to refuse application No. 12161064.

II. The decision was delivered in abridged form ("according to the state of the file"), and refers for its reasons to two communications: first to a communication under Rule 71(3) EPC (intention to grant a European patent), dated 3 November 2015, and secondly to the summons to oral proceedings, dated 19 March 2015. The annex to the summons contained reasons why the examining division considered the then main request not to be allowable, the communication under Rule 71(3) EPC indicated that the examining division intended to grant a patent on the basis of the then auxiliary request 3 (which was not approved by the applicant), and why the higher-ranking requests, filed in response to the summons, had not swayed the examining division's opinion expressed in the summons.

III. The following documents were referred to:

D1 Anonymous: "Epoll - I/O event notification facility", Linux Progammer's Manual, 1 February 2009, XP055046179, retrieved from the Internet: http://www.freebsd.org/cgi/man.cgi?query=epoll&apropos=0&sektion=0&manpath=SuSE Linux/i386 11.3&arch=default&format=pdf.

D2 Xiangyong Ouyang et al.: "Accelerating Checkpoint Operation by Node-Level Write Aggregation on Multicore Systems", International Conference on Parallel Processing 2009, ICPP '09, IEEE, Piscataway/NJ, USA, 22 September 2009, pages 34-41, XP031586421, ISBN: 978-1-4244-4961-3.

IV. A notice of appeal was received on 16 February 2017. The appeal fee was paid on the same day. With a statement of grounds of appeal received on 3 May 2017, the appellant requested that the decision under appeal be set aside and that a patent be granted on the basis of one of three sets of claim (1-16, 1-16 and 1-14, respectively) filed with as well, in combination with the documents as indicated in the decision under appeal.

V. In a communication dated 6 May 2022, the board raised an objection of lack of support by the description (Article 84 EPC) in respect of the main request. The board also indicated that it was minded to remit the case with an order to grant a patent based on auxiliary request 1.

VI. In a letter dated 23 June 2022, the appellant requested that the case be remitted with an order to grant a patent on the basis of auxiliary request 1, and withdrew its request for oral proceedings on the condition that the case be remitted with such an order.

VII. Further, in a letter of 8 July 2022, the appellant withdrew its main request.

VIII. Claim 1 of auxiliary request 1 reads as follows:

"1. A method for managing I/O event notifications in a data processing system (200), the data processing system comprising a plurality of applications (202, 203) and an operating system (201) having a kernel (209) and an I/O event notification mechanism (204) operable to maintain a plurality of I/O event notification objects (205, 206) each handling a set of file descriptors associated with one or more I/O resources, the method comprising:

directing a plurality of application-level configuration calls from an application to the I/O event notification mechanism for configuring an I/O event notification object; the method characterised by:

for each of said plurality of application-level configuration calls: intercepting at a user-level interface (210) a configuration call from said application to the I/O event notification mechanism for configuring said I/O event notification object; and

storing a set of parameters of the configuration call at a data structure (211), each set of parameters representing an operation on the set of file descriptors handled by the I/O event notification object;

and subsequently, on a predetermined criterion being met:

the user-level interface (210) making a first system call (212) from the user level to the kernel so as to pass the sets of parameters to a kernel library; and

the kernel library causing the plurality of configuration calls to be effected by making a configuration call for each set of parameters."

IX. Independent claim 16 of auxiliary request 1 relates to a data processing system comprising only means corresponding to the steps of claim 1.

X. In view of the board's decision, the wording of the claims according to auxiliary request 2 is immaterial.

Reasons for the Decision

1. *Summary of the invention*

The application relates to an improvement of an operating system program for managing input/output (I/O) event notifications, such as the system call "epoll" of the Linux operating system (OS) kernel (see https://en.wikipedia.org/wiki/Epoll) which is the only embodiment of the claimed subject-matter disclosed in the patent application. Epoll monitors sets of file descriptors for events to see if I/O is possible on resources/devices represented by any of these file descriptors. A resource/device may be a file, socket, FIFO, block or character device, or other I/O resources maintained by the OS (see A2 publication, [2], first sentence). Epoll is "relatively inefficient when the set of monitored file descriptors and I/O events changes because this additionally requires calls to epoll_ctl()" ([4], third sentence). The invention consists in *intercepting* epoll_ctl() system calls (figure 2: 208; [38], fifth sentence) at the user-level interface (210) in the user space (207). When several epoll_ctl() system calls (208) have been intercepted and their parameters stored, the user-level interface makes a single combined system call (212; [40], first sentence) to the kernel, needing only one context switch from user space (207) to kernel space (201). Then in the kernel space (201), the epoll_ctl() system calls (214) represented by the single system call (212) are *made within the kernel without the need for time-consuming context switches for each of them*, as would have been the case with the many original system calls (208; see also figure 1: 108) from user space to kernel space ([41]). In doing so, the applications (202, 203) that make the epoll_ctl() system calls do not have to be modified ([38], sentences 1-4) because the invention works transparently.

2. *Original disclosure*

2.1 The examining division did not raise any objections under Article 123(2) EPC in its decision and the board concurs that there was no reason to do so.

2.2 The amendments in independent claims 1 and 16 of auxiliary request 1 are based on original claim 5 and description paragraph [41], first sentence.

3. *Inventiveness*

3.1 In the two communications to which the examining division referred in its decision according to the state of the file (see point II above), it took D1 as the starting point for assessing inventive step (see summons, 3.1). The subject-matter of the then main request differed from D1 in the last three steps steps (i.e. intercepting configuration calls, storing parameter sets of the configuration calls, effecting configuration calls by a first system call). These steps were disclosed in D2 (summons, 3.5).

3.2 The board agrees with the appellant (grounds of appeal, 6.11) that D2 does not disclose an interception of calls (i.e. at runtime and transparent to the applications). Instead, according to the method of D2, the programmer of the applications is instructed to make sure that the application process (AP) stores its "medium size" data into a shared memory allocated by an IO process (IOP), see page 37, right column, third paragraph, third sentence:

> "The IOP creates a shared memory region that can be accessed by all application processes (APs)."

fourth paragraph, first sentence:

> "1) Application processes: When taking a checkpoint, an AP will perform a series of data writes as described below."

and sixth paragraph:

> "(3) If the data is regarded as “medium size”, then we [i.e. the AP] will store this data into the share memory created by IOP."

3.3 The AP furthermore sends a signal to the IOP to notify it of the pending data, see page 38, left column, second paragraph, fifth sentence:

> "After an AP finishes copying data to the shared memory, it checks if the share memory has accumulated enough dirty data. If it finds enough data there, then it sends a signal to the IOP to notify it of the pending dirty data."

3.4 It is furthermore not disclosed in D2 whether the IOP is in user space or kernel space. In the board's view, the skilled person would assume it to be in kernel space, since it is concerned with I/O. This additionally puts in question the argument of the examining division that the IOP corresponds to the "user-level interface", which is in user space (see the summons to oral proceedings of the examining division, dated 19 March 2015, page 3, firth paragraph).

3.5 However, without interception in user space, the method of D2 cannot produce the technical effect of reducing the context-switches between user space and kernel space in the same way as the invention, regardless of whether it reduces context switches at all.

3.6 As to the argument in the summons of the examining division (3.5, last but one paragraph) that the aggregation of data writes in D2 implied an interception in user space, the board notes that the aggregation is achieved by programming the applications accordingly (i.e. by construction), and not by interception (i.e. at run-time).

3.7 Therefore, the subject-matter of claim 1 of auxiliary request 1 involves an inventive step over the combination of D1 and D2, Article 56 EPC.

4. *Remittal*

4.1 The appellant did not file application documents adapted in view of the claims of auxiliary request 1. The board agrees that it is reasonable to postpone the adaptation of the description until after agreement on an allowable set of claims has been reached. The board also considers the common practice for boards of appeal to leave the adaptation of the description to the applicant (or proprietor) in dialogue with the department of first instance, as part of the grant (or maintenance) procedure that is to be carried out. The board considers this to be a special reason within the meaning of Article 11 RPBA 2020 for remitting this case to the first instance.

4.2 In this context, the board notes that dependent claim 5 of auxiliary request 1 seems to be redundant, since its additional features are contained in claim 1 of auxiliary request 1.

Order

For these reasons it is decided that:

1) The decision under appeal is set aside.

2) The case is remitted to the department of first instance with the order to grant a patent on the basis of claims 1-16 of auxiliary request 1 filed with the grounds of appeal, with the description and the drawings to be adapted as necessary.

The Registrar: The Chairman:

Beschwerdekammern des Europäischen Patentamts ✦ Boards of Appeal of the European Patent Office ✦ Chambres de recours de l'Office européen des brevets ✦

K. Götz-Wein Martin Müller

Decision electronically authenticated