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

Datasheet for the decision
of 4 May 2017

Case Number:	T 0923/12 - 3.5.06
Application Number:	09751301.4
Publication Number:	2165260
IPC:	G06F9/50, G06F9/45
Language of the proceedings:	EN

Title of invention:
PARALLEL PROCESSING OF DISTRIBUTED ARRAYS

Applicant:
The MathWorks, Inc.

Headword:
Data distribution schemes/MATHWORKS

Relevant legal provisions:
EPC Art. 83

Keyword:
Sufficiency of disclosure - (no)

Decisions cited:

EPA Form 3030
This datasheet is not part of the Decision.
It can be changed at any time and without notice.

Catchword:

Beschwerdekammern

Boards of Appeal

Chambres de recours

European Patent Office
D-80298 MUNICH
GERMANY
Tel. +49 (0) 89 2399-0
Fax +49 (0) 89 2399-4465

Case Number: T 0923/12 - 3.5.06

D E C I S I O N
of Technical Board of Appeal 3.5.06
of 4 May 2017

Appellant: (Applicant)	The MathWorks, Inc. 3 Apple Hill Drive Natick, MA 01760-2098 (US)
Representative:	Betten & Resch Patent- und Rechtsanwälte PartGmbB Maximiliansplatz 14 80333 München (DE)
Decision under appeal:	**Decision of the Examining Division of the European Patent Office posted on 6 December 2011 refusing European patent application No. 09751301.4 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman   W. Sekretaruk
Members: S. Krischer
                  A. Teale

Summary of Facts and Submissions

I.   The appeal is directed against the decision of the
        examining division, dated 6 December 2011, to refuse
        application No. 09751301.4 for added subject-matter
        (Article 123(2) EPC). Under "further remarks", the
        decision states that the invention is insufficiently
        disclosed (Article 83 EPC).

II.   A notice of appeal was received on 20 January 2012. The
        appeal fee was paid the same day. A statement of
        grounds of appeal was received on 5 April 2012. Claim
        sets according to a main and an auxiliary request were
        filed. Oral proceedings were conditionally requested.

III.   In its summons to oral proceedings, the board gave
        reasons for its preliminary opinion that the
        application lacked sufficient disclosure (Article 83
        EPC).

        The following documents were referred to:

        D1   N. Travinin Bliss et al.: "pMatlab parallel
              Matlab library", International Journal of High
              Performance Computing Applications, Sage Science
              Press, USA, vol. 21, no. 3, September 2007,
              pages 336-359, XP9122756, ISSN: 1078-3482.

        D2   S. Hiranandani et al.: "COMPILING FORTRAN D for
              MIMD Distributed-Memory Machines", Communi-
              cations of the ACM, USA, vol. 35, no. 8, August
              1992, pages 66-80, XP331490, ISSN: 0001-0782.

        D3   M. Hall et al.: "Interprocedural compilation of
              Fortran D for MIMD distributed-memory machines",
              Proceedings of the Supercomputing Conference,

Minneapolis, 16-20 November 1992, IEEE Comp.
Soc. Press, USA, vol. 5, pages 522-534,
XP10026946, ISBN: 978-0-8186-2630-2.

D4   B. Carpenter et al.: "Translation schemes for
      the HPJava parallel programming language", 14th
      Int. Workshop "Languages and Compilers for
      Parallel Computing" (LCPC 2001), Lecture Notes
      in Computer Science vol. 2624, Springer-Verlag,
      2003, pages 18-32, XP2547687,
      ISBN: 3-540-04029-3.

IV.   In a letter dated 4 April 2017, the appellant filed a
      claim set of a new second auxiliary request.

V.   Oral proceedings were held on 4 May 2017. At their
      conclusion, the board's decision was announced.

VI.   The appellant requests that the decision be set aside
      and a patent granted based on the main or first
      auxiliary request, both filed with the grounds of
      appeal, or the second auxiliary request, filed on
      4 April 2017.
      The further application documents on file are:
      description pages 1-38 as published; drawing sheets
      1-47 as published.

VII. Claim 1 of the main request reads as follows:

      "1. A computing device-implemented method for
      performing parallel processing, comprising:
         identifying, via a single programming language, one
      or more data distribution schemes for executing a
      program;

selecting a data distribution scheme of the one or
more identified data distribution schemes such as to
optimize a time to solution;

transforming, via the single programming language,
the program into a parallel program based on the
selected data distribution scheme;

allocating the parallel program to two or more labs
for parallel execution, each of the labs being a
hardware and/or software unit;

receiving one or more results associated with the
parallel execution of the parallel program from the two
or more labs; and

providing the one or more results to the program."

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

"1. A computing device-implemented method for
performing parallel processing, comprising:

simultaneously executing, to derive an efficient
data distribution scheme, the steps of identifying, via
a single programming language, one or more data
distribution schemes for executing a program and
selecting an appropriate data distribution scheme of
the one or more identified data distribution schemes
for executing the program, selecting a fastest
algorithm for executing the program, and selecting
appropriate resources for executing the program;

transforming, via the single programming language,
the program into a parallel program based on the
efficient data distribution scheme;

allocating the parallel program to two or more labs
for parallel execution, each of the labs being a
hardware and/or software unit;

receiving one or more results associated with the
parallel execution of the parallel program from the two
or more labs; and
    providing the one or more results to the program,
    wherein the efficient data distribution scheme is
derived such as to optimize a time to solution."

IX.   Claim 1 of the second auxiliary request reads as
       follows:

"1. A computing device-implemented method for
performing parallel processing, comprising:
    identifying, via a single programming language, one
or more data distribution schemes for executing a
program;
    selecting an appropriate data distribution scheme of
the one or more identified data distribution schemes
for each operation of a plurality of operations in the
program to optimize a time to solution;
    transforming, via the single programming language,
the program into a parallel program based on the
selected appropriate data distribution scheme for each
operation of the plurality of operations;
    allocating the parallel program to two or more labs
for parallel execution, each of the labs being a
hardware and/or software unit;
    receiving one or more results associated with the
parallel execution of the parallel program from the two
or more labs; and
    providing the one or more results to the program."

Reasons for the Decision

1.   *Overview of the invention*

        The application relates to specifying "data distri-
        bution schemes" by means of parallel processing
        constructs in a program (see original description
        page 15, lines 21-23; page 19, paragraph 3 and lines
        17-20 and 29-31; called a "single programming language"
        in the claims). The data to be distributed may be a
        "distributed array", i.e. an array (the data-structure)
        which is divided into segments for being processed in
        parallel on different processors or different cores of
        one processor (page 1, third paragraph; page 7, lines
        9-12; page 3, lines 29-30; page 15, paragraph 5; for
        "distributed arrays" see D1, page 1, right column,
        second paragraph). The parallel processing constructs
        (not explicitly claimed) of the single programming
        language are "Distribution scheme commands" (page 19,
        paragraph 5; figures 11-14) and "Distributed array
        commands" (paragraph 6; figures 16-21). The central
        notion of the "data distribution scheme" itself is not
        explained in the description. What is disclosed, is
        that the distribution scheme commands "provide a
        mechanism (e.g. a distribution property of a
        distributor object) for obtaining a specific
        distribution scheme object" (page 22, line 6-8; see
        also figure 14) and that the data distribution scheme
        may be user-defined (figure 15: 1530; page 22, lines
        23-24, 26-31; page 30, line 14-17; page 37, lines 1-5).
        Further, distribution scheme commands specify a layout
        of data onto a parallel resource set (e.g. labs;
        page 4, lines 2-9), and specify which parallel resource

set is to be used for a distribution (page 29, lines 27-29).

The claimed method performs the following steps: the parallel processing constructs of the program are *identified* (page 29, lines 14-17) and an *optimal* data distribution scheme is *selected* in order to optimize a time to solution (page 29, lines 18-19; and page 18, line 25-27 for the expression "time to solution"). Then, the program is *transformed* according to the selected data distribution scheme into a "parallelly" executable form (page 19, lines 20-22 and 33-35; see also page 20, lines 8-12; repeated on page 29, lines 31-35) and executed in parallel by several processors (figure 3) or cores of a processor (page 7, lines 9-12).

2.   *Insufficient disclosure (Article 83 EPC)*

2.1   The board agrees with the statements in the grounds of appeal (page 10, third paragraph) and the decision (sections 17.4 and 18.2) that the difference between the claimed invention and the prior art is the *selection of an optimal data distribution scheme*. None of documents D1 to D4 discloses a selection, let alone an optimal selection, of a data distribution scheme (decision page 5, lines 1-4; page 11, last paragraph).

2.2   The decision (sections 17.4 and 18.2) argues that this optimal selection is the core of the claimed invention.

2.3   In the grounds of appeal, this is confirmed by the formulation of the objective technical problem derived from this difference ("to improve the performance of

parallel executing of a particular program"; page 10,
paragraph 4).

2.4   The board agrees that the selection of an optimal data
        distribution scheme in order to improve the time to
        solution in executing a program is the core of the
        invention.

2.5   However, the application does not disclose how this
        core should be programmed by the software developer
        implementing the claimed invention. The description
        merely states *that* the invention automatically selects
        an optimal data distribution scheme which minimises the
        time to solution (page 18, lines 25-30; page 29, lines
        18-25), but not *what* data distribution schemes look
        like, *how* a set of possible data distribution schemes
        might be derived, and *how* the fastest one should be
        selected from such a set for any program to be
        parallelised and executed by the invention *before*
        actually executing that program.

2.6   The grounds of appeal (page 9, last paragraph) argue
        that, since D1 disclosed performance metrics, the
        skilled person would have been "aware of various
        techniques for implementing performance metrics and
        thus for optimizing the time to solution".

2.7   However, the performance metrics of D1 consist in
        measuring "execution time ... as compared with serial
        MATLAB, the underl[a]ying MatlabMPI communication
        library and C+MPI benchmarks" (page 4, table 1, line
        "Metrics" below the "Goal - High performance"; see also
        section 2, third sentence: "The performance metrics ...
        primarily look at the computation and memory overhead
        of programs written with pMatlab relative to serial

programs written using MATLAB and parallel programs
written using C with MPI."). This means that they
measure the execution time of three manually written
program variants (see figure 4 for the pMatlab program)
implementing the same mathematical function (e.g. Fast
Fourier Transform/FFT) to compute the result for the
same benchmark suite as input (Abstract, ninth
sentence: "Performance is evaluated by implementing the
HPC Challenge benchmark suite and comparing pMatlab
performance with the equivalent C+MPI codes.").

2.8   In the board's view, this does not enable a skilled
        person to implement an algorithm which selects the
        fastest data distribution scheme *before* executing the
        program to be analysed. Such an algorithm would have to
        estimate the execution time beforehand with possible
        (applicable) data distributions schemes. D1 merely
        shows how to *measure* the execution time of a program
        with a data distribution scheme given by the writer of
        that program.

2.9   The appellant argues that "given past performance
        metrics for programs and/or operations executed under
        known data distribution schemes, the skilled person may
        choose a known data distribution scheme for future
        execution of a program and/or operation that optimizes
        a time to solution for the program by selecting a data
        distribution scheme with optimized past performance
        metrics relative to other data distribution
        schemes" (letter dated 4 April 2017, page 2, last
        paragraph).

2.10   For the board, this sounds like an instruction as to
        how to make an invention. Instead of filling the gaps
        in the disclosure with common general knowledge, a

skilled person would need to make extensive experimentation, and also theoretical research in order to generalise the results of such experimentation to find heuristics to determine optimal data distribution schemes for as many typical programs as possible to be input to the claimed method. However such preparatory work before programming the claimed computer-implemented method would clearly impose an undue burden on the skilled person in carrying out the invention.

2.11   Furthermore, in view of the disclosure that distribution scheme commands specify a layout of data onto a parallel resource set and specify which parallel resource set is to be used for a distribution (page 29, lines 27-29), the question arises where the leeway for selecting an optimal data distribution scheme could lie, if the data distribution scheme is specified by the program constructs which the writer of the program to be parallelised and executed had chosen beforehand.

2.12   The appellant argues that "there is leeway in the number and types of labs to use" (letter dated 4 April 2017, page 3, third paragraph).

2.13   However, there is no disclosure of *how* to exploit this leeway. The skilled person would have to find the solution on his/her own.

2.14   This shows that the core of the invention, i.e. the selection of an optimal data distribution scheme, is not disclosed in a manner sufficiently clear and complete for it to be carried out by a person skilled in the art, contrary to Article 83 EPC.

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

B. Atienza Vivancos   W. Sekretaruk

Decision electronically authenticated