BESCHWERDEKAMMERN   BOARDS OF APPEAL OF   CHAMBRES DE RECOURS
DES EUROPÄISCHEN   THE EUROPEAN PATENT   DE L'OFFICE EUROPÉEN
PATENTAMTS   OFFICE   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 23 February 2016

Case Number: T 2343/10 - 3.5.07

Application Number: 03729000.4

Publication Number: 1506501

IPC: G06F17/30

Language of the proceedings: EN

Title of invention:
Associative database searching using FPGA devices

Applicant:
IP Reservoir, LLC

Headword:
Data matching using FPGA/IP RESERVOIR

Relevant legal provisions:
EPC Art. 123(2), 56
RPBA Art. 13(1)

Keyword:
Amendments - added subject-matter (yes) main request and
auxiliary requests 1-9
Inventive step - (no) auxiliary requests 10 and 11

Decisions cited:

Catchword:

Case Number: **T 2343/10 - 3.5.07**

D E C I S I O N
of Technical Board of Appeal 3.5.07
of 23 February 2016

Appellant:	IP Reservoir, LLC
(Applicant)	349 Marshall Avenue, Suite 100
	St. Louis, MO 63119 (US)
Representative:	Naylor, Matthew John
	Mewburn Ellis LLP
	City Tower
	40 Basinghall Street
	London EC2V 5DE (GB)
Decision under appeal:	Decision of the Examining Division of the European Patent Office posted on 12 July 2010 refusing European patent application No. 03729000.4 pursuant to Article 97(2) EPC.

Composition of the Board:

Chairman	R. Moufang
Members:	M. Rognoni
	R. de Man

Summary of Facts and Submissions

I.   The former applicant (appellant), Washington University, appealed against the decision of the Examining Division to refuse European patent application no. 03729000.4, filed as international application PCT/US03/15638 and published as WO 03/100662.

II.   The Examining Division based its decision essentially on the following objections against all requests filed with telefax dated 10 May 2010:

-       the subject-matter of claim 1 according to the main request extended beyond the content of the application as originally filed (Article 123(2) EPC);
-       claim 1 of the main request was neither clear nor supported by the description as required by Article 84 EPC;
-       claims 1 according to auxiliary requests 1 to 7 did not comply with Articles 123(2) and 84 EPC.

        As an Obiter Dictum, the Examining Division noted that the subject-matter of a claim amended to overcome the objections of clarity and added subject-matter would not be new over the following prior-art document:

        D7: Nwodoh T. A., *et al*.: "A Processing System for Real-Time Holographic Video Computation", Proceedings of the SPIE Conference on Reconfigurable Technology: FPGAs for Computing and Applications, Boston, September 1999, SPIE Vol. 3844, pages 129 to 140.

        Furthermore, in the Examining Division's opinion, the subject-matter of this hypothetical claim would not

involve an inventive step in view of the common knowledge generally available in the art.

III.   With the statement of grounds of appeal the appellant filed new claim sets according to a main request and auxiliary requests 1 to 9.

IV.   Registration of the transfer of the present application to IP Reservoir, LLC (current appellant) took effect on 13 January 2014.

V.   With a communication dated 9 September 2015, the appellant was summoned to oral proceedings scheduled to take place on 23 February 2016.

VI.   In a communication pursuant to Article 15(1) EPC, dated 11 December 2015, the Board expressed, *inter alia*, the preliminary opinion that claim 1 according to the main request and to auxiliary requests 1 to 9 did not appear to comply with Article 123(2) EPC. The Board expressed also doubts as to the inventive step of the claimed invention and, in this respect, drew the appellant's attention to the following prior art document:

D9:   Leung, K., Ercegovac, M. and Muntz, R.: "Exploiting Reconfigurable FPGA for Parallel Query Processing in Computation Intensive Data Mining Applications", UC MICRO Technical Report, 1999.

VII.   With letter dated 8 February 2016, the appellant submitted new auxiliary requests 10 and 11.

VIII.   Oral proceedings before the Board were held as scheduled on 23 February 2016. At the end of these proceedings, the Chairman pronounced the Board's decision.

IX.   The appellant requested that the decision under appeal
       be set aside and a patent be granted on the basis of the
       claims of the main request filed with the statement of
       grounds of appeal or, in the alternative, on the basis
       of the claims of one of auxiliary requests 1 to 9 filed
       with the statement of grounds of appeal or one of
       auxiliary requests 10 and 11 filed with the letter dated
       8 February 2016.

X.   Claim 1 according to the appellant's <u>main request</u> reads
       as follows:

       "A system for performing data reduction, the system
       comprising:
           at least one re-configurable logic device (21) that
       is configured to (i) store a data key, (ii) receive a
       stream of target data, the target data comprising a
       plurality of data fields, (iii) match the stream of
       target data against the data key to find matched data
       within the stream of target data with respect to a first
       field of the target data, and (iv) summarize the matched
       data in aggregate form according to a second field of
       the target data."

       Claims 2 to 18 are dependent on claim 1. Claim 19 is
       directed to a method essentially comprising the steps
       which the system of claim 1 is configured to perform.
       Claims 20 to 36 are dependent on claim 19.

       Claim 1 according to <u>auxiliary request 1</u> reads as
       follows:

       "A system for performing data reduction, the system
       comprising:
           at least one re-configurable logic device (21) that
       is configured to (i) store a data key, (ii) receive a

stream of target data, the target data comprising a
plurality of portions, (iii) match the stream of target
data against the data key to find matched data within
the stream of target data, and (iv) summarize the
matched data in aggregate form such that the matching
operation and the aggregate summarization operation are
performed on different portions of the target data."

Claim 1 according to <u>auxiliary request 2</u> reads as
follows:

"A system for performing data reduction, the system
comprising:
     at least one re-configurable logic device (21) that
is configured to (i) store a data key, (ii) receive a
stream of financial information, (iii) match the stream
of financial information against the data key to find
matched data within the stream of financial information,
the matched data comprising data representative of a
plurality of prices for a stock, and (iv) summarize the
matched data in aggregate form according to the prices
for the stock."

Claim 1 according to <u>auxiliary request 3</u> differs from
claim 1 of auxiliary request 2 in that features (iii)
and (iv) read as follows:

"(iii) match the stream of financial information against
the data key to find matched data within the stream of
financial information, the matched data comprising data
representative of a plurality of prices and times for a
stock, and (iv) summarize the matched data in aggregate
form according to the prices for the stock by computing
in parallel paths a minimum price, a maximum price, and
a latest price for the stock prices within the matched
data as the matched data streams through a data shift

register (102) resident on the at least one reconfigurable logic device."

Claim 1 of auxiliary request 4 reads as follows:

"An apparatus comprising:
    at least one re-configurable logic device (21) for performing data reduction searching on financial information, the at least one re-configurable logic device configured to (i) store a data key, (ii) receive a stream of financial information, the financial information comprising a plurality of stock prices and associated times, (iii) match the stream of financial information against the data key to find a stock within the stream of financial information, (iv) return a plurality of stock prices and times for the found stock, (v) stream the stock prices and times for the found stock through a data shift register (102) resident on the at least one re-configurable logic device, and (vi) summarize the stock prices and times for the found stock in an aggregate form by identifying a minimum price, a maximum price, and a latest price of the found stock using decision logic resident on the at least one re-configurable logic device as the stock prices and times for the found stock stream through the data shift register, the decision logic comprising parallel paths for identifying the minimum price, maximum price, and latest price of the found stock."

Claim 1 according to auxiliary request 5 differs from claim 1 of auxiliary request 4 in that the last feature of the latter ("the decision logic comprising parallel paths for identifying the minimum price, maximum price, and latest price of the found stock") has been replaced by the following wording:

"wherein the decision logic comprises:

a first data register (104) configured to store a
most recent minimum stock price;

a second data register (106) configured to store a
most recent maximum stock price;

a third data register (108) configured to store a
most recent latest stock price;

a fourth data register (116) configured to store a
most recent time;

a first comparator (110);

a second comparator (112); and

a third comparator (114);

wherein the first, second and third comparators are
configured to operate in parallel with each other;

wherein the first comparator is configured to
compare the current stock price in the data shift
register with the stored most recent minimum stock price
to determine which is lower; and

wherein the decision logic is further configured to
update the stored most recent minimum stock price in the
first data register with the current stock price in
response to a determination by the first comparator that
the current stock price is lower;

wherein the second comparator is configured to
compare the current stock price in the data shift
register with the stored most recent maximum stock price
to determine which is higher;

wherein the decision logic is further configured to
update the stored most recent maximum stock price in the
second data register with the current stock price in
response to a determination by the second comparator
that the current stock price is higher;

wherein the third comparator is configured to
compare the current time in the data shift register with
the stored most recent time to determine which is later;
and

wherein the decision logic is further configured to
update the stored most recent latest stock price in the
third data register with the current stock price in the
data shift register in response to a determination by
the third comparator that the current time is later."

Claim 1 according to <u>auxiliary request 6</u> differs from
claim 1 according to auxiliary request 3 in that feature
(iv) reads as follows:

"(iv) summarize the matched data in aggregate form,
wherein the at least one re-configurable logic device
comprises:
     a data shift register (102);
     a first data register (104) configured to store a
most recent minimum stock price;
     a second data register (106) configured to store a
most recent maximum stock price, the second data
register configured to operate in parallel with the
first data register;
     a third data register (108) configured to store a
most recent latest stock price, the third data register
configured to operate in parallel with the first data
register and the second data register;
     a fourth data register (116) configured to store a
most recent time, the third data register configured to
operate in parallel with the first data register and the
second data register [*sic*];
     a first comparator (110);
     a second comparator (112) configured to operate in
parallel with the first comparator; and
     a third comparator (114) configured to operate in
parallel with the first comparator and the second
comparator;
     wherein the at least one re-configurable logic
device is further configured to stream the data

representative of the stock prices and times within the matched data through the data shift register;

wherein the first comparator is configured to compare the current stock price in the data shift register with the stored most recent minimum stock price to determine which is lower; and

wherein the at least one re-configurable logic device is further configured to update the stored most recent minimum stock price in the first data register with the current stock price in response to a determination by the first comparator that the current stock price is lower;

wherein the second comparator is configured to compare the current stock price in the data shift register with the stored most recent maximum stock price to determine which is higher;

wherein the at least one re-configurable logic device is further configured to update the stored most recent maximum stock price in the second data register with the current stock price in response to a determination by the second comparator that the current stock price is higher;

wherein the third comparator is configured to compare the current time in the data shift register with the stored most recent time to determine which is later; and

wherein the at least one re-configurable logic device is further configured to update the stored most recent latest stock price in the third data register with the current stock price in the data shift register in response to a determination by the third comparator that the current time is later."

Claim 1 according to <u>auxiliary request 7</u> reads as follows:

"A system for performing data reduction, the system
comprising:

at least one re-configurable logic device (21) that
is configured to (i) store a data key, (ii) receive a
stream of financial information, the financial
information comprising a plurality of stock symbols, a
plurality of stock prices associated with the stock
symbols, and a plurality of times associated with the
stock prices, (iii) match the stream of financial
information against the data key to find matched data
within the stream of financial information with respect
to a stock symbol, the matched data comprising a
plurality of stock prices and times for the stock
symbol, and (iv) summarize the matched data in aggregate
form according to price, wherein the at least one re-
configurable logic device comprises:

a data shift register (102);

a first data register (104) configured to store
a most recent minimum stock price;

a second data register (106) configured to store
a most recent maximum stock price;

a third data register (108) configured to store
a most recent latest stock price;

a fourth data register (116) configured to store
a most recent time;

a first comparator (110);

a second comparator (112); and

a third comparator (114);

wherein the at least one re-configurable logic
device is further configured to stream the stock prices
and times within the matched data through the data shift
register;

wherein the first comparator is configured to
compare the current stock price in the data shift
register with the stored most recent minimum stock price
to determine which is lower; and

wherein the at least one re-configurable logic
device is further configured to update the stored most
recent minimum stock price in the first data register
with the current stock price in response to a
determination by the first comparator that the current
stock price is lower;

wherein the second comparator is configured to
compare the current stock price in the data shift
register with the stored most recent maximum stock price
to determine which is higher:

wherein the at least one re-configurable logic
device is further configured to update the stored most
recent maximum stock price in the second data register
with the current stock price in response to a
determination by the second comparator that the current
stock price is higher;

wherein the third comparator is configured to
compare the current time in the data shift register with
the stored most recent time to determine which is later;
and

wherein the at least one re-configurable logic
device is further configured to update the stored most
recent latest stock price in the third data register
with the current stock price in the data shift register
in response to a determination by the third comparator
that the current time is later."

Claim 1 according to auxiliary request 8 reads as
follows:

"A system for performing data reduction, the system
comprising:

at least one re-configurable logic device (21) that
is configured to (i) store a data key, (ii) receive a
stream of target data, the target data comprising a
plurality of data fields, (iii) match the stream of

target data against the data key to find a plurality of
matched data occurrences within the stream of target
data, wherein the matched data occurrences comprise the
target data deemed a match against the data key as well
as other target data within a bounding field, and (iv)
summarize the matched data occurrences in aggregate form
such that the matching operation and the aggregate
summarization operation are performed on different data
fields of the target data."

Claim 1 according to auxiliary request 9 differs from
claim 1 of auxiliary request 8 in that the expression
"data fields" in features (ii), (iii) and (iv) has been
replaced with "data portions" or "portion", to read as
follows:

"(ii) receive a stream of target data, the target data
comprising a plurality of portions, (iii) .... wherein
the matched data occurrences comprise the target data
deemed a match against the data key as well as other
target data portions within a bounding field, and
(iv) ... such that the matching operation and the
aggregate summarization operation are performed on
different portions of the target data."

Claim 1 according to auxiliary request 10 reads as
follows:

"A method for performing data reduction searching, the
method comprising:
    streaming financial information through
reconfigurable logic that is configured to perform a
matching operation on data according to a data key;
    matching, via the reconfigurable logic configured to
perform the matching operation, the streaming financial
information against the data key to find matched data

within the streaming financial information, the matched
data comprising a plurality of prices and times for a
stock;

   storing the matched data in a memory;

   streaming the matched data from the memory through
reconfigurable logic that is configured to perform an
aggregate summarization operation via a plurality of
parallel paths; and

   summarizing, via the reconfigurable logic
configured to perform the aggregate summarization
operation, the streaming matched data according to the
prices for the stock by computing in the parallel paths
a maximum, minimum, and latest price for the stock as
the matched data streams through the reconfigurable
logic configured to perform the aggregate summarization
operation."

Claim 1 according to auxiliary request 11 reads as
follows:

   "A method for performing data reduction searching,
the method comprising:

   streaming financial information through
reconfigurable logic that is configured to perform a
matching operation on data according to a data key;

   matching, via the reconfigurable logic configured to
perform the matching operation, the streaming financial
information against the data key to find a plurality of
matched data occurrences within the streaming financial
information, wherein the matched data occurrences
comprise a plurality of prices and times for a stock
within a bounding field of the streaming financial
information;

   storing the matched data occurrences in a memory;

   streaming the stock prices and times within the
matched data occurrences from the memory through

reconfigurable logic that is configured to perform an
aggregate summarization operation via a plurality of
parallel paths; and

summarizing, via the reconfigurable logic configured
to perform the aggregate summarization operation, the
matched data occurrences according to the prices for the
stock by computing in the parallel paths a maximum,
minimum, and latest price for the stock as the stock
prices and times stream through the reconfigurable logic
configured to perform the aggregate summarization
operation."

XI.   The appellant's arguments relevant to the present
        decision may be summarised as follows:

        The content of a patent application encompassed all the
        subject-matter that the skilled person, using common
        general knowledge, would directly and unambiguously
        derive from the original application documents. Where
        features were disclosed in different embodiments within
        the application, as in the present case, the relevant
        question was whether the skilled person would seriously
        contemplate combining the different features.

        The present invention was capable of performing data
        reduction searching. According to page 26, lines 29 to
        34 of the application, reduction searching involved both
        matching data and summarising the matched data in some
        aggregate form. This was an explicitly clear and
        unambiguous disclosure of a data reduction embodiment
        which included both matching and aggregate
        summarisation. Indeed, the very general nature of the
        statement made it clear that this embodiment was not
        limited to one of the examples of matching previously
        described in the application, but rather could be
        implemented using any of the disclosed examples.

Claim 1 of the main request combined the disclosure in
the application relating to the use of reconfigurable
logic to perform data matching with the disclosure
relating to the use of reconfigurable logic to perform
aggregate summarisation.

Figure 27 disclosed a specific embodiment for the use of
reconfigurable logic to perform aggregate summarisation
of financial information. Many of the application's
preceding figures disclosed embodiments for using
reconfigurable logic to perform matching operations.
Given that the patent application already set out
extensive descriptions of how matching data against a
key could be performed in reconfigurable logic, the
description at pages 26 and 27 for the data reduction
searching embodiment focused on the aggregate
summarisation operation and did not explicitly referred
to matching operations again.

Although Figure 27 disclosed a specific application of
aggregate summarisation of financial information, the
patent application made it clear that the scope of the
disclosure for the data reduction searching embodiment
was not to be limited solely to the processing of
financial information. Furthermore, the application
clearly implied that all embodiments of the invention
could be deployed on at least one logic device. The
person skilled in the art would therefore understand
these disclosures to mean that if a particular
implementation required more logic resources than were
available on a given reconfigurable logic chip, then
multiple reconfigurable logic chips could be used
together.

It was furthermore well understood within the financial
industry that stock market information was distributed
as streaming feeds of stock prices for various stocks,
and that these feeds were comprised of various fields
such as stock symbol, stock price, and a time associated
with each stock price. This basic elemental
understanding was so well-known in the art that it
needed not be elaborated upon in the patent application.
Figure 27 expressly showed that the aggregate
summarisation logic for the data reduction searching
embodiment operated on the price and time fields of
streaming financial information of a stock. Hence,
financial information was disclosed as having multiple
fields.

Consequently, for the skilled person the present
application taught that key-based matching was operated
on streaming financial information comprised of multiple
fields in order to find the price data of a stock, where
the stock price data included at least a stock price
field and an associated time field, and furthermore that
the stock prices and time were processed by the
aggregate summarisation logic to compute the desired
aggregated price summarisation, namely its minimum,
maximum and latest prices. Accordingly, a skilled person
would understand that the matching operation and the
aggregate summarisation operation operated in different
fields of the financial information.

Claim 1 according to the main request and according to
the auxiliary requests 1 to 9 did not specify whether
the streamed target data was a live data stream or read
from a mass storage. Although the description explicitly
referred to target data streamed from a mass storage, it
was implicit to the skilled reader that the present
invention was not limited to the explicit embodiments

and, in particular, could be applied to live data
streams.

As to the question of inventive step, it was worth
reviewing the context in which the present invention had
been conceived back in 2002. Whilst the invention was
originally aimed primarily at the use of FPGAs for text
mining purposes, as discussed in the context of the
"matching alone" embodiments, the inventors also
realised that such methods could be applied alongside
data reduction searching in respect of financial market
data. This embodiment satisfied a substantial need in
the financial technology sector for a lower latency
"ticker plant" system, designed to handle the collection
and throughput of massive data streams, as generated by
financial markets, to allow the display of "real time"
prices and, more particularly, to feed computerised
trading systems.
The data reduction searching functionality according to
the present invention could provide trading applications
with updated information because of the speed with which
the relevant information could be extracted from the
financial data stream. This could translate into the
effective advantage of placing such trading applications
"ahead of the market" in relation to other applications
which relied on the output from conventional (slower)
ticker plant systems.

The use of FPGAs in the processing of streamed financial
data allowed the appellant's ticker plant appliances to
operate over ten times faster than commodity servers
running software.

None of the cited prior art taught the combination of
steps recited in claim 1 of any of the requests. In
particular, document D9 merely described a generalised

capability for using FPGAs to perform "stream oriented, signal processing type applications". No teachings were found that related to using FPGA logic to perform matching operations followed by aggregate summarisation operations, much less the specific aggregate summarisation operations recited in claim 1 according to the auxiliary requests 10 and 11. Instead, the focus of document D9 was on how to compute processing times for various operations so that a decision could be made as to whether such operations should be deployed in reconfigurable logic or on a host processor. Document D9 was completely lacking in any specific teachings that would guide the skilled person to the present invention.

Reasons for the Decision

1.   The appeal is admissible.

2.   The present application relates*, inter alia*, to the use of a reconfigurable logic device, such as a field programmable gate array (FPGA), in the context of data matching and data aggregation.

2.1   As explained at page 5, lines 3 to 11 of the published application, an advantage offered by FPGAs is the speed of reconfiguration. In fact, FPGAs "*may be quickly programmed in the field to optimize the search methodology using a template, the template having been prepared in advance and merely communicated to the FPGA's over a connecting bus. Should it then be desired to search using a different methodology, the FPGA's may then be quickly and conveniently re-programmed with another prepared template in a minimal number of clock cycles and the second search started immediately*".

Main request

3.   Claim 1 according to the main request relates to a
        "system for performing data reduction" and comprises the
        following features:

        -       at least one re-configurable logic device that is
                configured to

                (i) store a data key,

                (ii)receive a stream of target data,
        -       the target data comprising a plurality of
                data fields,

                (iii)match the stream of target data against the
                     data key to find matched data within the
                     stream of target data with respect to a
                     first field of the target data, and

                (iv)summarize the matched data in aggregate form
                    according to a second field of the target
                    data.

        Claim 1 of the main request differs from the
        corresponding claim of the main request considered by
        the Examining Division only by the deletion of the term
        "continuously" in feature (iii) and of the expression
        "to thereby achieve a data reduction for the matched
        data" in feature (iv), as shown in the following:

                (iii)~~continuously~~ match the stream of target
                     data against the data key to find matched
                     data within the stream of target data with
                     respect to a first field of the target data,
                     and

(iv) summarize the matched data in aggregate form
     according to a second field of the target
     data ~~to thereby achieve a data reduction for
     the matched data~~.

4.   In paragraph 12 of the contested decision, the Examining
        Division pointed out that the features "to find matched
        data [...] with respect to a **first field**" and
        "summarize the matched data [...] according to a **second
        field**" had no basis in the application as filed (Article
        123(2) EPC).

4.1.1   In particular, the Examining Division noted in paragraph
        13.1 of the contested decision that the sole example of
        the claimed "data reduction" given in the description
        (see published application page 26, line 32 to page 27,
        line 16) merely related to calculating/searching the
        minimum, maximum and latest price of a stock using a
        couple of shift registers and processors. It did not
        disclose matching a data field against a data key or
        using such matched data in any way. In fact, said
        example of "aggregate data reduction" operated on "raw"
        data read from a mass storage medium, not on matched
        data fields.

4.1.2   Furthermore, the Examining Division considered that the
        different embodiments of the invention disclosed in the
        application related either to "matching" or to "data
        reduction" and did not suggest a system combining both
        operations.

4.2   The appellant essentially argued that the application
        conveyed the following unambiguous information:

- the data reduction embodiment shown at pages 26, lines 32 to page 27, line 16 of the description comprised a matching operation and a downstream aggregate summarisation operation that was undoubtedly performed on the matched data;

- the financial information example of the data reduction embodiment operated to find the minimum, maximum and latest price for a stock.

4.2.1   In the appellant's view, the only reasonable conclusion that could be derived from the interplay of the above aspects of the invention was that the matching operation was directed to finding a stock within the financial information, while the downstream aggregate summarisation operation computed the min/max/latest price of the matching stock. This necessarily implied that the matching operation and the aggregate summarisation operation were performed on different data fields.

5.   In view of the reasons given by the Examining Division for refusing the then main request, the first question to be considered in the present appeal is whether the subject-matter according to claim 1 of the main request can be directly and unambiguously derived from the application as originally filed.

5.1   The application deals primarily with various forms of data matching (*i.e.* exact, approximate, analogue and digital).

        For instance, it is specified at page 5, line 36 to page 6, line 13 of the published application that, "*[g]enerally, the invention may be described as a technique for data retrieval through approximate*

matching of a data key with a continuous reading of data
as stored on a mass storage medium, using FPGA's to
contain the template for the search and do the
comparison, all in hardware and at essentially line
speed. By utilizing FPGA's, the many advantages and
features commonly known are made available. These
include the ability to arrange the FPGA's in a
"pipeline" orientation, in a "parallel" orientation, or
even in an array incorporating a complex web overlay of
interconnecting data paths allowing for complex
searching algorithms. In its broadest, and perhaps most
powerful, embodiment, the data key may be an analog
signal and it is matched with an analog signal generated
by a typical read/write device as it slews across the
mass storage medium. In other words, the steps taught to
be required in the prior art of not only reading the
analog representation of digital data stored on the mass
storage medium but also the conversion of that signal to
its digital format prior to being compared are
eliminated" (underlining added).

5.1.1   A second aspect of the invention, "data reduction",
         apart from being mentioned in a statement at page 4,
         lines 8 to 12 ("... *given the vast amount of information*
         available to be searched, data reduction operations
         (i.e., the ability to summarize data in some aggregate
         form) has [sic] *become critical. Oftentimes, the ability*
         to quickly perform data reduction functions can provide
         a company with a significant competitive advantage"), is
         more extensively treated in relation to the embodiment
         shown in Figure 27.

5.1.2   In particular, the description (page 26, line 29 to page
         27, line 16) specifies in connection with Figure 27 the
         following:

"*The present invention is also capable of performing data reduction searching. <u>Such searching involves matching as previously described herein</u>, but includes summarizing the <u>matched data</u> in some aggregate form. For example, in the financial industry, one might want to search financial information to identify a minimum, maximum, and latest price of a stock. A re-configurable logic device for computing such aggregate data reductions is illustrated as 100 in Figure 27. Here, a data shift register 102 <u>reads target data from a mass storage medium containing stock price information</u>. In the example shown, three data reduction searches are shown, namely calculating the minimum price, the maximum price, and the latest price. As target data is fed into the data shift register 102, decision logic computes the desired data reduction operation. In particular, the stock price is fed to a minimum price comparator 110 and maximum price comparator 112 and stored therein. Each time a stock price is fed to comparator 110, it compares the last stored stock price to the stock price currently being fed to it and whichever is lower is stored in data register 104. Likewise, each time a stock price is fed to comparator 112, it compares the last stored stock price to the stock price currently being fed to it and whichever is higher is stored in data register 106. In order to compute the latest price, the stock price is fed into a data register 108 and the current time is fed into a comparator 114. Each time a time value is fed into comparator 114, it compares the last stored time with the current time and which ever is greater is stored in data register 116. Then, at the end of the desired time interval for which a calculation is being made, the latest price is determined*" (underlining added).

5.1.3   In other words, the data reduction referred to in the
        description relates to *"target data"* read from a *"mass
        storage medium containing stock price information"* and
        is aimed at determining the minimum price, the maximum
        price and the latest price. It is further noted at page
        27, lines 17 to 23 that the present invention can
        perform data reduction searching for a variety of
        different applications of varying complexity requiring
        such functionality.

5.1.4   According to the disclosure at page 11, lines 3 to 13, a
        re-configurable logic device *"may itself include <u>a
        plurality of functional logic elements</u> including a data
        shift register and possibly a microprocessor, or they
        could be on separate chips, or the individual logic
        elements could be configured in a pipeline or parallel
        orientation as shown in some of the other figures
        herein. In any event, re-configurable logic refers to
        any logic technology whose form and function can be
        significantly altered (i.e., reconfigured) in the field
        post-manufacture. Examples of re-configurable logic
        devices include without limitation programmable logic
        devices (PLDs)"* (underlining added).

5.2   In general, the Board agrees with the appellant that the
        application as originally filed at least implicitly
        discloses that a re-configurable logic device may
        perform more than one logical function and that a system
        according to the invention can be configured to perform
        both data matching and data reduction on the previously
        matched data.

5.3   All independent claims 1, 10, 14, 28, 32 and 33 of the
        original application specify that the data to be matched
        or otherwise processed is stored in a mass storage
        device, in particular a magnetic storage device. In the

appellant's opinion, this limitation of the claimed
subject-matter was unnecessary, because the skilled
reader understood from the disclosure that the present
invention could be applied to a live data stream.

5.3.1   According to page 13, lines 5 to 11, of the published
        application, *"the present invention may be used to
        perform a variety of different types of matching or data
        reduction operations on the target data. [...] For all
        operations, however, it will be assumed that the target
        data is written onto the magnetic mass storage medium 26
        with sufficient formatting information attached so that
        the logical structure of the target data can be
        extracted"* (underlining added).

5.3.2   In particular, the present invention can be *"implemented
        in a stand-alone computer or computer system"* and
        comprises *"at least one re-configurable logic device 21
        coupled to at least one magnetic mass storage medium 26,
        with that re-configurable logic device being an FPGA"*
        (published application, page 10, line 37 to page 11,
        line 3 and Figure 1 - underlining added).

5.3.3   Thus, the mass storage medium provides storage for the
        large amount of data, referred to as *"target data"*,
        which is fed to the re-programmable logic device (see
        description, page 12, lines 3 to 5). In fact, it is
        reiterated throughout the application that the target
        data, on which the data matching or the data reduction
        operation is performed, is streamed from a mass storage
        device (cf. for instance page 15, lines 7 to 8 and lines
        22 to 24, page 19, lines 33 to 36, page 20, lines 27 to
        31, page 22, lines 30 to 31, page 23, lines 20 to 23).

5.3.4   As pointed out by the appellant, claim 1 of the main
        request does not exclude the possibility that between

the matching operation and the data reduction operation matched data is stored. However, even if it may make technical sense to buffer the matched data before data reduction, the skilled person need not implement the invention as claimed using a <u>mass storage medium</u> as a buffer. In other words, the mass storage medium referred to in the application cannot be regarded as an implicit feature of the claimed subject-matter.

5.3.5   In the Board's view, claim 1 clearly encompasses embodiments of a system which may perform both data matching and data reduction "on the fly", *i.e.* without data buffering between the two operations.

These embodiments are neither explicitly nor implicitly disclosed in the original application which, as shown above, consistently specifies that the "target data" is data streamed from a mass storage and, in particular, that *"a shift register 102 reads target data from a mass storage medium containing stock price information"* (page 26, lines 36 to 37).

5.3.6   Hence, as far as the claim covers the above embodiments, it generalizes and extends the teaching contained in the application as originally filed. It does therefore not comply with Article 123(2) EPC.

5.4   As to the Examining Division's objection that claim 1 then on file defined data matching and data summarisation with respect to undisclosed first and second data "fields" (see features (iii) and (iv)), the Board notes that the application as filed does not specify any format for the target data, although, in the case of financial information relating to stock prices, it may be assumed that target data would comprise different data "types", such as company identifiers and

corresponding stock prices at given times, and that data
matching would be performed on the basis of a company
identifier, while corresponding stock prices would be
the object of summarisation. However, this does not
necessarily imply that the target data comprises "*a
plurality of data fields*" and that the data processing
is performed on the basis of a first field and a second
field, as recited in claim 1 of the present main
request.

5.4.1   Indeed, according to the description (page 27, lines 17
        to 23), data reduction is not limited to financial
        information, but it can be performed for a variety of
        different applications of varying complexity requiring
        such functionality. Furthermore, the application
        provides several examples of data matching with
        reference to different kinds of target data, such as
        text data, image data etc. (cf. published application,
        page 4, lines 13 to 20, page 5, lines 21 to 28, page 7,
        lines 1 to 26).

5.4.2   If data matching consists essentially in detecting
        certain binary patterns in a data stream, data reduction
        involves also the extraction of data somehow linked with
        the matched patterns.

        In the case of financial information, data matching is
        related to the recognition of a company or stock
        identifier and data reduction to the extraction of other
        data associated with the matched pattern.

        On the other hand, the general statement at page 27,
        lines 17 to 23, could be interpreted as hinting at the
        possibility of performing data matching and reduction
        on, for instance, image data so as to identify images
        with a certain pattern (*e.g.* a face) and discarding from

the identified image all the pixels which do not belong
to the matched pattern. In this case, the data matching
and reduction would be performed on target data which
could not be divided into predefined "fields" or
"portions".

5.5   In the Board's opinion, defining the target data in
       claim 1 as comprising a plurality of data fields and
       specifying data matching and summarisation with respect
       to first and second fields constitute an undisclosed
       intermediate generalisation of the invention lying
       between the specific embodiment shown in Figure 27 and
       the general hint at other possible applications given,
       for instance, at page 27, lines 17 to 23.

5.6   Hence, the Board comes to the conclusion that claim 1
       according to the main request does not comply with
       Article 123(2) EPC.

Auxiliary requests 1 to 9

6.   Claims 1 according to the auxiliary requests 1 to 9
       encompass embodiments of a system which processes a
       stream of data "on the fly" to match the data to a
       stored data key and summarise the matched data.

6.1   However, as pointed out above, the application as filed
       specifies, in particular, that the "target data" to be
       "summarized" in some aggregate form according to the
       embodiment described at page 26, line 29 to page 27,
       line 16, is read "*from a mass storage medium*".

6.2   Thus, for the reasons given above, auxiliary requests 1
       to 9 relate to subject-matter extending beyond the
       content of the application as originally filed (Article
       123(2) EPC).

Admission of auxiliary requests 10 and 11

7.   In reply to the Board's preliminary opinion, the
        appellant filed auxiliary requests 10 and 11 with letter
        dated 8 February 2016. Although these requests were
        submitted late in the appeal proceedings, the Board
        considered that they constituted a promising attempt on
        the part of the appellant to overcome at least some of
        the Board's concerns.

7.1   Hence, the Board, exercising it discretion under
        Article 13(1) RPBA, decided to admit these late-filed
        requests into the proceedings.

Auxiliary request 10

8.   Claim 1 according to auxiliary request 10 relates to a
        "method for performing data reduction searching", which
        comprises the following features itemised by the Board:

        (a) streaming financial information through
            reconfigurable logic that is configured to perform
            a matching operation on data according to a data
            key;

        (b) matching, via the reconfigurable logic configured
            to perform the matching operation, the streaming
            financial information against the data key to find
            matched data within the streaming financial
            information,
                (i) the matched data comprising a plurality of
                    prices and times for a stock;

        (c) storing the matched data in a memory;

(d) streaming the matched data from the memory through reconfigurable logic that is configured to perform an aggregate summarization operation via a plurality of parallel paths; and

(e) summarizing, via the reconfigurable logic configured to perform the aggregate summarization operation, the streaming matched data according to the prices for the stock by computing in the parallel paths a maximum, minimum, and latest price for the stock as the matched data streams through the reconfigurable logic configured to perform the aggregate summarization operation.

9.   Steps (d) and (e) are directly supported by the embodiment shown in Figure 27 and described at page 26, line 29 to page 27, line 16. The description also specifies that "reduction searching" involves data matching and summarising the matched data, and that the data to be summarised are read from a mass storage medium. This implies that a method for performing the data reduction searching disclosed in the application also comprises steps (a), and (b) recited in claim 1.

       As to step (c), the Board notes that the term "*memory*" does not define exclusively a "*mass storage device*", but may also apply to a data buffer. It is therefore broader than the terms used throughout the application, *i.e. "mass storage medium"* and *"magnetic storage medium"*.

9.1   It may therefore be questionable whether claim 1 according to auxiliary request 10 complies with Article 123(2) EPC. However, in view of the fact that a possible objection under Article 123(2) EPC could be easily overcome and thus would not prejudice the grant of a

patent, the Board has decided not to pursue this issue
further, but to direct its attention to the question of
inventive step of the claimed subject-matter.

10.   According to the appellant, none of the available prior
       art documents disclosed a method for performing data
       reduction searching which combined data matching with
       data summarising performed on the matched data. There
       was also no suggestion in the prior art that it might be
       advantageous to provide a method based on the
       combination of these two operations. In particular,
       there was no disclosure hinting at the possibility of
       performing these operations on financial information for
       computing a maximum, minimum and latest price for a
       stock.

10.1   Furthermore, in the appellant's opinion, it was
       particularly advantageous to perform these operations by
       means of reconfigurable logic which, in view of its
       flexibility and processing speed, provided up-to-date
       financial information and thus best satisfied the
       demands of today's financial world.

11.   The claimed method is directed to the processing of
       streamed financial information relating to stock prices.
       The data stream is processed to identify a particular
       stock and the corresponding stock price at a given time
       in order to determine the maximum price, the minimum
       price and the latest price.

11.1   It is undisputed that financial information in the form
       of a report of the price for certain securities, updated
       continuously throughout the trading session by the
       various stock exchanges (*i.e.* "stock ticker"), was known
       and available before the filing date of the present
       application.

11.2   In the Board's opinion, it would be the task of a
        financial analyst to decide what kind of information
        should be extracted from a stock ticker and how it
        should be presented. In any case, even the lay investor
        should be interested in receiving regular updates on the
        maximum, minimum and latest prices reached by selected
        stocks during a particular trading session.

11.3   Starting from this general background, the claimed
        method addresses the problem of extracting information
        of general and obvious interest from stock price reports
        delivered as a data stream.

11.4   The skilled person confronted with the above problem
        will realise that it can be solved by performing a few
        basic logic operations. Firstly, in order to extract the
        price of a stock, the stock must be matched with a stock
        identifier. Secondly, the minimum and maximum prices are
        found by comparing the latest price with the stored
        maximum or minimum, respectively.

11.5   In the Board's opinion, the mere idea of combining
        simple logic operations, which are anyway necessary for
        extracting required information, cannot involve an
        inventive step.

11.6   A trivial way of performing the logic operations
        required to extract the minimum, maximum and latest
        stock prices from financial information would be by
        means of a software program running on a general purpose
        computer.

11.7   On the other hand, document D9, which is concerned with
        exploiting re-configurable FPGAs for query processing in
        computation intensive data mining application, points

out (see "Introduction") that General Purpose Processors (GPPs) do not provide enough computation power for processing and manipulating highly complex data. Specialized hardware (ASICs) provides considerable gains in performance. A further solution is provided by FPGAs based on static RAMs (SRAM-based FPGA) which *"have bridged two extreme worlds (ASICs) and (GPPs) of computer architecture and provide not only the flexibility of binding the functionality of hardware dynamically, but also achieve a significant performance gain in certain applications"* (sentence bridging columns 1 and 2).

11.7.1 Furthermore, according to section 2.1.2 (first paragraph), "*[a]vailability of SRAM-based FPGA coprocessors creates a new design space and optimization problem for high performance parallel query processing.*"

In fact, "*[t]he applications that have been reported to perform well on FPGAs are stream oriented, signal processing type applications. ... The data moves through the logic and the results are returned to users as a stream without having to wait for the completion of the evaluation of the entire task*".

11.8   In the light of the known advantages that FPGAs provide when performing logic operations on a data stream, the Board considers that it would have been obvious to a skilled person, wishing to implement a system for extracting stock prices from a data stream of financial information and determining the minimum, maximum and latest prices of a certain stock, to consider the use of reconfigurable logic.

As the minimum, maximum and latest prices for a stock are determined on the basis of three independent

operations performed on the same data, it would have been obvious to exploit a further known advantage provided by FPGAs (see section 2.1.1 of D9) and perform the corresponding operations via parallel paths.

11.9   As to the step of storing the matched data in a memory prior to the summarising step (see feature (c) of claim 1), the Board considers that storing data in a buffer memory between two operations is a feature generally known in the art in order to avoid synchronisation problems between different data operations and data loss. On the other hand, if the term "memory" is interpreted as a mass storage medium, feature (c) would render the summarising operation time independent of the data matching and allow further data analysis, as the matched data would be permanently stored.

In the Board's opinion, a skilled person, in the light of the data operations to be performed and of general technical considerations, would have contemplated the possibility of storing the matched data in a data buffer or in a mass storage medium before the data summarising step (e). In doing so, the skilled person would have arrived at a system falling within the terms of claim 1 of the auxiliary request 10 (Article 56 EPC).

Auxiliary request 11

12.   Claim 1 according to auxiliary request 11 differs from claim 1 of auxiliary request 10 essentially in the way it specifies the streamed financial information. In particular, claim 1 defines the matched data as "*matched data occurrences*" which "*comprise a plurality of prices and times for a stock within a bounding field of the streaming financial information*".

12.1   The above definitions aim at clarifying that the
        "matched data" comprise not only data which matches the
        data key (*e.g.* the stock ID), but also related data,
        such as stock prices and times, on which the aggregate
        summarisation operation is performed.

12.2   In the Board's opinion, the differences between claims 1
        of the auxiliary requests 10 and 11 do not contain any
        technical features which might contribute to the
        inventive step of the claimed subject-matter. Hence, for
        the same reasons given above, the subject-matter of
        claim 1 does not satisfy the requirement of Article 56
        EPC.

13.   As none of the appellant's requests provides a basis for
        granting a patent, the appeal has to be dismissed.

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

I. Aperribay   R. Moufang

Decision electronically authenticated