



US007809927B2

(12) **United States Patent**
Shi et al.

(10) **Patent No.:** **US 7,809,927 B2**
(45) **Date of Patent:** **Oct. 5, 2010**

(54) **COMPUTATION PARALLELIZATION IN
SOFTWARE RECONFIGURABLE ALL
DIGITAL PHASE LOCK LOOP**

(75) Inventors: **Fuqiang Shi**, Allen, TX (US); **Roman
Staszewski**, McKinney, TX (US);
Robert B. Staszewski, Delft (NL)

(73) Assignee: **Texas Instruments Incorporated**,
Dallas, TX (US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 243 days.

(21) Appl. No.: **11/949,310**

(22) Filed: **Dec. 3, 2007**

(65) **Prior Publication Data**

US 2009/0070568 A1 Mar. 12, 2009

Related U.S. Application Data

(63) Continuation-in-part of application No. 11/853,575,
filed on Sep. 11, 2007.

(51) **Int. Cl.**
G06F 9/30 (2006.01)

(52) **U.S. Cl.** **712/36; 708/303; 708/320**

(58) **Field of Classification Search** **712/36;**
708/303, 320

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,495,591 A * 1/1985 Loomis, Jr. 708/320
4,811,263 A * 3/1989 Hedley et al. 708/316
5,081,604 A * 1/1992 Tanaka 708/319
5,243,551 A * 9/1993 Knowles et al. 708/603
5,432,723 A * 7/1995 Chen et al. 708/300
5,523,962 A * 6/1996 Yoshino et al. 708/319

5,636,150 A * 6/1997 Okamoto 708/300
5,636,353 A * 6/1997 Ikenaga et al. 712/218
5,745,396 A * 4/1998 Shanbhag 708/322
5,761,104 A * 6/1998 Lloyd et al. 708/517
5,805,479 A * 9/1998 Tang 708/316
6,108,680 A * 8/2000 Liu et al. 708/317

(Continued)

OTHER PUBLICATIONS

Parhi (Chapter 10: Pipelined and Parallel Recursive and Adaptive
Filters) <http://web.archive.org/web/20030823062135/www.ece.umn.edu/users/parhi/SLIDES/chap10.pdf>*

(Continued)

Primary Examiner—Eddie P Chan


Assistant Examiner—Keith Vicary

(74) *Attorney, Agent, or Firm*—Ronald O. Neerings; Wade
James Brady, III; Frederick J. Telecky, Jr.

(57) **ABSTRACT**

A novel and useful apparatus for and method of software based phase locked loop (PLL). The software based PLL incorporates a reconfigurable calculation unit (RCU) that is optimized and programmed to sequentially perform all the atomic operations of a PLL or any other desired task in a time sharing manner. An application specific instruction-set processor (ASIP) incorporating the RCU includes an instruction set whose instructions are optimized to perform the atomic operations of a PLL. A multi-stage data stream based processor incorporates a parallel/pipelined architecture optimized to perform data stream processing efficiently. The multi-stage parallel/pipelined processor provides significantly higher processing speeds by combining multiple RCUs wherein input data samples are input in parallel to all RCUs while computation results from one RCU are used by adjacent downstream RCUs. A register file provides storage for historical values while local storage in each RCU provides storage for temporary results.

22 Claims, 26 Drawing Sheets



U.S. PATENT DOCUMENTS

6,115,580	A *	9/2000	Chuprun et al.	455/1
6,122,653	A *	9/2000	Kuroda	708/320
6,272,616	B1 *	8/2001	Fernando et al.	712/20
6,438,570	B1 *	8/2002	Miller	708/625
6,809,598	B1	10/2004	Staszewski et al.	
7,584,342	B1 *	9/2009	Nordquist et al.	712/22
2002/0049799	A1 *	4/2002	Wang	708/320
2004/0093366	A1 *	5/2004	Gagnon et al.	708/300
2006/0033582	A1	2/2006	Staszewski et al.	
2006/0038710	A1	2/2006	Staszewski et al.	
2006/0195498	A1 *	8/2006	Dobbek et al.	708/300
2007/0260856	A1 *	11/2007	Tran et al.	712/217
2009/0030961	A1 *	1/2009	Matsuyama et al.	708/320

OTHER PUBLICATIONS

Parhi et al. (Pipeline Interleaving and Parallelism in Recursive Digital Filters—Part I: Pipelining Using Scattered Look-Ahead And Decomposition) This paper appears in: Acoustics, Speech and Signal Processing, IEEE Transactions on; Jul. 1989; vol. 37, Issue: 7; On pp. 1099-1117.*

Parhi et al. (Pipeline Interleaving and Parallelism in Recursive Digital Filters—Part II: Pipelined Incremental Block Filtering) This paper appears in: Acoustics, Speech and Signal Processing, IEEE Transactions on; Jul. 1989; vol. 37, Issue: 7; On pp. 1118-1134.*

Wu et al. (Application-Specific CAD of VLSI Second-Order Sections) This paper appears in: Acoustics, Speech and Signal Processing, IEEE Transactions on; Publication Date: May 1988; vol. 36, Issue: 5; On pp. 813-825.*

Lu et al. (Fast Recursive Filtering with Multiple Slow Processing Elements) This paper appears in: Circuits and Systems, IEEE Transactions on; Publication Date: Nov. 1985; vol. 32, Issue: 11; On pp. 1119-1129.*

Sung et al. (Efficient Multi-Processor Implementation of Recursive Digital filters) This paper appears in: Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP '86.; Publication Date: Apr. 1986; vol. 11, On pp. 257-260.*

Lorca et al. (Efficient ASIC and FPGA Implementations of IIR Filters for Real Time Edge Detection) This paper appears in: Image Processing, 1997. Proceedings., International Conference on; Publication Date: Oct. 26-29, 1997; On pp. 406-409 vol. 2.*

Moyer (An Efficient Parallel Algorithm for Digital IIR Filters); This paper appears in: Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP '76.; Publication Date: Apr. 1976; vol. 1, On pp. 525-528.*

Wong et al. (Computer-Aided Design of Pipelined IIR Digital Filters) This paper appears in: Circuits and Systems, 1992., Proceedings of the 35th Midwest Symposium on; Publication Date: Aug. 9-12, 1992; On pp. 795-799 vol. 2.*

Ljung (Embedded.com: How to Create Fixed- and floating-point IIR filters for FPGAs) Published May 31, 2006, Accessed Dec. 10, 2009.*

Staszewski et al. (VHDL Simulation and Modeling of an All-Digital RF Transmitter) This paper appears in: System-on-Chip for Real-Time Applications, 2005. Proceedings. Fifth International Workshop on; Publication Date: Jul. 20-24, 2005; On pp. 233-238.*

Krenik et al. (Digital RF and Handset Integration) Wireless Design & Development | Nov. 1, 2005 | Accessed Dec. 10, 2009.*

Curtin et al. (Phase-Locked Loops for High-Frequency Receivers and Transmitters—Part 1) Analog Dialogue: vol. 33, No. 3, Mar. 1999.*

U.S. Appl. No. 11/853,575, filed Sep. 11, 2007, Staszewski et al.

U.S. Appl. No. 11/853,588, filed Sep. 11, 2007, Staszewski et al.

* cited by examiner




FIG. 1
(PRIOR ART)




FIG. 2
(PRIOR ART)





FIG. 3
(PRIOR ART)






FIG. 5




FIG. 6




FIG. 7





FIG. 9

FIG. 8






FIG. 10









FIG. 11

$$y_k = (1-a) \cdot y_{k-1} + a \cdot x_k$$

$$H(z) = \frac{a}{1 - (1-a) \cdot z^{-1}}$$

FIG. 12A








FIG. 15






FIG. 16







FIG. 19






FIG. 21




FIG. 22







FIG. 23

FIG. 24










FIG. 27




FIG. 28




FIG. 29




FIG. 30




FIG. 32





FIG. 31

FIG. 33






FIG. 34




FIG. 37




FIG. 36





FIG. 39



1

COMPUTATION PARALLELIZATION IN SOFTWARE RECONFIGURABLE ALL DIGITAL PHASE LOCK LOOP

REFERENCE TO PRIORITY APPLICATION

This application is a continuation-in-part (CIP) of U.S. application Ser. No. 11/853,575, filed Sep. 11, 2007, entitled "Software Reconfigurable Digital Phase Lock Loop Architecture", incorporated herein by reference in its entirety.

CROSS-REFERENCE TO RELATED APPLICATION

This application is related to U.S. application Ser. No. 11/853,588, filed Sep. 11, 2007, entitled "Computation Spreading for Spur Reduction in a Digital Phase Lock Loop", incorporated herein by reference in its entirety.

FIELD OF THE INVENTION

The present invention relates to the field of data communications and more particularly relates to a software reconfigurable all digital phase locked loop (ADPLL) architecture and application specific instruction-set processor related thereto.

BACKGROUND OF THE INVENTION

Phase locked loop (PLL) circuits are well known in the art. A block diagram illustrating an example prior art phase locked loop (PLL) circuit is shown in FIG. 1. The typical PLL circuit, generally referenced 170, comprises phase detector 172, loop filter or low pass filter (LPF) 174 and voltage controlled oscillator (VCO) 176.

In operation, a frequency reference clock signal, often derived from a crystal oscillator, is input to the phase detector along with the VCO output signal (often divided down). The phase detector, typically implemented as a charge pump or mixer, generates a phase error (PHE) proportional to the phase difference between the reference clock input signal and the VCO output clock signal. The resultant PHE signal is then low pass filtered to yield a slow varying frequency command signal that controls the frequency of the VCO. The frequency command signal is input to a VCO or digitally controlled oscillator (DCO) such that the VCO output frequency/phase is locked to the reference clock with a certain fixed relationship. This oscillator generates an RF signal whose frequency depends on the frequency command signal.

In wireless communication systems, e.g., GSM, UMTS, Bluetooth, WiFi, etc., the RF synthesizer is a fundamental block that is used to provide a high quality, high frequency RF carrier for the transmitter and a local oscillator clock for the receiver, whose output frequency can range from several hundreds of MHz to several GHz. Different applications with different standards require different RF frequencies with different RF performance requirements. The RF clock generating the RF carrier plays a critical role in the entire wireless communication system. The quality of the RF clock directly affects the communication performance and often is the determining factor whether the system meets standards specifications.

Typically, the RF synthesizer is implemented using a phase locked loop (PLL) typically using a pure hardwired (i.e. fixed hardware with limited reconfigurability) design approach. All digital phase locked loops (ADPLLs) for RF synthesizer construction targeting wireless communications are known in the

2

art. Conventional ADPLL circuits, however, are implemented as purely fixed hardware based (or hardwired) with very limited reconfigurability. It is thus difficult for one design to support multi-standard wireless applications, e.g., GSM, GPRS, EDGE, WCDMA, etc. as well as wireless data networks, such as Bluetooth, WiFi and WiMAX.

Once a hardwired circuit design is committed to a physical implementation, there is little that can be changed regarding the transfer function or operation of the ADPLL. Any modification requiring logic and interconnect change results in numerous time consuming steps within the ASIC creation process (i.e. timing closure, physical design, etc.) typically requiring significant engineering resources and months of delay to launch a product. In addition, once the silicon is manufactured, any change to the ADPLL architecture makes an even costlier impact, making such changes virtually impractical.

In general, a main difference between a hardwired implementation and a microprocessor based implementation is that the microprocessor implementation uses shared hardware running at higher speed, while the hardwired implementation uses dedicated hardware running at lower speed. A block diagram illustrating an example prior art generalized processing block using a dedicated hardware implementation is shown in FIG. 2. The hardwired implementation, generally referenced 10, comprises a plurality of dedicated hardware blocks 12 for each function 14. The circuit provides memory (Mem1, Mem2, Mem3, Mem4) and dedicated hardware for each function (F1, F2, F3, F4), wherein each block runs at the data path speed f_s .

A block diagram illustrating an example prior art generalized processing block using a processor based implementation is shown in FIG. 3. The circuit, generally referenced 16, comprises instruction memory 18, instruction fetch 20, instruction decode 22, ALU 24, data bus 29, register file 26 and data memory 28. The processor based solution has one shared hardware block ALU that can be configured to execute any of the four functions (F1, F2, F3, F4). The ALU is programmed by the instructions stored in instruction memory 18 and the ALU is adapted to run four times faster ($4f_s$) to complete the data processing within the data path speed of f_s .

It is important to note that the ALU 24 typically has a set of general purpose instructions which precludes its applicability in many applications, especially its use in low-power RF synthesizer circuits.

With CMOS process technology currently advancing from 65 nm to 45 nm to 32 nm, transistors are becoming faster and faster. The interconnections, however, are becoming more and more dominant in SOC design regarding the delay and area contribution. The interconnections in a hardwired design having a large area will significantly slow the circuit speed while adding a significant silicon area overhead. Since processor based solutions run at higher speed with shared hardware, resulting in smaller area, advancements in semiconductor technology will make processor based solutions more and more attractive. This further favors use of multiple but smaller processors with a dedicated instruction set rather than one processor with a more general instruction set.

Furthermore, in conventional ADPLL circuits, the digital part of local oscillator (DLO) (i.e. a portion of the ADPLL) is implemented using dedicated random logic gates. Thus, all computations are launched on the rising edge of the ADPLL system reference clock and latched on the next rising edge. Since a majority of the circuit switching activity is centered on the rising edge of the system reference clock, most of the digital current is being switched at that point as well, creating large current transients. These digital current surges find their

way into on-chip DCO and PA circuit nodes via various coupling mechanisms, e.g., capacitive, etc. These disturbances at the system clock rate have strong sub-harmonics that are upconverted into sensitive areas of the RF spectrum, resulting in unacceptable RF spurs.

It is thus desirable to have a processor based PLL architecture that is software based and programmable. The programmable PLL should provide an on-the-fly reconfiguration capability which eases silicon debugging and development tasks and provides multi-standard operation capability. Further, the software based PLL architecture should create significantly lower concentration of current transients thus reducing the generation of spurs in the output spectrum. At the same time, the unavoidable spurious energy that is generated by the logic activity and coupled into RF circuits should be pushed higher in frequency where they lie outside of or can be easily filtered out of critical frequency bands.

SUMMARY OF THE INVENTION

The present invention is a novel and useful apparatus for and method of software based phase locked loop (PLL). The processor-based PLL (i.e. all digital phase-locked loop or ADPLL) architecture described herein can be used for RF frequency synthesis in radio transceivers and other applications.

The software based phase locked loop of the present invention incorporates a reconfigurable calculation unit (RCU) that can be programmed to sequentially perform all the atomic operations of a phase locked loop or any other desired task. The RCU is a key component in a so called application specific instruction-set processor (ASIP). The ASIP includes an instruction set that comprises instructions optimized to perform the individual atomic operations of a phase locked loop.

The reconfigurable computational unit (RCU) is time shared for all computations within the phase locked loop. The reconfigurable computational unit and related configuration control logic replaces the dedicated and distributed random logic inside the conventional digital PLL. The reconfigurable computational unit is controlled via microcode stored in on-chip memory (e.g., RAM or ROM). Since the computational unit is time shared among all operations, it is operated at an oversampled rate that is high enough to insure the proper implementation of the phase locked loop. In order to achieve this, the reconfigurable computational unit is optimized to perform all computations of the phase locked loop atomic operations within a single reference clock cycle.

In one embodiment, the instruction set is implemented in microcode that is stored in volatile or non-volatile memory. Thus, the ASIP can easily be reconfigured to implement customized designs for different applications, such as multiple cellular standards, including GSM, GPRS, EDGE, WCDMA, Bluetooth, WiFi, etc., as well as wireless data network standards, including Bluetooth, WiFi, WiMAX, etc. The ASIP can be configured on the fly to handle the different RF frequency and performance requirements of each communication standard. The software based PLL of the present invention provides the flexibility for a more unified design that fits different applications.

In a second embodiment, the phase locked loop task is partitioned into a plurality of atomic operations. The ASIP is adapted to spread the computation of the atomic operations out over and completed within an entire PLL reference clock period. Each computation being performed at a much higher processor clock frequency than the PLL reference clock rate. This functions to significantly reduce the per cycle current

transient generated by the computations. Further, the frequency content of the current transients is at the higher processor clock frequency. This results in a significant reduction in spurs within sensitive portions of the output spectrum.

In other embodiments, a data stream based processor incorporating a combination parallel/pipelined architecture is optimized to perform data stream processing in an efficient manner. The parallel/pipelined processor provides for significantly higher processing speeds by combining multiple RCUs such that input data samples are input in parallel to all RCUs while computation results from one RCU are used by adjacent downstream RCUs. A register file provides storage for historical values while local storage in each RCU provides storage for temporary results.

An example application is provided of the software based phase locked loop incorporated in a single chip radio, e.g., Bluetooth, GSM, etc., that integrates the RF circuitry with the digital base band (DBB) circuitry on the same die.

Advantages of the software reconfigurable phase locked loop of the present invention include the following. Firstly, the invention enables all phase domain calculations to be performed within one reference clock cycle due to the use of the reconfigurable calculation unit optimized for performing PLL calculations serially at high frequency. Secondly, defining the ASIP instruction set in microcode stored in volatile or non-volatile memory makes it inherently software reconfigurable, permitting the microcode to be replaced without changing any lithography masks. This enables easier silicon debugging and multi-standard radio support.

Thirdly, the invention permits a significant reduction in silicon area. The invention trades the rate of operation for the amount of active implementation area required by the process of oversampling and function sharing. For an X factor increase in operational frequency, there is a complimentary X factor decrease in the required computational combinatorial logic area. An additional area is needed due to the overhead of computational unit multiplexing. While the storage area is constant, the net result is a significant reduction in overall implementation area required.

Fourthly, the invention enables a significant reduction of RF spurs in the sensitive frequency bands of a radio by changing the frequency of the switching logic gates. Prior art solutions perform PLL computations at relatively low rates, e.g., FREF of 26-38.8 MHz. The resulting switching current transients are mixed with the carrier and appear as frequency spurs at sensitive radio frequency bands. Considering GSM, for example, the most sensitive RX band is approximately 20 to 80 MHz away from the carrier. The invention performs the bulk of computations at oversampled rates, resulting in spurs outside sensitive regions. The amount of oversampling can be controlled (e.g., increased or decreased) to provide any desired frequency planning by changing the frequency of the processing clock. Fifthly, reduction in silicon area provided by the invention enables power routing and decoupling capacitance requirements to be relaxed.

Note that some aspects of the invention described herein may be constructed as software objects that are executed in embedded devices as firmware, software objects that are executed as part of a software application on either an embedded or non-embedded computer system such as a digital signal processor (DSP), microcomputer, minicomputer, microprocessor, etc. running a real-time operating system such as WinCE, Symbian, OSE, Embedded LINUX, etc. or non-real time operating system such as Windows, UNIX, LINUX, etc., or as soft core realized HDL circuits embodied in an Application. Specific Integrated Circuit (ASIC) or Field

5

Programmable Gate Array (FPGA), or as functionally equivalent discrete hardware components.

There is thus provided in accordance with the invention, a processor for use in a software based phase locked loop (PLL), comprising a first adder/subtractor operative to receive input data, a shifter operative to shift the output of the first adder/subtractor by a predetermined amount, a second adder/subtractor operative to receive the output of the shifter, a latch operative to store the output of the second adder/subtractor, a plurality of data paths connecting the first adder/subtractor, the shifter, the second adder/subtractor and the latch, the plurality of data paths configurable in accordance with one or more control signals and wherein the processor having an instruction set for controlling the first adder/subtractor, the shifter, the second adder/subtractor, the latch and the plurality of data paths.

There is also provided in accordance with the invention, a processor for use in a software based phase locked loop (PLL), comprising one or more computation units optimized for performing computations within a phase locked loop, wherein the one or more computation units are time-shared among all phase locked loop computations, data memory coupled to the one or more computation units, instruction memory coupled to the one or more computation units and operative to store instructions for implementing the phase locked loop, the instructions part of an instruction set and a decoder operative to generate one or more control signals for controlling the operation of the one or more computation units.

There is further provided in accordance with the invention, a processor based phase locked loop (PLL), comprising an oscillator operative to generate a radio frequency (RF) signal having a frequency determined in accordance with a tuning command input thereto, a processor operative to generate the tuning command, the processor comprising a reconfigurable calculation unit (RCU) operative to perform atomic operations required to implement the phase locked loop, data memory coupled to the reconfigurable calculation unit for storing phase locked loop state information, program memory coupled to the reconfigurable calculation unit for storing a plurality of instructions that when executed on the processor implement the phase locked loop and the processor having an instruction set, wherein each instruction is operative to perform an atomic operation of the phase locked loop.

There is also provided in accordance with the invention, a radio comprising a transmitter coupled to an antenna, the transmitter comprising a software based phase locked loop (PLL), the phase locked loop comprising an oscillator operative to generate a radio frequency (RF) signal having a frequency determined in accordance with a tuning command input thereto, a processor operative to generate the tuning command, the processor comprising a reconfigurable calculation unit (RCU) operative to perform atomic operations required to implement the phase locked loop, data memory coupled to the reconfigurable calculation unit for storing phase locked loop state information, program memory coupled to the reconfigurable calculation unit for storing a plurality of instructions that when executed on the processor implement the phase locked loop, the processor having an instruction set, wherein each instruction is operative to perform an atomic operation of the phase locked loop, a receiver

6

coupled to the antenna and a baseband processor coupled to the transmitter and the receiver.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is herein described, by way of example only, with reference to the accompanying drawings, wherein:

FIG. 1 is a block diagram illustrating a simplified block diagram of an example prior art phase locked loop (PLL) circuit;

FIG. 2 is a block diagram illustrating an example prior art generalized processing block using a dedicated hardware implementation;

FIG. 3 is a block diagram illustrating an example prior art generalized processing block using a processor based implementation;

FIG. 4 is a block diagram illustrating a single chip polar transceiver radio incorporating a software based ADPLL mechanism of the present invention;

FIG. 5 is a simplified block diagram illustrating an example mobile communication device incorporating the software based ADPLL mechanism of the present invention;

FIG. 6 is a block diagram illustrating functions of an example ADPLL-based polar transmitter suitable for use with the present invention;

FIG. 7 is a simplified block diagram illustrating an embodiment of the software based ADPLL incorporating a processor based phase domain calculator;

FIG. 8 is a block diagram illustrating an example embodiment of the phase domain calculator of the present invention in more detail;

FIG. 9 is a timing diagram illustrating the processing clock and frequency reference timing;

FIG. 10 is a block diagram illustrating an instruction view of the software based ADPLL architecture of the present invention;

FIG. 11 is a block diagram illustrating an example processor based software ADPLL architecture of the present invention;

FIG. 12A is a diagram illustrating the output and transfer function equations for the infinite impulse response (IIR) filter portion of the ADPLL;

FIG. 12B is an equivalent block diagram implementing the output equation shown in FIG. 12A;

FIG. 12C is an equivalent block diagram implementing the output equation shown in FIG. 12A whereby the multiplication operations have been replaced with shift operations;

FIG. 12D is a diagram illustrating the resultant reconfigurable calculation unit (RCU) implementing the output equation shown in FIG. 12A;

FIG. 13 is a block diagram illustrating an example RCU unit for implementing the ADPLL circuit;

FIG. 14 is a block diagram illustrating the RCU unit for implementing the F_Diff() instruction;

FIG. 15 is a block diagram illustrating the RCU unit for implementing the F_PheAcc() instruction;

FIG. 16 is a block diagram illustrating the RCU unit for implementing the F_IIR() instruction;

FIG. 17 is a block diagram illustrating the RCU unit for implementing the F_SHR() instruction;

FIG. 18 is a block diagram illustrating the RCU unit for implementing the F_IntAcc() instruction;

FIG. 19 is a block diagram illustrating the RCU unit for implementing the F_SHRAdd() instruction;

FIG. 20 is a timing diagram illustrating several ADPLL processing clock options and the current spikes resulting therefrom;

FIG. 21 is a timing diagram illustrating an example RF spectrum generated by a legacy ADPLL;

FIG. 22 is a timing diagram illustrating an example RF spectrum generated by the software based ADPLL of the present invention;

FIG. 23 is a flow diagram illustrating the RF spur reduction method of the present invention;

FIG. 24 is a block diagram illustrating an example ASIP based implementation of the software ADPLL of the present invention;

FIG. 25 is a block diagram illustrating an example DSP based implementation of the software ADPLL of the present invention;

FIG. 26 is a block diagram illustrating an example hardware state machine based implementation of the software ADPLL of the present invention;

FIG. 27 is a block diagram illustrating a first example implementation of the state machine used to implement the software ADPLL of the present invention;

FIG. 28 is a block diagram illustrating a second example implementation of the state machine used to implement the software ADPLL of the present invention;

FIG. 29 is a diagram illustrating an example superscalar pipeline architecture;

FIG. 30 is a block diagram illustrating the single instruction multiple data (SIMD) technique;

FIG. 31 is a block diagram illustrating an example data stream processing model;

FIG. 32 is a block diagram illustrating an example RCU based data stream processor for implementing the model of FIG. 31;

FIG. 33 is a block diagram illustrating an example two-stage parallel/pipelined RCU architecture of a data stream processor;

FIG. 34 is a block diagram illustrating an example data stream processing path incorporating forking and merging;

FIG. 35 is a block diagram illustrating an example parallel/pipelined architecture incorporating forking handling capability;

FIG. 36 is a block diagram illustrating control aspect of an example RCU;

FIG. 37 is a block diagram illustrating control aspect of an example three-stage pipelined RCU;

FIG. 38 is a block diagram illustrating an example three-stage parallel/pipelined RCU architecture; and

FIG. 39 is a diagram illustrating the data processing within the three-stage parallel/pipeline processor of FIG. 38.

DETAILED DESCRIPTION OF THE INVENTION

Notation Used Throughout

The following notation is used throughout this document.

Term	Definition
AC	Alternating Current
ACL	Asynchronous Connectionless Link
ACW	Amplitude Control Word
ADC	Analog to Digital Converter
ADPLL	All Digital Phase Locked Loop
ALU	Arithmetic Logic Unit
AM	Amplitude Modulation
ASIC	Application Specific Integrated Circuit
ASIP	Application Specific Instruction-set Processor
AVI	Audio Video Interface

-continued

Term	Definition
5 AWS	Advanced Wireless Services
BIST	Built-In Self Test
BMP	Windows Bitmap
BPF	Band Pass Filter
CMOS	Complementary Metal Oxide Semiconductor
CPU	Central Processing Unit
10 CU	Control Unit
CW	Continuous Wave
DAC	Digital to Analog Converter
dB	Decibel
DBB	Digital Baseband
DC	Direct Current
DCO	Digitally Controlled Oscillator
DCXO	Digitally Controlled Crystal Oscillator
DPA	Digitally Controlled Power Amplifier
DRAC	Digital to RF Amplitude Conversion
DRP	Digital RF Processor or Digital Radio Processor
DSL	Digital Subscriber Line
DSP	Digital Signal Processor
EDGE	Enhanced Data Rates for GSM Evolution
EDR	Enhanced Data Rate
EEPROM	Electrically Erasable Programmable Read Only Memory
EPROM	Erasable Programmable Read Only Memory
eSCO	Extended Synchronous Connection-Oriented
25 FCC	Federal Communications Commission
FCW	Frequency Command Word
FIB	Focused Ion Beam
FM	Frequency Modulation
FPGA	Field Programmable Gate Array
FSM	Finite State Machine
GMSK	Gaussian Minimum Shift Keying
GPRS	General Packet Radio Service
GPS	Global Positioning System
GSM	Global System for Mobile communications
HB	High Band
HDL	Hardware Description Language
35 HFP	Hands Free Protocol
I/F	Interface
IC	Integrated Circuit
IEEE	Institute of Electrical and Electronics Engineers
IIR	Infinite Impulse Response
JPG	Joint Photographic Experts Group
40 LAN	Local Area Network
LB	Low Band
LDO	Low Drop Out
LNA	Low Noise Amplifier
LO	Local Oscillator
LPF	Low Pass Filter
MAC	Media Access Control
MAP	Media Access Protocol
MBOA	Multiband OFDM Alliance
MIM	Metal Insulator Metal
Mod	Modulo
MOS	Metal Oxide Semiconductor
MP3	MPEG-1 Audio Layer 3
50 MPG	Moving Picture Experts Group
MUX	Multiplexer
NZIF	Near Zero IF
OFDM	Orthogonal Frequency Division Multiplexing
PA	Power Amplifier
PAN	Personal Area Network
55 PC	Personal Computer
PCI	Personal Computer Interconnect
PD	Phase Detector
PDA	Personal Digital Assistant
PE	Phase Error
PHE	Phase Error
60 PLL	Phase Locked Loop
PM	Phase Modulation
PPA	Pre-Power Amplifier
QoS	Quality of Service
RAM	Random Access Memory
RCU	Reconfigurable Calculation Unit
RF	Radio Frequency
65 RFBIST	RF Built-In Self Test
RMS	Root Mean Squared

-continued

Term	Definition
ROM	Read Only Memory
SAM	Sigma-Delta Amplitude Modulation
SAW	Surface Acoustic Wave
SCO	Synchronous Connection-Oriented
SEM	Spectral Emission Mask
SIM	Subscriber Identity Module
SoC	System on Chip
SRAM	Static Read Only Memory
SYNTH	Synthesizer
TDC	Time to Digital Converter
TDD	Time Division Duplex
TV	Television
UART	Universal Asynchronous Transmitter/Receiver
UGS	Unsolicited Grant Services
UMTS	Universal Mobile Telecommunications System
USB	Universal Serial Bus
UWB	Ultra Wideband
VCO	Voltage Controlled Oscillator
WCDMA	Wideband Code Division Multiple Access
WiFi	Wireless Fidelity
WiMAX	Worldwide Interoperability for Microwave Access
WiMedia	Radio platform for UWB
WLAN	Wireless Local Area Network
WMA	Windows Media Audio
WMAN	Wireless Metropolitan Area Network
WMV	Windows Media Video
WPAN	Wireless Personal Area Network
XOR	Exclusive Or
ZIF	Zero IF

DETAILED DESCRIPTION OF THE INVENTION

The present invention is a novel and useful apparatus for and method of software based phase locked loop (PLL). The processor-based (or alternatively software-based or highly reconfigurable) PLL (i.e. all digital phase-locked loop or ADPLL) architecture can be used for RF frequency synthesis in radio transceivers and other applications. The software based phase locked loop of the present invention incorporates a reconfigurable calculation unit (RCU) that is programmed to sequentially perform all the atomic (or substantially atomic wherein the level of 'atomicity' is conveniently determined during architectural optimization) operations of a phase locked loop or of any other desired task. The RCU is a key component in an application specific instruction-set processor (ASIP). The ASIP includes an instruction set that comprises instructions optimized to perform the individual atomic operations of a phase locked loop. The PLL operations perform data stream processing of various PLL signals, such as frequency command word, frequency error, phase error, tuning word, etc.

An example application is provided of the software based phase locked loop incorporated in a single chip radio, e.g., Bluetooth, GSM, etc., that integrates the RF circuitry with the digital base band (DBB) circuitry on the same die.

Although the software based phase locked loop mechanism is applicable to numerous wireless communication standards and can be incorporated in numerous types of wireless or wired communication devices such a multimedia player, mobile station, cellular phone, PDA, DSL modem, WPAN device, etc., it is described in the context of a digital RF processor (DRP) based transceiver that may be adapted to comply with a particular wireless communications standard such as GSM, Bluetooth, EDGE, WCDMA, WLAN, WiMax, etc. It is appreciated, however, that the invention is not limited to use with any particular communication standard and may be used in optical, wired and wireless applications. Further,

the invention is not limited to use with a specific modulation scheme but is applicable to numerous modulation schemes.

Note that throughout this document, the term communications device is defined as any apparatus or mechanism adapted to transmit, receive or transmit and receive data through a medium. The term communications transceiver is defined as any apparatus or mechanism adapted to transmit and receive data through a medium. The communications device or communications transceiver may be adapted to communicate over any suitable medium, including wireless or wired media. Examples of wireless media include RF, infrared, optical, microwave, UWB, Bluetooth, WiMAX, WiMedia, WiFi, or any other broadband medium, etc. Examples of wired media include twisted pair, coaxial, optical fiber, any wired interface (e.g., USB, Firewire, Ethernet, etc.). The term Ethernet network is defined as a network compatible with any of the IEEE 802.3 Ethernet standards, including but not limited to 10 Base-T, 100Base-T or 1000Base-T over shielded or unshielded twisted pair wiring. The terms communications channel, link and cable are used interchangeably. The notation DRP is intended to denote either a Digital RF Processor or Digital Radio Processor. References to a Digital RF Processor infer a reference to a Digital Radio Processor and vice versa.

The term multimedia player or device is defined as any apparatus having a display screen and user input means that is capable of playing audio (e.g., MP3, WMA, etc.), video (AVI, MPG, WMV, etc.) and/or pictures (JPG, BMP, etc.). The user input means is typically formed of one or more manually operated switches, buttons, wheels or other user input means. Examples of multimedia devices include pocket sized personal digital assistants (PDAs), personal media player/recorders, cellular telephones, handheld devices, and the like.

Some portions of the detailed descriptions which follow are presented in terms of procedures, logic blocks, processing, steps, and other symbolic representations of operations on data bits within a computer memory. These descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. A procedure, logic block, process, etc., is generally conceived to be a self-consistent sequence of steps or instructions leading to a desired result. The steps require physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared and otherwise manipulated in a computer system. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, bytes, words, values, elements, symbols, characters, terms, numbers, or the like.

It should be born in mind that all of the above and similar terms are to be associated with the appropriate physical quantities they represent and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as 'processing,' 'computing,' 'calculating,' 'determining,' 'displaying' or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

The invention can take the form of an entirely hardware embodiment, an entirely general-purpose software embodi-

ment or an embodiment containing a combination of hardware and software elements. In one embodiment, a portion of the mechanism of the invention is implemented in software, which includes but is not limited to firmware, resident software, object code, assembly code, microcode, etc.

Furthermore, the invention can take the form of a computer program product accessible from a computer-usable or computer-readable medium providing program code for use by or in connection with a computer or any instruction execution system. For the purposes of this description, a computer-usable or computer readable medium is any apparatus that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, or device, e.g., floppy disks, removable hard drives, computer files comprising source code or object code, flash semiconductor memory (USB flash drives, etc.), ROM, EPROM, or other semiconductor memory devices.

Single Chip Radio

A block diagram illustrating a single chip radio incorporating a software based ADPLL mechanism of the present invention is shown in FIG. 4. For illustration purposes only, the transmitter, as shown, is adapted for the GSM/EDGE/WCDMA cellular standards. It is appreciated, however, that one skilled in the communication arts can adapt the transmitter and receiver illustrated herein to other modulations and communication standards as well without departing from the spirit and scope of the present invention.

The radio, generally referenced 30, comprises a radio integrated circuit 31 coupled to a crystal 38, RF front end module 46 coupled to an antenna 44, and battery management circuit 32 coupled to battery 68. The radio chip 31 comprises a script processor 60, digital baseband (DBB) processor 61, memory 62 (e.g., static RAM), TX block 42, RX block 58, digitally controlled crystal oscillator (DCXO) 50, slicer 51, power management unit 34 and RF built-in self test (BIST) 36. The TX block comprises high speed and low speed digital logic block 40 including $\Sigma\Delta$ modulators (not shown), phase domain calculator (ASIP) 52 for performing data stream processing of the PLL signals, digitally controlled oscillator (DCO) 56, accumulator 59, sampler 69 and digitally controlled power amplifier (DPA) 48. The RX block comprises a low noise transconductance amplifier 63, current sampler 64, discrete time processing block 65, analog to digital converter (ADC) 66 and digital logic block 67. Note that the data stream processing is not limited to the phase domain calculations of the PLL, but could be applied to the transmit modulation operation, which could be performed in the script processor 60 or internally in the transmitter 42. The data stream processing principles could also be applied to the receive signals in the RX digital logic block 67.

The Digital RF Processor (DRP) principles presented herein have been used to develop three generations of a Digital RF Processor (DRP): single-chip Bluetooth, GSM and GSM/EDGE radios realized in 130 nm, 90 nm and 65 nm digital CMOS process technologies, respectively. This architecture is also used as the foundation for a UMTS single-chip radio manufactured using a 45 nm CMOS process. The common architecture is highlighted with features added specific to the cellular radio. The all digital phase locked loop (ADPLL) based transmitter employs a polar architecture with all digital phase/frequency and amplitude modulation paths. The receiver employs a discrete-time architecture in which the RF signal is directly sampled and processed using analog and digital signal processing techniques.

A key component is the digitally controlled oscillator (DCO) 56, which avoids any analog tuning controls. A digitally-controlled crystal oscillator (DCXO) generates a high-quality base station-synchronized frequency reference such that the transmitted carrier frequencies and the received symbol rates are accurate to within 0.1 ppm. Fine frequency resolution for both DCO and DCXO is achieved through high-speed $\Sigma\Delta$ dithering of their varactors. Digital logic built around the DCO realizes an all-digital PLL (ADPLL) that is used as a local oscillator for both the transmitter and receiver. The polar transmitter architecture utilizes the wideband direct frequency modulation capability of the ADPLL and a digitally controlled power amplifier (DPA) 48 for the amplitude modulation. The DPA operates in near-class-E mode and uses an array of nMOS transistor switches to regulate the RF amplitude. It is followed by a matching network and an external front-end module 46, which comprises a power amplifier (PA), a transmit/receive switch for the common antenna 44 and RX surface acoustic wave (SAW) filters. Fine amplitude resolution is achieved through high-speed $\Sigma\Delta$ dithering of the DPA nMOS transistors.

The receiver 58 employs a discrete-time architecture in which the RF signal is directly sampled at the Nyquist rate of the RF carrier and processed using analog and digital signal processing techniques. The transceiver is integrated with a script processor 60, dedicated digital base band processor 61 (i.e. ARM family processor and/or DSP) and SRAM memory 62. The script processor handles various TX and RX calibration, compensation, sequencing and lower-rate data path tasks and encapsulates the transceiver complexity in order to present a much simpler software programming model.

The frequency reference (FREF) is generated on-chip by a 26 MHz (or any other desired frequency, such as 13 or 38.4 MHz) digitally controlled crystal oscillator (DCXO) 50, which provides negative resistance to sustain the oscillations. The output of the DCXO is coupled to slicer 51. The output of the slicer is input to the phase domain calculator which comprises a software based PLL in accordance with the invention and described in more detail infra.

An integrated power management (PM) system 34 is connected to an external battery management circuit 32 that conditions and stabilizes the supply voltage. The PM comprises multiple low drop out (LDO) regulators that provide internal supply voltages and also isolate supply noise between circuits, especially protecting the DCO. The RF built-in self-test (RFBIST) 36 performs autonomous phase noise and modulation distortion testing, various loopback configurations for bit-error rate measurements and implements the DPA calibration and BIST mechanism. The transceiver is integrated with the digital baseband, SRAM memory in a complete system-on-chip (SoC) solution. Almost all the clocks on this SoC are derived from and are synchronous to the RF oscillator clock. This helps to reduce susceptibility to the noise generated through clocking of the massive digital logic.

The transmitter comprises a polar architecture in which the amplitude and phase/frequency modulations are implemented in separate paths. Transmitted symbols generated in the digital baseband (DBB) processor are first pulse-shape filtered in the Cartesian coordinate system. The filtered in-phase (I) and quadrature (Q) samples are then converted through a CORDIC algorithm into amplitude and phase samples of the polar coordinate system. The phase is then differentiated to obtain frequency deviation. The polar signals are subsequently conditioned through signal processing to

13

sufficiently increase the sampling rate in order to reduce the quantization noise density and lessen the effects of the modulating spectrum replicas.

A more detailed description of the operation of the ADPLL can be found in U.S. Patent Publication No. 2006/0033582A1, published Feb. 16, 2006, to Staszewski et al., entitled "Gain Calibration of a Digital Controlled Oscillator;" U.S. Patent Publication No. 2006/0038710A1, published Feb. 23, 2006, Staszewski et al., entitled "Hybrid Polar/Cartesian Digital Modulator" and U.S. Pat. No. 6,809,598, to Staszewski et al., entitled "Hybrid Of Predictive And Closed-Loop Phase-Domain Digital PLL Architecture," all of which are incorporated herein by reference in their entirety.

Mobile Device/Cellular Phone/PDA System

A simplified block diagram illustrating an example mobile communication device incorporating the software based ADPLL mechanism of the present invention is shown in FIG. 5. The communication device may comprise any suitable wired or wireless device such as a multimedia player, mobile station, mobile device, cellular phone, PDA, wireless personal area network (WPAN) device, Bluetooth EDR device, etc. For illustration purposes only, the communication device is shown as a cellular phone or smart phone. Note that this example is not intended to limit the scope of the invention as the software based ADPLL mechanism of the present invention can be implemented in a wide variety of wireless and wired communication devices.

The cellular phone, generally referenced **70**, comprises a baseband processor or CPU **71** having analog and digital portions. The basic cellular link is provided by the RF transceiver **94** and related one or more antennas **96, 98**. A plurality of antennas is used to provide antenna diversity which yields improved radio performance. The cell phone also comprises internal RAM and ROM memory **110**, Flash memory **112** and external memory **114**.

In accordance with the invention, the RF transceiver **94** comprises the software reconfigurable ADPLL of the present invention. In operation, the software reconfigurable ADPLL mechanism may be implemented as dedicated hardware, as software executed as a task on the baseband processor **71** or dedicated processor or a combination of hardware and software. Implemented as a software task, the program code operative to implement the software reconfigurable ADPLL mechanism of the present invention is stored in one or more memories **110, 112, 114** or in on-chip volatile or non-volatile memory.

Several user interface devices include microphone **84**, speaker **82** and associated audio codec **80**, a keypad for entering dialing digits **86**, vibrator **88** for alerting a user, camera and related circuitry **100**, a TV tuner **102** and associated antenna **104**, display **106** and associated display controller **108** and GPS receiver **90** and associated antenna **92**.

A USB interface connection **78** provides a serial link to a user's PC or other device. An FM receiver **72** and antenna **74** provide the user the ability to listen to FM broadcasts. WLAN radio and interface **76** and antenna **77** provide wireless connectivity when in a hot spot or within the range of an ad hoc, infrastructure or mesh based wireless LAN network. A Bluetooth EDR radio and interface **73** and antenna **75** provide Bluetooth wireless connectivity when within the range of a Bluetooth wireless network. Further, the communication device **70** may also comprise a WiMAX radio and interface **123** and antenna **125**. SIM card **116** provides the interface to a user's SIM card for storing user data such as address book entries, etc. The communication device **70** also comprises an

14

Ultra Wideband (UWB) radio and interface **83** and antenna **81**. The UWB radio typically comprises an MBOA-UWB based radio.

Portable power is provided by the battery **124** coupled to battery management circuitry **122**. External power is provided via USB power **118** or an AC/DC adapter **120** connected to the battery management circuitry which is operative to manage the charging and discharging of the battery **124**.

Example ADPLL Polar Transmitter

A block diagram illustrating an example ADPLL-based polar transmitter for wireless applications is shown in FIG. 6. The example ADPLL shown is used as the basis for a software based ADPLL mechanism described in more detail infra. A more detailed description of the operation of the ADPLL can be found in U.S. Patent Publication No. 2006/0033582A1, published Feb. 16, 2006, to Staszewski et al., entitled "Gain Calibration of a Digital Controlled Oscillator;" U.S. Patent Publication No. 2006/0038710A1, published Feb. 23, 2006, Staszewski et al., entitled "Hybrid Polar/Cartesian Digital Modulator" and U.S. Pat. No. 6,809,598, to Staszewski et al., entitled "Hybrid Of Predictive And Closed-Loop Phase-Domain Digital PLL Architecture," all of which are incorporated herein by reference in their entirety.

For illustration purposes only, the transmitter, as shown, is adapted for the GSM/EDGE/WCDMA cellular standards. It is appreciated, however, that one skilled in the communication arts can adapt the transmitter illustrated herein to other modulations and communication standards as well without departing from the spirit and scope of the present invention.

The transmitter, generally referenced **130**, is well-suited for a deep-submicron CMOS implementation. The transmitter comprises a complex pulse shaping filter **168**, amplitude modulation (AM) block **169** and ADPLL **132**. The circuit **130** is operative to perform complex modulation in the polar domain in addition to the generation of the local oscillator (LO) signal for the receiver. All clocks in the system are derived directly from this source. Note that the transmitter is constructed using digital techniques that exploit the high speed and high density of the advanced CMOS, while avoiding problems related to voltage headroom. The ADPLL circuit replaces a conventional RF synthesizer architecture (based on a voltage-controlled oscillator (VCO) and a phase/frequency detector and charge-pump combination), with a digitally controlled oscillator (DCO) **148** and a time-to-digital converter (TDC) **162**. All inputs and outputs are digital and some even at multi-GHz frequency.

The core of the ADPLL is a digitally controlled oscillator (DCO) **148** adapted to generate the RF oscillator clock CKV. The oscillator core (not shown) operates at a multiple of the 1.6-2.0 GHz (e.g., 4) high band frequency or at a multiple of the 0.8-1.0 GHz low band frequency (e.g., 8). Note that typically, the multiple is a power-of-two but any other suitable integer or even fractional frequency relationship may be advantageous. The output of the DCO is then divided for precise generation of RX quadrature signals, and for use as the transmitter's carrier frequency. The single DCO is shared between transmitter and receiver and is used for both the high frequency bands (HB) and the low frequency bands (LB). In addition to the integer control of the DCO, at least 3-bits of the minimal varactor size used are dedicated for $\Sigma\Delta$ dithering in order to improve frequency resolution. The DCO comprises a plurality of varactor banks, which may be realized as n-poly/n-well inversion type MOS capacitor (MOSCAP) devices or Metal Insulator Metal (MIM) devices that operate in the flat regions of their C-V curves to assist digital control. The

15

output of the DCO is a modulated digital signal at f_{RF} . This signal is input to the pre-power amplifier (PPA) **152**. It is also input to the RF low band pre-power amplifier **154** after divide by two via divider **150**.

The expected variable frequency f_V is related to the reference frequency f_R by the frequency command word (FCW).

$$FCW[k] = \frac{E(f_V[k])}{f_R} \quad (1)$$

The FCW is time variant and is allowed to change with every cycle $T_R=1/f_R$ of the frequency reference clock. With $W_F=24$ the word length of the fractional part of FCW, the ADPLL provides fine frequency control with 1.5 Hz accuracy, according to:

$$\Delta f_{res} = \frac{f_R}{2^{W_F}} \quad (2)$$

The number of integer bits $W_I=8$ has been chosen to fully cover the GSM/EDGE and partial WCDMA band frequency range of $f_V=1,600$ -2,000 MHz with an arbitrary reference frequency $f_R \geq 8$ MHz.

The ADPLL operates in a digitally-synchronous fixed-point phase domain as follows: The variable phase accumulator **156** determines the variable phase $R_V[i]$ by counting the number of rising clock transitions of the DCO oscillator clock CKV as expressed below.

$$R_V[i] = \sum_{l=0}^i 1 \quad (3)$$

The index i indicates the DCO edge activity. The variable phase $R_V[i]$ is sampled via sampler **158** to yield sampled FREF variable phase $R_V[k]$, where k is the index of the FREF edge activity. The sampled FREF variable phase $R_V[k]$ is fixed-point concatenated with the normalized time-to-digital converter (TDC) **162** output $\epsilon[k]$. The TDC measures and quantizes the time differences between the frequency reference FREF and the DCO clock edges. The sampled differentiated (via block **160**) variable phase is subtracted from the frequency command word (FCW) by the digital frequency detector **138**. The frequency error $f_E[k]$ samples

$$f_E[k] = FCW - [(R_V[k] - \epsilon[k]) - (R_V[k-1] - \epsilon[k-1])] \quad (4)$$

are accumulated via the frequency error accumulator **140** to create the phase error $\phi_E[k]$ samples

$$\phi_E[k] = \sum_{l=0}^k f_E[l] \quad (5)$$

which are then filtered by a fourth order IIR loop filter **142** and scaled by a proportional loop attenuator α . A parallel feed with coefficient ρ adds an integrated term to create type-II loop characteristics which suppress the DCO flicker noise.

16

The IIR filter is a cascade of four single stage filters, each satisfying the following equation:

$$y[k] = (1-\lambda)y[k-1] + \lambda x[k] \quad (6)$$

wherein

$x[k]$ is the current input;

$y[k]$ is the current output;

k is the time index;

λ is the configurable coefficient;

The 4-pole IIR loop filter attenuates the reference and TDC quantization noise with an 80 dB/dec slope, primarily to meet the GSM/EDGE spectral mask requirements at 400 kHz offset. The filtered and scaled phase error samples are then multiplied by the DCO gain K_{DCO} normalization factor f_R $\hat{K}_{DCO}$ via multiplier **146**, where f_R is the reference frequency and $\hat{K}_{DCO}$ is the DCO gain estimate, to make the loop characteristics and modulation independent from K_{DCO} . The modulating data is injected into two points of the ADPLL for direct frequency modulation, via adders **136** and **144**. A hitless gear-shifting mechanism for the dynamic loop bandwidth control serves to reduce the settling time. It changes the loop attenuator α several times during the frequency locking while adding the $(\alpha_1/\alpha_2-1)\phi_1$ dc offset to the phase error, where indices **1** and **2** denote before and after the event, respectively. Note that $\phi_1=\phi_2$, since the phase is to be continuous.

The frequency reference FREF is input to the retimer **166** and provides the clock for the TDC **162**. The FREF input is resampled by the RF oscillator clock CKV via retimer block **166** which may comprise a flip flop or register clocked by the reference frequency FREF. The resulting retimed clock (CKR) is distributed and used throughout the system. This ensures that the massive digital logic is clocked after the quiet interval of the phase error detection by the TDC. Note that in the example embodiment described herein, the ADPLL is a discrete-time sampled system implemented with all digital components connected with all digital signals.

Software Based ADPLL Architecture

A simplified block diagram illustrating an embodiment of the software based ADPLL incorporating a processor based phase domain calculator is shown in FIG. 7. The ADPLL circuit, generally referenced **180**, comprises a phase domain calculator **184**, DCO **186**, integer feedback block **188**, fractional feedback block **189** and programmable fractional-N clock divider **182**.

In operation, the phase domain calculator replaces the conventional ADPLL circuit with a software based ADPLL. As with the conventional ADPLL, it is operative to generate the DLO update that is input to the DCO **186** which in turn generates the RF output frequency clock CKV. The phase domain calculator receives the FCW commands, variable phase information (i.e. integer and fractional feedback) and the reference frequency clock FREF, which typically ranges between 13 and 52 MHz. The processing clock output of the programmable clock divider **182** runs at a frequency significantly higher than FREF, such as in the range 200 to 600 MHz, for example.

In accordance with the invention, the phase domain calculator performs the ADPLL operations serially rather than in parallel. In order to complete the ADPLL computation within reference clock cycle, the much faster processor clock is used to clock the phase domain calculator internal circuitry.

The solution uses a reconfigurable computational unit (RCU) or ALU (described infra) that is time shared for most

or all computations within the ADPLL. The RCU and its related configuration control logic (constituting a special purpose microcomputer) replaces the dedicated and distributed random logic within a conventional ADPLL. The RCU is controlled via microcode stored in on-chip memory such as random access memory (RAM), read only memory (ROM), Flash memory, etc. Since the computational unit is time shared among most operations, it is operated at a much higher clock rate than the conventional ADPLL which performs all calculations in parallel using dedicated hardware circuits. The RCU circuitry is optimized to perform all the required ADPLL atomic computations within one reference clock cycle.

A block diagram illustrating an example embodiment of the phase domain calculator of the present invention in more detail is shown in FIG. 8. The phase domain calculator, generally referenced **190**, comprises an ALU (or RCU) **202**, instruction and data memory **192**, register file **194**, sequencer **196**, latches **208**, **206**, multiplexers **209**, **198**, **200**, **204**.

The instructions for implementing the ADPLL operation are stored in the instruction memory. Instructions are input to the sequencer which performs the instruction decoding and generates the appropriate signals to execute each instruction. The register file stores intermediate values calculated by the ALU.

A timing diagram illustrating the processing clock and reference frequency timing is shown in FIG. 9. As shown, the processing clock **210**, used to clock the memory, sequencer and register file, is at a significantly higher clock rate than the reference clock FREF **212**. This is required in order the complete an operation cycle of the ADPLL within a single reference clock period.

A block diagram illustrating an instruction view of the software based ADPLL architecture of the present invention is shown in FIG. 10. The circuit, generally referenced **220**, comprises a phase calculation unit **222**, gain normalization **238**, DCO **240**, gain calibration **242**, accumulator **244** and variable phase sampler **246**. The phase calculation unit shows an instruction view of the ADPLL architecture performed in software. In particular, the phase calculation unit **222** comprises a phase detector **224**, frequency error accumulator **226**, IIR loop filter **228**, a proportional gain multiplier **230**, IIR filter **232**, adder **234**, integral gain multiplier **236**, offset phase error adder **248**, phase error accumulator **250** and integral or p gain multiplier **252**. Both proportional and integral multipliers preferably use power-of-two arithmetic so that their respective multipliers could be implemented as bit-shift operators. For non-power-of-two arithmetic, full multipliers can be used with the consequent increase in circuit cost and complexity.

The ADPLL circuit shown is an ADPLL architecture that is commonly used in wireless applications for RF frequency generation. In this digital architecture, the traditional VCO is replaced with a digitally controlled oscillator (DCO) and the oscillating frequency of the DCO is controlled by a frequency command word (FCW) instead of the reference clock, as described in detail supra. The phase detecting and filtering parts are all digital with intensive digital signal processing involved as highlighted.

In operation, the ADPLL operation is partitioned into a plurality of atomic operations, wherein each atomic operation performs a complete processing step within the ADPLL. For example, an adding operation representing adder **224** comprises one atomic operation. Similarly, accumulation block **226** and each of four elemental first-order IIR operations in IIR filter block **228** also comprise a single atomic operation each. Each atomic operation is performed by a separate

instruction. Individual instructions are provided for each atomic operation in the ADPLL (or any other task).

The computation elements within the phase calculation block **222** are the elements implemented and executed as software instructions. The computations within block **222** can be expressed as pseudo code as shown in the following Listing 1.

Listing 1: Sequential pseudo-code for ADPLL phase computation

```

L1: y16 = y15 - S1          // S1 is the store element
    S1 = y15                // delay
    y1 = FCW - y16
    y2 = y2 + y1            // accumulation
    y3 = IIR (a1, y2, y3)   // 4th order infinite impulse response(IIR)
                                filter
    y4 = IIR (a2, y3, y4)
    y5 = IIR (a3, y4, y5)
    y6 = IIR (a4, y5, y6)
    y7 = SHL (y6, alpha)   // shift operation
    y11 = IIR (gs, y7, y11)
    y9 = y6 - PhE
    y10 = y10 + y9          // implies accumulation
    y12 = SHL (y10, rho)
    y13 = y11 + y12
    y14 = y13 * gain
    jump L1                // operates in a continuous loop
  
```

As described supra, in prior art ADPLL circuits, all ADPLL phase computations are implemented using dedicated hardware (i.e. a hardwired design), which limits the ability to adjust the ADPLL algorithm. In accordance with the software based architecture of the present invention, these ADPLL functions are integrated within a processor using shared hardware thereby providing significant flexibility to the ADPLL algorithm.

Processor Based ADPLL Architecture

A block diagram illustrating an example processor based software ADPLL architecture of the present invention is shown in FIG. 11. The ADPLL circuit, generally referenced **260**, comprises an application specific instruction-set processor (ASIP) **262**, which is used to realize data stream processing of the present invention, DCO **286**, accumulator **288** and variable phase sampler **289**. The ASIP **262** comprises instruction memory **264**, fetch block **266**, decode block **268**, data bus **278**, RCU **270**, S-unit **272**, L-unit **274**, A-unit **276**, register file **280**, data memory **282** and interface **284**.

All ADPLL computations as delineated in Listing 1 above are incorporated into a so-called Application Specific Instruction-set Processor (ASIP) **262**. It is appreciated that a general purpose processor may also be used to perform the ADPLL operation. An ASIP, however, is far more efficient due to the instruction set being adapted to perform a small but dedicated set of atomic (or substantially atomic) operations.

In operation, the ASIP processor stores the ADPLL software instructions in the instruction memory **264**. The instructions are then fetched from the instruction memory via fetch block **266** and fed into the decoding block **268**. All required control signals are generated through the decoding block to control the operation of the various computational units, including the A-Unit **276** for performing arithmetic operations, e.g., addition, subtraction, etc., L-Unit **274** for performing logic operations, e.g., AND, OR, XOR, etc. and S-Unit **272** for performing data storage and movement operations. A reconfigurable calculation unit (RCU) **270** is constructed to provide application specific instructions for the ADPLL. The

RCU, for example, is operative to implement the 1-stage IIR filtering operation within a single instruction. It is noted that the application driven customized instruction set is what differentiates an ASIP from a general purpose processor which performs ADPLL computations much less efficiently to the extent that it may not even be able to complete the necessary computations within the reference clock period.

The ASIP processor is operative to read the FCW and variable phase (PhV) inputs, sequentially perform all the computations (i.e. atomic operations) required for the ADPLL as presented in Listing 1 within one system reference clock cycle and send the resulting tuning word DCO_TUNE (i.e. DLO update) to the DCO which in turn uses the tuning word to adjust its output frequency. It is important to note that all the computations are performed via the programmed software stored in instruction memory of the ASIP. Note also that the majority of the computations are performed by the RCU, which is designed specifically to implement the atomic operations of the targeted ADPLL application.

Reconfigurable Calculation Unit (RCU)

The structure of the RCU will now be described in more detail. As described above, all the computations in Listing 1 are described in terms of atomic arithmetic operations, such as additions, subtractions, shifting, multiplications, etc. and as more complicated operations, such as IIR filtering. A prior art hardwired implementation simply instantiates the number of hardware operators equal to the number of atomic operations required by the ADPLL algorithm. This, however, has its drawbacks as discussed supra.

The ASIP based design of the present invention utilizes one or more reconfigurable computational units that are used to perform all arithmetic operations implementing the ADPLL algorithm. This computational unit is "recycled" sequentially among all the arithmetic operations within a single cycle of the ADPLL system clock. The internal state of the ADPLL is stored between clock cycles in internal storage elements (i.e. register file, data memory, etc.). An important aspect of the RCU design is the greatly increased application efficiency along with a maximization of resource reuse.

Detailed knowledge of the task to be implemented (e.g., ADPLL algorithm) is important in creating the functionality of the computational unit in order to optimize its complexity and system throughput requirements. In the case of an ADPLL algorithm, all the elemental ADPLL computations (except for IIR filtering) are either add/subtract or power-of-two multiply operation. Therefore, the invention provides for a single IIR computation that has been optimized to be represented as a single operation in the computational unit.

FIGS. 12A, 12B, 12C and 12D illustrate the mapping process and the resulting configuration of the RCU. In particular, FIG. 12A is a diagram illustrating the output and transfer function equations for the infinite impulse response (IIR) filter portion of the ADPLL. FIG. 12B is an equivalent block diagram implementing the output equation shown in FIG. 12A. The circuit, generally referenced 290, comprises multipliers 292, 298, adders 294, 299 and unit delay 296.

FIG. 12C is an equivalent block diagram implementing the output equation shown in FIG. 12A whereby the multiplication operations have been replaced with shift operations. The circuit, generally referenced 300, comprises shift operations 302, 309, adders 304, 308 and unit delay 306. FIG. 12D is a diagram illustrating the resultant reconfigurable calculation unit (RCU) implementing the output equation shown in FIG. 12A. The RCU, generally referenced 310, comprises shifter 312, 318 and adders 314, 316. The unit delay element is

replaced with data lines Rb and Rd for reading and writing to an external memory such as the register file or data memory.

Thus, for application to an ADPLL, the multiplication operation in the IIR filter is simplified with a shifting operation. The final RCU in 12D is a pure computation unit without any storage element. To map the IIR equation below

$$y_k = (1 - a) \cdot y_{k-1} + a \cdot x_k \quad (7)$$

$$H(Z) = \frac{a}{1 - (1 - a) \cdot z^{-1}}$$

to the RCU unit, the following applies:

- x_k maps to RCU input Ra;
- y_{k-1} maps to RCU input Rb;
- y_k maps to RCU output Rd;

The RCU unit is made further configurable to accommodate all the main computations for the ADPLL in an efficient manner. A block diagram illustrating an example RCU unit 322 for implementing the ADPLL circuit way is shown in FIG. 13. The RCU 322 and related circuitry, generally referenced 320, comprises a plurality of elements as follows: two data inputs (Ra, Rb), two data outputs (Rd_st, Rd), three atomic computation units (first addition/subtraction 324, shifter 326, second addition/subtraction 328), two latch/storage elements (input latch S_d 336 and output latch Rd 330), local registers for data storage, e.g., shift amount a, FCW 340, PhE 342, multiplexers for data steering 332, 334, 338 and control/configure signals.

In operation, the RCU takes input data Ra and goes through the first addition/subtraction followed by a shifting operation followed by a second addition/subtraction. The data is then sent to the outside register file Rd_st or latched (Rd) for the next computation. An input latching element (S_d) is included as part of a differentiation operation. All the computation units including the data paths inside the RCU are configurable with the control/configure signals generated by the ASIP decoding block.

Table 1 below shows the mapping of the customized instruction set provided by the RCU and their corresponding targeted computations in the ADPLL.

TABLE 1

Customized instructions and their corresponding computations in the ADPLL	
Computations in ADPLL	Customized Instructions
L1: $y16 = y15 - S_d$	F_Diff(y15, y16)
$S_d = y15$	
$y1 = FCW - y16$	F_PheAcc(y1, y2)
$y2 = y2 + y1$	
$y3 = IIR(a1, y2, y3)$	F_IIR1(y2, y3)
$y4 = IIR(a2, y3, y4)$	F_IIR2(y3, y4)
$y5 = IIR(a3, y4, y5)$	F_IIR3(y4, y5)
$y6 = IIR(a4, y5, y6)$	F_IIR4(y5, y6)
$y7 = SHR(y6, \alpha)$	F_SHR(y6, y7)
$y11 = IIR(gs, y7, y11)$	F_IIRgs(y7, y11)
$y9 = y6 - PhE$	F_IntAcc(y6, y10)
$y10 = y10 + y9$	
$y12 = SHR(y10, \rho)$	F_SHRAdd(y10, y11, y13)
$y13 = y11 + y12$	
$y14 = y13 * gain$	F_Gain()
jump L1	NA

Detailed descriptions for most of the instructions in Table 1 above are provided below. For each instruction, a corresponding figure is provided illustrating the data paths and

21

computation elements in the RCU used in executing the instruction. Elements in each of the figures described below operate as described above in connection with FIG. 13. Further, the bold lines or arrows in each figure highlight the data path for that particular instruction.

F_Diff Instruction

F_Diff(y15, y16):
 $y16 = y15 - S_d$
 $S_d = y15$

A block diagram illustrating the RCU unit 350 for implementing the F_Diff() instruction is shown in FIG. 14. The F_Diff instruction implements a differentiation operation. The bold arrow lines highlight the data path in the RCU. Input y15 is received and the first addition/subtraction unit is bypassed (e.g., the second or negating input to the first adder is set to zero). The shifting unit is also bypassed and a subtraction (via the second addition/subtraction unit) is performed with local register S_d. The result y16 is output via Rd_st. The instruction also updates S_d with input y15 via register latch S_d.

F_PheAcc Instruction

F_PheAcc(y16, y2):
 $y1 = FCW - y16$
 $y2 = y2 + y1$

A block diagram illustrating the RCU unit 352 for implementing the F_PheAcc() instruction is shown in FIG. 15. This instruction implements a subtraction plus an accumulation operation. It takes the input y16, performs the accumulation operation using the first addition/subtraction unit, bypasses any shifting (i.e. the shifter is configured for pass through operation). A subtraction operation with local register FCW is then performed. The accumulation is done on the Rd register latch.

F_IIR Instruction

F_IIR():
 $Rd = Rb \gg (1-a) + Ra \gg a$

A block diagram illustrating the RCU unit 354 for implementing the F_IIR() instruction is shown in FIG. 16. This instruction implements an IIR filtering operation. The shift amount 'a' is pre-set locally within the RCU. Thus, the RCU is optimized to have the capability of performing an IIR filter operation in a single instruction cycle. This permits an efficient computation of the atomic operations needed to implement the ADPLL within a single reference clock period.

F_SHR Instruction

F_SHR(Ra, Rd):
 $Rd = Ra \gg a$

A block diagram illustrating the RCU unit 356 for implementing the F_SHR() instruction is shown in FIG. 17. This instruction implements a shifting operation. The shift amount 'a' is pre-set locally within the RCU. The two addition/subtraction units are bypassed for this instruction.

F_IntAcc Instruction

F_IntAcc(y6, y10):
 $y9 = y6 - PhE$
 $y10 = y10 + y9$

A block diagram illustrating the RCU unit 358 for implementing the F_IntAcc() instruction is shown in FIG. 18. This instruction implements a subtraction and an accumulation

22

operation. It takes the input y16, performs an accumulation in the first addition/subtraction unit and then bypasses the shifter. A subtraction is then performed using local register PhE.

F_SHRAdd Instruction

F_SHRAdd(y10, y11, y13):
 $y12 = SHR(y10, rho)$
 $y13 = y11 + y12$

A block diagram illustrating the RCU unit 360 for implementing the F_SHRAdd() instruction is shown in FIG. 19. This instruction implements a shifting operation followed by an addition operation. The shift amount RHO is set locally in the RCU. Input y10 is received, the first addition/subtraction unit is bypassed and then a shifting operation is performed. The shifter output then undergoes an addition operation via the second addition/subtraction unit.

F_Gain Instruction

The F_Gain() instruction performs a multiplication by a gain value. In the case where the gain value is a power of two, the shift operation is used to perform this instruction. For non power of two gain values, a multiplier in the RCU is used (not shown).

In addition to the instructions described in detail hereinabove, the RCU comprises other customized instructions that are needed for general purpose applications, such as for setting RCU local register values, etc.

RF Spur Reduction

The effects of the software based ADPLL on RF spur reduction will now be described. As described supra, in prior art single-chip radios, the phase domain calculation portion of the ADPLL signal processing is traditionally implemented using dedicated random logic gates. In such an implementation, all computations are initiated on a rising edge of the ADPLL system clock and latched on the subsequent rising edge of the clock. This is shown in FIG. 20 which illustrates a timing diagram of several ADPLL processing clock options and the current spikes resulting therefrom. Trace 370 is the FREF system reference clock while trace 374 represents the processor clock.

Since the majority of circuit switching activity in the PLL (and also other close-in circuitry) is centered on a rising edge of the FREF system clock, most of the digital current is being switched at this edge as well, as indicated by trace 372. These digital current surges find their way into on-chip DCO and PA circuit nodes in the transmitter and LNA and mixer nodes in the receiver via various parasitic coupling mechanisms. The current rush energy due to digital processing at the system clock rate gets upconverted into the RF spectrum by the DCO, resulting in unacceptable RF spurs that are close in frequency to the carrier. The Federal Communications Commission (FCC) rules and numerous wireless communication standards place very low limits on the energy outside of the information carrying frequency range that is allowed to be radiated from wireless terminals.

These low frequency RF spurs are normally very difficult to filter out in a wireless terminal before they are radiated by the antenna because of their proximity in the spectrum to the carrier as shown in FIG. 21. The FREF clock spurs 384 are shown around the carrier 382. The filtering envelope 380 will not sufficiently attenuate the spurs to meet the various wireless standards. This is because a high order filtering is required to block the undesired energy, i.e. a steep filter envelope.

23

lope is required. The typical combined filtering effects of the (power amplifier), SAW filter and antenna filter out only a portion of the energy of these undesired RF spurs, thus making FCC rules and wireless standard compliance extremely difficult to meet.

In contrast, the mechanism of the present invention is operative to perform the atomic operations serially at the much faster processor rate. At each processor cycle, only a single instruction is executed resulting in reduced current transients being generated, as indicated in trace 376. This results in significantly reduced RF spur generation. In the present invention, the effects these generated spurs have on RF performance are two-fold: (1) the spurs are shifted to higher frequencies where they are easier to filter out by the oscillator's LC tank, DPA matching network, PA matching network, bandpass filtering in the RF front-end module, as well as overall parasitic RC (resistor-capacitor) network; and (2) the energy of each spur is reduced.

In accordance with the invention, the software based ADPLL significantly reduces the generation of RF spurs in sensitive frequency bands of the DRP by changing the frequency plan of the switching logic gates. The ASIP/RCU performs the bulk of computations at the processor clock rate which is much higher in frequency than that of the system reference clock (FREF). This result in the RF spurs being shifted outside the sensitive regions close to the carrier frequency as shown in FIG. 22. The processing clock spurs 392 are now far away from the carrier 392 and the filtering envelope 390 is able to remove these spurs with ease.

It is noted that the total amount of energy consumed in performing computations in the conventional ADPLL and the software based ADPLL of the invention is substantially the same. The frequency content of this energy, however, is significantly different. In case of the conventional ADPLL, the computation energy is concentrated at the reference clock edges and therefore has strong low frequency harmonic.

In case of the software based ADPLL, the computation energy is spread out between the reference clock edges (which may or may not coincide with the processor clock edges). The spreading out of the computation energy over the reference clock period serves to create a much higher frequency harmonic. When these harmonics mix with the carrier, they are offset in frequency around the carrier. The filtering requirements to remove the undesired RF spurs close to the carrier are excessive, as in the case of the conventional ADPLL. The further away from the carrier the RF spurs are located, the more relaxed the filtering requirements become. Changing the frequency content of the computation energy to improve the RF spectrum characteristics is referred to as frequency planning. Such frequency planning also helps to relax circuit design constraints, such as power routing resistance and decoupling capacitance values, which contribute to the magnitude of the energy surges at the clock edges.

Further, the frequency planning can be modified by varying the amount of oversampling (i.e. the ratio of processor clock frequency to the reference frequency FREF). The frequency planning can be lowered as long as the minimum required processing throughput in the processor is maintained. Otherwise, the required atomic operations of the ADPLL may not be completed within a reference clock cycle.

It is appreciated that the application of the software based mechanism described herein is not limited to an ADPLL. The mechanism can be applied to any computing or processing task that can benefit from reduced spur generation. A flow diagram illustrating the RF spur reduction method of the present invention is shown in FIG. 23. In the general case, the task is first partitioned into a plurality of atomic operations

24

(step 740). The computation of the atomic operations required to complete the task are spread out over an entire reference clock period (step 742). The computation of each atomic operation occurs at the higher processor clock rate, rather than the slower reference clock rate (step 744).

ALTERNATIVE EMBODIMENTS

Several embodiments of the processor based mechanism will now be presented. Each is an alternative implementation of the software based ADPLL. A block diagram illustrating an example ASIP based implementation of the software ADPLL of the present invention is shown in FIG. 24. The software based ADPLL, generally referenced 550, comprises a phase domain calculator 552 that performs the data stream processing, DCO 556, programmable fractional-N clock divider 554, integer feedback block 578 and fractional feedback block 579. The phase domain calculator 552 comprises an ALU (RCU) 572, instruction and data memory 558, register file 560, sequencer 562, latches 566, 576, multiplexers 564, 568, 570, 574.

A block diagram illustrating an example DSP based implementation of the software ADPLL of the present invention is shown in FIG. 25. The software based ADPLL, generally referenced 580, comprises a phase domain calculator 582, DCO 586, programmable fractional-N clock divider 584, integer feedback block 588 and fractional feedback block 590. The phase domain calculator 582 comprises memory block (RAM/ROM) 594, a DSP core 595, external bus interface (I/F) 592, UART 596 and timer 598.

A block diagram illustrating an example hardware state machine based implementation of the software ADPLL of the present invention is shown in FIG. 26. The software based ADPLL, generally referenced 600, comprises a phase domain calculator 602, DCO 606, programmable fractional-N clock divider 604, integer feedback block 608 and fractional feedback block 610. The phase domain calculator 602 is similar in construction and operation to the phase calculation 222 (FIG. 10). In particular, the phase domain calculator 602 comprises flip-flops 612, 616, 624, 630, 634, 656, 650, 638, subtractors 614, 654, adder 628, multipliers 622, 632, 652, accumulators 618, 658, differentiator 636, IIR blocks 620, 626, Lambda register 640, PHE offset register 642, Gain1 register 646, Gain2 register 648 and Gain3 register 644.

A block diagram illustrating a first example implementation of the state machine used to implement the software ADPLL of the present invention is shown in FIG. 27. The circuit, generally referenced 660, comprises a clock generator circuit 662 and a plurality of finite state machines (FSMs) 664, 666, 668, 670. The clock generator is operative to generate the clock signals for the functional units, i.e. FSM1, FSM2, FSM3, FSM4. Although only four units are shown for the sake of simplicity, it is appreciated that any number of state machines may be implemented.

The functional units are executed in a sequential manner to provide the desired overall functionality of a subsystem such as a PLL or other circuit/task. An input of a functional unit is connected to the output of the previous unit with the last unit feeding the first unit. Only a simple circular arrangement is shown although the technique is applicable to more complicated arrangements as well, such as forking and merging.

The higher frequency clock CKVD 672 is generated by performing frequency or edge division of the variable clock CKV. The CKVD clock is used to generate internal clocks CK1, CK2, CK3, CK4, each one providing significant edges to its respective functional units FSM1, FSM2, FSM3, FSM4. The frequency of the CKVD clock is chosen so as to guaran-

tee that all the elementary functional operations performed by functional units FSM1, FSM2, FSM3, FSM4 are executed within a single reference clock FREF cycle 674. This is accomplished by requiring the CKVD clock to be (in this example) at least four times faster (i.e. the number of functional units) than the FREF clock. In the example presented in FIG. 27, the frequency ratio is approximately 4.6.

A block diagram illustrating a second example implementation of the state machine used to implement the software ADPLL of the present invention is shown in FIG. 28. The circuit, generally referenced 680, comprises clock and phase generator 682 and FSM 684. This figure presents the embodiment wherein the functionality of the various functional units (i.e. FSM1, FSM2, FSM3, FSM4) are merged into a single reconfigurable finite state machine (FSM). The FSM unit reprograms itself such that its functionality becomes one of the four units FSM1, FSM2, FSM3, FSM4 of FIG. 27. The determination of which unit the FSM 684 becomes is controlled by the "Phase" signal. The input and output connections between the functional units in FIG. 27 are now stored as state variables internal to the FSM, which now operates at a higher clock rate 688. The single clock line CK comprises gaps due to the non-integer frequency ratio of CKVD to FREF 689.

RCU Implementation Using Parallelization and Pipelining

As an aid in understanding the principles of the present invention, as applied exemplarily to the RCU implementation, its environment as well as its control mechanism, a brief overview of parallel processing concepts is provided hereinbelow. Numerous CPU technologies exist that address instruction level parallelism and data level parallelism. Examples of such CPU technologies include superscalar, very long instruction word (VLIW), explicitly parallel instruction computing (EPIC), single instruction multiple data (SIMD), etc.

A superscalar architecture executes more than one instruction in a single pipeline stage by pre-fetching multiple instructions and simultaneously dispatching them to redundant functional units on the processor, as shown in FIG. 29, which illustrates superscalar pipeline instruction parallelization. In a superscalar CPU, the dispatcher reads instructions from memory and determines which instructions can be run in parallel, dispatching them to redundant functional units contained inside a single CPU. A superscalar processor can, therefore, be considered to have multiple parallel pipelines, each of which processes instructions simultaneously from a single instruction thread.

The performance improvement obtainable from using superscalar techniques is limited, however, by two key areas: (1) the degree of parallelism intrinsic within the instruction stream wherein the amount of instruction level parallelism may be limited due to data dependencies; and (2) the complexity and time cost of the dispatcher and associated dependency checking logic.

VLIW and EPIC techniques are similar to the superscalar technique discussed supra in that all the above parallelism techniques have multiple pipelines running in parallel on different functional units. Thus, their efficiencies are all limited by the degree of intrinsic parallelism in the instruction stream.

Single instruction multiple data (SIMD) is a technique employed to achieve data level parallelism. A block diagram illustrating the single instruction multiple data (SIMD) technique is shown in FIG. 30. The system, generally referenced

870, comprises an instruction pool 874, data pool 872 and a plurality of processing units (PUs) 876. A single instruction is applied to multiple data to perform the same operation.

The parallel implementation of the ASIP of the present invention utilizes data stream based parallelization as opposed to instruction level parallelization. Due to data dependencies, however, these parallel CPU architecture techniques are not efficient. The invention provides a novel architecture comprising mixed parallel and pipeline techniques that address the disadvantages of the prior art. The architecture of the invention is optimized to perform stream data processing in a more efficient manner.

An example stream data processing model will now be described. A block diagram illustrating an example stream data based processing model is shown in FIG. 31. The model, generally referenced 880, comprises four processing function blocks 882 labeled F1, F2, F3, F4. In operation, the input samples (s(1), s(2), s(3), . . .) are fed into the processing function blocks at a fixed sampling rate, e.g., 38.4 MSPs. Each sample is processed through sequential function blocks (F1, F2, F3, F4) with the final result (o(1), o(2), o(3), . . .) being output from final functional block (F4).

A functional processing block may be of one of two types: (1) recursive or (2) non-recursive. In the case of a recursive functional processing block type, the current output is not only related to the current input, e.g., x(n), but is also related to historical output, e.g., y(n-1). An example of a recursive functional processing block is an Infinite Impulse Response (IIR) filter:

$$y(n)=a \cdot x(n)+b \cdot y(n-1) \quad (8)$$

Note, that the IIR filter coefficient naming (i.e. 'a' and 'b') in this section is different than in the previous examples, Equation 7 in particular. An example of a non-recursive functional processing block, for which the output is not related to historical output of the block, is an addition/subtraction block:

$$y(n)=a \cdot x(n)+b \quad (9)$$

Note that there are several ways to implement the data processing functions as described in connection with FIG. 31. A pure hardwired solution comprises dedicated hardware for every function block (i.e. F1, F2, F3, F4). A processor based solution comprises a reconfigurable computation unit (RCU) controlled by instructions adapted to handle different functions (F1, F2, F3, F4).

A block diagram illustrating example processor based data processing for implementing the model of FIG. 31 is shown in FIG. 32. The data processor, generally referenced 890, comprises an RCU 896 in communication with a control and configuration block 892 and a register file 894. The RCU comprises a data interface to the register file for temporary result storage via read ports 1 and 2 and a write port.

As an aid to illustrating the operation of the data stream processor 890, an example is provided to show how the processor 890 handles the data processing performed in FIG. 31. For purposes of this example, assume blocks 882 comprises recursive functional processing blocks F1, F2, F3, F4 that implement four IIR filters. The four filters can be expressed mathematically as follows:

$$F1: y(n)=a1 \cdot x(n)+b1 \cdot y(n-1)$$

$$F2: y(n)=a2 \cdot x(n)+b2 \cdot y(n-1)$$

$$F3: y(n)=a3 \cdot x(n)+b3 \cdot y(n-1)$$

$$F4: y(n)=a4 \cdot x(n)+b4 \cdot y(n-1)$$

(10)

where $y(n)$ is the current local output of a given filter, $x(n)$ is the current local input to the filter and $y(n-1)$ is the previous local output of the filter. Note that the processor comprises eight registers in register file 894, labeled R1 through R8. The pseudo code for the data processing in FIG. 31 is presented below in Listing 2 (assuming a1 through a4 and b1 through b4 is stored within the RCU).

Listing 2: Pseudo-code for the data processing of FIG. 31

```

R1=a1*s(n) + b1*R1; //R1 stores the history value for F1;
R2=a2*R1 + b2*R2; //R2 stores the history value for F2;
R3=a3*R2 + b3*R3; //R3 stores the history value for F3;
R4=a4*R3 + b4*R4; //R4 stores the history value for F4; result is
                    output o(n).
R1=a1*s(n+1) + b1*R1;
R2=a2*R1 + b2*R2;
R3=a3*R2 + b3*R3;
R4=a4*R3 + b4*R4; //result is output as o(n+1).
...

```

A Data Stream Processing Model

With reference to FIG. 32, one instruction controls the RCU to complete one function within one processor clock cycle. The processing power of a processor constructed using the RCU is limited by the speed of operation the RCU can sustain. Multiple RCUs operating in parallel, however, can speed up the process. The invention provides a processing technique using multiple RCUs that have the following capabilities. Note that it is assumed that RCU comprises two read ports and one write port to and from the Register File.

1. A mechanism for controlling multiple RCUs whereby the size of the instruction does not increase with the use of multiple RCUs.
2. Historical values for recursive operations are stored using only one write port to the Register File.
3. Communications between multiple RCUs is handled locally.

In accordance with the invention, a multi-RCU processing model is provided comprising a mixed parallel/pipelined architecture. A block diagram illustrating an example two-stage parallel/pipelined RCU architecture is shown in FIG. 33. The architecture, generally referenced 690, comprises a two-stage parallel/pipelined processor 691 in communication with a control/configuration block 708 and register file 710/718. Note that register file 710/718 are the same register file drawn twice for clarity sake. The processor 691 comprises multiplexers 692, 696, 712, 716, registers 694, 704, 714, 706, delay 698 and RCU1 700, RCU2 702.

The local registers L1 704, L2 706 function to store the temporal result for each RCU 700, 702, respectively. Note that in operation, only the first RCU (i.e. RCU1) receives input from the register file 710. Similarly, only the last RCU (i.e. RCU2) writes results back to the register file 710/718. The input sample pairs (i.e. $s(1)$, $s(2)$, ...) are fed to the RCUs in parallel, thus comprising the parallel aspect of the architecture. Each input sample is processed in only one RCU (i.e. either RCU1 or RCU2). This constitutes the parallel aspect of the architecture. RCU2 is operative to always perform the function processing that RCU1 performed in the previous cycle. This constitutes the pipeline aspect of the architecture. RCU2 is operative to use the historical value from the output of RCU1. This also constitutes the pipeline aspect of the architecture.

Further, RCU2 is operative to always complete the processing of a sample one cycle later than RCU1. Note that due to the use of local registers associated with each RCU output, the architecture needs only a single register file read port. Each RCU unit comprises (1) a feedback path 699, 701 for RCU1, RCU2, respectively, that is used to perform sequential functional processing of data samples and (2) a forward path 703 for transferring processing results between logically adjacent RCUs in pipeline fashion. In the feedback path, the data sample is processed through various functions in sequential fashion, wherein the results output of one function are used as input to the next function. In the forward path, processing results are transferred from one RCU to another in pipeline fashion. Normally, the RCUs are logically adjacent to each other (i.e. may or may not be physically adjacent) whereby the processing result output from one RCU are input to a downstream (i.e. a logically forward adjacent RCU) RCU for further processing.

As an illustrative example, the processing performed by the two-stage parallel/pipeline RCU architecture 690 of FIG. 33 is presented below in Table 2. The example is operative to implement four IIR filters. Note that Rd1, Rd2 are outputs of RCU1, RCU2, respectively. Note that with reference to Equation 10, in RCU1, Rd1 is analogous to $y(n)$, $s(1)$, $s(3)$, etc. is analogous to $x(n)$ and R1, R2, etc. is analogous to $y(n-1)$; in RCU2, Rd2 is analogous to $y(n)$ and Rd1 is analogous to $y(n-1)$.

TABLE 2

Processing in the two-stage parallel/pipeline RCU architecture		
RCU1	RCU2	Comment
Rd1 = a1 * s (1) + b1 * R1	NA	RCU2 waits for pipeline to start
Rd1 = a2 * Rd1 + b2 * R2	Rd2 = a1 * s (2) + b1 * Rd1	Rd2 is also stored in R1
Rd1 = a3 * Rd1 + b3 * R3	Rd2 = a2 * Rd2 + b2 * Rd1	Rd2 is also stored in R2
Rd1 = a4 * Rd1 + b4 * R4	Rd2 = a3 * Rd2 + b3 * Rd1	Rd2 is also stored in R3
Rd1 = a1 * s (3) + b1 * R1	Rd2 = a4 * Rd2 + b4 * Rd1	Rd2 is also stored in R4
Rd1 = a2 * Rd1 + b2 * R2	Rd2 = a1 * s (4) + b1 * Rd1	Rd2 is also stored in R1
Rd1 = a3 * Rd1 + b3 * R3	Rd2 = a2 * Rd2 + b2 * Rd1	Rd2 is also stored in R2
Rd1 = a4 * Rd1 + b4 * R4	Rd2 = a3 * Rd2 + b3 * Rd1	Rd2 is also stored in R3
Rd1 = a1 * s (5) + b1 * R1	Rd2 = a4 * Rd2 + b4 * Rd1	Rd2 is also stored in R4
Rd1 = a2 * Rd1 + b2 * R2	Rd2 = a1 * s (6) + b1 * Rd1	Rd2 is also stored in R1
...	...	...

A generalized description of the processing in the parallel architecture of FIG. 33 is provided below in connection with Table 3. With reference to FIGS. 31 and 33, the sequence of functions performed is $F1 \rightarrow F2 \rightarrow F3 \rightarrow F4 \rightarrow F1 \rightarrow F2 \dots$. The sample sequence comprises $s(1) \rightarrow s(2) \rightarrow s(3) \rightarrow s(4) \dots$. The time sequence is $T1 \rightarrow T2 \rightarrow T3 \rightarrow T4 \dots$. The processor comprises a two-stage parallel/pipeline comprising RCU1 and RCU2. The term ' $s(i) \rightarrow Fx$ ' means conduct function Fx for sample $s(i)$. The term ' $s(i) \rightarrow Fx(st)$ ' means conduct function Fx for sample $s(i)$ and store the result in the register file.

TABLE 3

Parallel/Pipeline Calculations		
Time	RCU1	RCU2
T1	S (1)→F1	NA
T2	S (1)→F2	s (2)→F1 (st)
T3	S (1)→F3	s (2)→F2 (st)
T4	S (1)→F4	s (2)→F3 (st)
T5	S (3)→F1	s (2)→F4 (st)
T6	S (3)→F2	s (4)→F1 (st)
T7	S (3)→F3	s (4)→F2 (st)
T8	S (3)→F4	s (4)→F3 (st)
T9	S (5)→F1	s (4)→F4 (st)
T10	S (5)→F2	s (6)→F1 (st)

To achieve a pipelined process, the processor only configures the first RCU unit (i.e. RCU1), the second RCU unit (i.e. RCU2) performs the same processing function with one cycle delay. For example, at time T3, the processor configures RCU1 to perform function F3 for sample s (1), while RCU2

In2: defines the second input for RCU1 (from the register file for the historical value of a function unit) and for RCU2 (always comes from the output of RCU1);

5 Out: defines the resulting output of RCU2, which is stored in the register file;

Note that Fx_d is the historical value of functional unit Fx that is stored in the register Ri.

10

Let us assume the RCU is performing four IIR functions repeatedly. Consider time T3, for example, RCU1 performs function F3 for S1, it uses its own latched output (Rd1) as input along with the historical value of function F3 (F3_d) from the register file. The result is not stored in the register file but rather is locally latched. RCU2 performs function F2 for s (2), it uses the output from RCU1 (Rd1) along with the local latched result (Rd2) from the previous cycle. The result is stored in the register file.

15

20

TABLE 4

Pipelined Processing Within the RCUs					
Time	RCU1	RCU2	RCU1(Ra, Rb, Rd)	RCU2(Ra, Rb, Rd)	RCU(Ra, Rb, Rd)
T1	s(1)→F1		F1(s(1), F1_d, Rd1)		F1(s(1), R1, NA)
T2	s(1)→F2	s(2)→F1(st)	F2(Rd1, F2_d, Rd1)	F1(s(2), Rd1, R1)	F2(s(2), R2, R1)
T3	s(1)→F3	s(2)→F2(st)	F3(Rd1, F3_d, Rd1)	F2(Rd2, Rd1, Rd2)	F3(NA, R3, R2)
T4	s(1)→F4	s(2)→F3(st)	F4(Rd1, F4_d, Rd1)	F3(Rd2, Rd1, Rd2)	F4(NA, R4, R3)
T5	s(3)→F1	s(2)→F4(st)	F1(s(3), F1_d, Rd1)	F4(Rd2, Rd1, Rd2)	F1(s(3), R1, R4)
T6	s(3)→F2	s(4)→F1(st)	F2(Rd1, F2_d, Rd1)	F1(s(4), Rd1, Rd2)	F2(s(4), R2, R1)
T7	s(3)→F3	s(4)→F2(st)	F3(Rd1, F3_d, Rd1)	F2(Rd2, Rd1, Rd2)	F3(NA, R3, R2)
T8	s(3)→F4	s(4)→F3(st)	F4(Rd1, F4_d, Rd1)	F3(Rd2, Rd1, Rd2)	F4(NA, R4, R3)
T9	s(5)→F1	s(4)→F4(st)	F1(s(5), F1_d, Rd1)	F4(Rd2, Rd1, Rd2)	F1(s(5), R1, R4)

is configured to perform function F2 for sample s (2). The result of function F2 (F2_d) is stored in the register file. Only the last RCU (i.e. RCU2 in this example) stores its result in the register file.

Table 4 presented below provides additional detail on the internal operation of RCU1 and RCU2 and an example processor command for the RCUs. First, we define the notation for the functions that RCU1 and RCU2 perform:

Fx(Ra, Rb, Rd):

Fx (): defines the type of operation performed by RCU1 and RCU2;

Ra: defines the first input for RCU1 and RCU2, which may come from the sample input or RCU1, RCU2 output, i.e. Rd1, Rd2, respectively;

Rb: defines the second input for RCU1 (from the register file) and for RCU2 (always comes from the output of RCU1);

Out: defines the resulting output;

For the RCU as a whole (i.e. includes both RCU1, RCU2), the meaning is slightly different than that above:

Fx(In1, In2, Out):

Fx (): defines the type of operation RCU1 and RCU2 will perform in the next cycle.

In1: defines the first input for RCU1, RCU2, which may come from the sample input or RCU1, RCU2 output, i.e. Rd1, Rd2, respectively;

40

45

50

55

60

65

Characteristics for the two-stage parallel/pipelined process are presented hereinbelow. Two data samples are fed into the RCU (the term RCU denotes the combination of all RCUs, i.e. RCU1 and RCU2) which are processed in a pipelined manner for all functions (i.e. F1→F2→F3...). The first RCU unit (RCU1) processes the first sample and the second RCU unit (RCU2) processes the second sample. The processor configures only RCU1, whereby RCU2 performs the processing function that RCU1 performed in the previous cycle. For processing functions comprising a memory unit, only the result from the last RCU unit that processes the last sample (i.e. RCU2 in this example embodiment) needs to be stored in the register file. The value stored in the register file is used in the subsequent cycle by the first RCU (i.e. RCU1). After receiving a new sample data, each RCU unit takes its own output as input for the subsequent function. The RCU2 unit takes the output from the RCU1 unit as a historical value.

Several example instructions for the RCU are provided below:

F1(In, R_1, R_last):

This instruction is operative to input data Si (where "i" is an odd number) to RCU1 to perform function F1. The historical value for function F1 is retrieved from register R_1. RCU2 performs the function last for sample "i-1" and stores the result in register R_last.

F2(In, R_2, R_1):

RCU1 performs function F2 for sample "i-1", wherein the historical value for function F2 is taken from R_2. This

31

instruction is operative to input data S_i (wherein “i” is even number) to RCU2 to perform function F1. The historical value for function F1 is taken from RCU1. The result is stored in register $R_{_1}$.

$F_i(NA, R_{_j}, R_{_k})$:

For this instruction, RCU1 performs function F1 with input coming from the local output of RCU1 and the historical value for function F1 is taken from register $R_{_j}$. RCU2 performs function F_{i-1} with historical value taken from RCU1. The result is stored in register $R_{_k}$.

It is appreciated that from the above analysis, one skilled in the art can build a fixed relation between function index and the input and output register locations. In this manner, the instruction coding space required can be reduced significantly.

Data Path Forking

With reference to FIGS. 6 and 10, there are two data paths operating in parallel in the ADPLL loop filter path: one is a proportional filter and the second is an integral filter. Thus, this is not a single path data processing mode as is shown in FIG. 31, but rather a data stream comprising forking and merging. A block diagram illustrating an example data path incorporating forking and merging is shown in FIG. 34. The example data path, generally referenced 720, comprises a plurality of function blocks 722 (i.e. F1, F2, F3, F4, F5, F6, F7). Function block F6 is a processing function that does not have a memory element (i.e. no historical value is needed for its calculation).

To handle the data path forking and merging, an additional latch unit is added for each RCU to latch the result for the other data path as shown in FIG. 35 which illustrates a block diagram of an example parallel/pipelined architecture incorporating forking handling capability. The additional loop back data path from the output of each RCU unit (751, 761) to its respective input serves to handle fork merging. This architecture allows the RCU internal pipeline to keep running without being affected by forking and merging in the data path. Note that loopback path 751 also serves as a forward path for transferring processing results from one RCU to a logically adjacent RCU.

The architecture, generally referenced 730, comprises a two-stage parallel/pipelined processor 731 in communication with a control/configuration block (not shown) and register file 742, 766. The processor 731 comprises multiplexers 732, 738, 752, 754, 740, 760, 762, 764, registers 734, 736, L11 748, L12 750, L21 756, L22 758 and RCU1 744, RCU2 746.

32

to register L12. At time T6, RCU1 needs to use the data contents of both registers L11 and L12 for the fork merge.

The control aspect of the RCU unit will be described in more detail. A block diagram illustrating control of an example RCU is shown in FIG. 36. The circuit, generally referenced 770, comprises an RCU control block 771, instruction fetching block 786 and instruction decoding block 788. The RCU block 771 comprises multiplexers 774, 782, 784, registers 772, L11 778, L12 780 and RCU unit 776. The circuit of FIG. 36 highlights the control signals within the ASIP for controlling the RCU. All the control signals for the RCU are provided by the instruction decoding block 788.

A block diagram illustrating control of an example three-stage pipelined/parallel RCU is shown in FIG. 37. The circuit, generally referenced 790, comprises three-stage pipeline RCU and control block 791, instruction fetching block 792 and instruction decoding block 794. The three-stage pipeline RCU and control block 791 comprises control signal delay latches 796, 798 and RCU units 800, 802, 804. An advantage of the parallel RCU structure of the present invention is that the control of the sub RCUs, i.e. RCU1, RCU2, RCU3, is the same but with a pipelined delay. The pipelined delay is handled through the delay latching of Ctl_L1 (delay latch 796) and Ctl_L2 (delay latch 798).

Three-Stage Pipeline

A block diagram illustrating an example three-stage parallel/pipelined RCU architecture is shown in FIG. 38. The circuit, generally referenced 810, comprises three-stage pipelined/parallel processor 811 and register files 812, 814. The three-stage pipelined/parallel processor comprises multiplexers 816, 820, 828, 830, 836, 844, 846, 850, 858, 860, 832, registers 818, L11 824, L12 826, 834, L21 840, L22 842, 848, L31 854, L32 856 and RCU units RCU1 822, RCU2 838, RCU3 852.

A diagram illustrating the data processing within the three-stage parallel/pipeline processor of FIG. 38 is shown in FIG. 39. The diagram shows how data samples are processed within the three-stage parallel/pipelined processor 811 (FIG. 38). The zigzagged cross line highlights the data stream processing for function F1. Samples S1, S2, S3, . . . are successively processed. The processed result generated by RCU3 is stored in the register file for use by RCU1 two cycles later. It is appreciated that this example processor can be modified to comprise any number of RCU units arranged in the same pipeline/parallel arrangement as shown in FIG. 38.

It is intended that the appended claims cover all such features and advantages of the invention that fall within the spirit

TABLE 5

Forking Processing Inside the Two-Stage RCU				
Time	RCU1	RCU2	RCU1(Ra, Rb, Rd)	RCU2(Ra, Rb, Rd) RCU(Ra, Rb, Rd)
T1	s(1)→F1		F1(S1, F1_d, L11)	F1(S1, R_1, NA)
T2	s(1)→F2	S2→F1(st)	F2(L11, F2_d, L11)	F2(S2, R_2, R_1)
T3	s(1)→F3	S2→F2(st)	F3(L11, F3_d, L12)	F3(NA, R_3, R_2)
T4	s(1)→F5	S2→F3(st)	F5(L11, F5_d, L11)	F5(NA, R_5, R_3)
T5	s(1)→F4	S2→F5(st)	F4(L12, F4_d, L12)	F4(NA, R_4, R_5)
T6	s(1)→F6	S2→F4(st)	F6(L12, L11, L11)	F6(NA, NA, R_4)
T7	s(1)→F7	S2→F6(st)	F7(L11, F7_d, L11)	F7(NA, R_7, NA)
T8	s(3)→F1	S2→F7(st)	F1(S3, F1_d, L11)	F1(S3, R_1, R_7)
T9	s(3)→F2	S4→F1(st)	F2(L11, F1_d, L11)	F2(S4, R_2, R_1)

For RCU1, forking begins at time T3 and merges at time T6. At time T3, RCU1 takes data from register L11 but writes

and scope of the present invention. As numerous modifications and changes will readily occur to those skilled in the art,

33

it is intended that the invention not be limited to the limited number of embodiments described herein. Accordingly, it will be appreciated that all suitable variations, modifications and equivalents may be resorted to, falling within the spirit and scope of the present invention.

What is claimed is:

1. A processor, comprising:

A. a plurality of reconfigurable calculation units coupled both in parallel wherein input data samples are fed in parallel thereto and in a pipelined configuration wherein results from one reconfigurable calculation unit are input to a reconfigurable calculation unit logically adjacent thereto;

B. each reconfigurable calculation unit including:

i. a feedback path to perform sequential functional processing of input data samples in which an input data sample is processed through various functions in sequential fashion and the results output of one function are used as input to the next function; and

ii. a forward path to transfer processing results between logically adjacent reconfigurable calculation units;

C. a hardware register file to store historical result values generated by a last reconfigurable calculation unit in the pipeline; and

D. an instruction decoder to generate the same control and configuration signals for all reconfigurable calculation units but with a pipelined delay therebetween so that an upstream unit receives a control and configuration signal before a downstream unit.

2. The processor according to claim 1, wherein each reconfigurable calculation unit comprises local registers for storing temporary computation results.

3. The processor according to claim 1, wherein only an upstream reconfigurable calculation unit reads previously written result data from the register file.

4. The processor according to claim 1, wherein an input data sample is processed in a single reconfigurable calculation unit.

5. The processor according to claim 1, wherein a downstream processing reconfigurable calculation unit processes a same processing function processed by an upstream reconfigurable calculation unit logically adjacent thereto in a previous cycle.

6. The processor according to claim 1, wherein a downstream processing reconfigurable calculation unit uses a historical value output by an upstream reconfigurable calculation unit logically adjacent thereto in a previous cycle.

7. The processor according to claim 1, wherein each reconfigurable calculation unit utilizes its own output as input for a subsequent processing function.

8. A data stream processor, comprising:

A. a plurality of reconfigurable calculation units coupled both in parallel and in a pipelined configuration wherein input data samples are fed in parallel to said the reconfigurable calculation units and processing results are transferred from one reconfigurable calculation unit to another in the pipeline configuration;

B. each reconfigurable calculation unit including:

i. a feedback path to perform sequential functional processing of the input data samples in which an input data sample is processed through various functions in sequential fashion and the results output of one function are used as input to the next function;

ii. a forward path to transfer processing results between logically adjacent reconfigurable calculation units; and

34

iii. a loopback path to support forking and merging operations;

C. a hardware register file to store historical result values generated by a last reconfigurable calculation unit in the pipeline; and

D. an instruction decoder to generate the same control and configuration signals for all reconfigurable calculation units but with a pipelined delay therebetween so that an upstream unit receives a control and configuration signal before a downstream unit.

9. The data stream processor according to claim 8, wherein the loopback path comprises one or more latches coupling a reconfigurable calculation unit output to its input thereby latching results for forking and merging operations.

10. The data stream processor according to claim 8, wherein each reconfigurable calculation unit comprises local registers for storing temporary computation results.

11. The data stream processor according to claim 8, wherein only an upstream reconfigurable calculation unit reads previously written result data from the register file.

12. The data stream processor according to claim 8, wherein an input data sample is processed in a single reconfigurable calculation unit.

13. The data stream processor according to claim 8, wherein a downstream processing reconfigurable calculation unit processes a same processing function processed by an upstream reconfigurable calculation unit logically adjacent thereto in a previous cycle.

14. The data stream processor according to claim 8, wherein a downstream processing reconfigurable calculation unit uses a historical value output by an upstream reconfigurable calculation unit logically adjacent thereto in a previous cycle.

15. The data stream processor according to claim 8, wherein each reconfigurable calculation unit utilizes its own output as input for a subsequent processing function after processing a new input sample data.

16. A data stream processing method, the method comprising the steps of:

A. coupling a plurality of reconfigurable calculation units both in parallel and in a pipelined configuration, including feeding input data samples in parallel to the reconfigurable calculation units and processing them in the pipelined configuration;

B. performing sequential functional processing of the input data samples within each reconfigurable calculation unit;

C. transferring processing results between logically adjacent reconfigurable calculation units; and

D. generating the same control and configuration signals in an instruction decoder for all reconfigurable calculation units but with a pipelined delay therebetween so that an upstream unit receives a control and configuration signal before a downstream unit.

17. The data stream processing method according to claim 16, further including the step of performing forking and merging operations using a loopback path in each reconfigurable calculation unit.

18. The data stream processing method according to claim 16, further including the step of reading historical result values from a register file by an upstream reconfigurable calculation unit and writing results values to the register file by a downstream reconfigurable calculation unit.

19. A processor based phase locked loop (PLL), comprising:

35

- A. an oscillator to generate a radio frequency (RF) signal having a frequency determined in accordance with a tuning command input thereto;
 - B. a processor to generate the tuning command including a plurality of reconfigurable calculation units coupled both in parallel and in a pipelined configuration wherein input data samples are fed alternately to the reconfigurable calculation units and processed in the pipelined configuration;
 - C. hardware program memory coupled to the reconfigurable calculation units for storing a plurality of instructions that when executed on the processor implement the phase locked loop; and
 - D. an instruction decoder to generate the same control and configuration signals for all reconfigurable calculation units but with a pipelined delay therebetween so that an upstream unit receives a control and configuration signal before a downstream unit.
20. The processor based phase locked loop according to claim 19, wherein each reconfigurable calculation unit comprises a loopback path to support forking and merging operations.
21. The processor based phase locked loop according to claim 19, wherein an input data sample is processed in a single reconfigurable calculation unit.
22. A radio having a transmitter coupled to an antenna, the transmitter including a software based phase locked loop (PLL), and the phase locked loop comprising:
- A. an oscillator to generate a radio frequency (RF) signal having a frequency determined in accordance with a tuning command input thereto;
 - B. a processor to generate the tuning command, the processor including:
 - C. a plurality of reconfigurable calculation units coupled both in parallel and in a pipelined configuration wherein

36

- input data samples are fed in parallel to the reconfigurable calculation units and processed in the pipelined configuration, each reconfigurable calculation unit performing atomic operations required to implement said the phase locked loop and including:
- D. a feedback path to perform sequential functional processing of input data samples in which a data sample is processed through various functions in sequential fashion and the results output of one function are used as input to the next function;
- E. a forward path to transfer processing results between logically adjacent reconfigurable calculation units; and
- F. a loopback path to support forking and merging operations;
- G. a hardware register file communicatively coupled to an upstream reconfigurable calculation unit and a downstream reconfigurable calculation unit;
- H. program memory coupled to the reconfigurable calculation units for storing a plurality of instructions that when executed on the processor implement the phase locked loop;
- I. the processor having an instruction set, wherein each instruction is to perform an atomic operation of the phase locked loop;
- J. a receiver coupled to said the antenna; and
- K. a baseband processor coupled to the transmitter and the receiver; and
- L. an instruction decoder to generate the same control and configuration signals for all reconfigurable calculation units but with a pipelined delay therebetween so that an upstream unit receives a control and configuration signal before a downstream unit.

* * * * *