



sf32bl

32-bit microprocessor

IMA (Implementation Architecture)
Reference Manual

Revision 0.9
18 March 2014

Author: Martin Raubuch



Revision History

Revision	Date	
0.9	18Mar2014	First version

Table of contents

1	Overview	4
1.1	Introduction	4
1.2	Feature Summary	4
1.3	Scope of this manual	4
2	I/O Overview	5
3	Interface Details	7
3.1	Instruction Fetch	7
3.2	Data Access	9
3.3	Interrupts	11
3.4	Debug	13
3.5	Reset	15
4	Instruction Execution Timing	17
4.1	Effective Execution Times	17
4.2	Stall Conditions	18
4.3	Loop Cache	19
5	Compatibility	20
5.1	Software	20
5.2	Hardware	20
5.3	Replacement Options	20

1 Overview

1.1 Introduction

The sf32 family of 32-bit microprocessors is targeted at general purpose control and computing applications. With fixed length 32-bit instruction coding the architectural focus is on high clock rates and small core implementations.

The sf32 family defines two ISAs (Instruction Set Architectures), a base (b) ISA for general purpose control & computing and a (d) DSP ISA extension for 32-bit DSP applications. This manual is the IMA (Implementation Architecture) reference of the sf32bl, the (l) light implementation of the sf32 (b) base ISA.

1.2 Feature Summary

The following list summarizes the sf32bl's main features

- Focus on performance and moderate core size
- 32-bit wide instruction and data interfaces
- Register file with 3/2 read/write ports
- Decoupled unit for instruction fetch and flow instruction execution
- Branch Speculation
- Loop cache, zero cycles loop branch from 2nd iteration
- Separate execution pipelines for computation and load/store instructions
- Single cycle effective execution of all computation instructions except multiplies
- Single cycle effective execution of all load/store instructions except load/store multiple
- Barrel shifter for single cycle effective execution of all shift instructions
- Average IPC (Instructions Per Cycle) of ~0.8 for typical code sequences
- 2 x 32-bit instruction pre-fetch buffer
- Fully synchronous design, all flip-flops are triggered with the rising edge of the clock input
- Clock rates up to 110MHz on low end FPGAs
- Clock rates >300MHz with deep sub-micron std-cell generic and low-power technologies

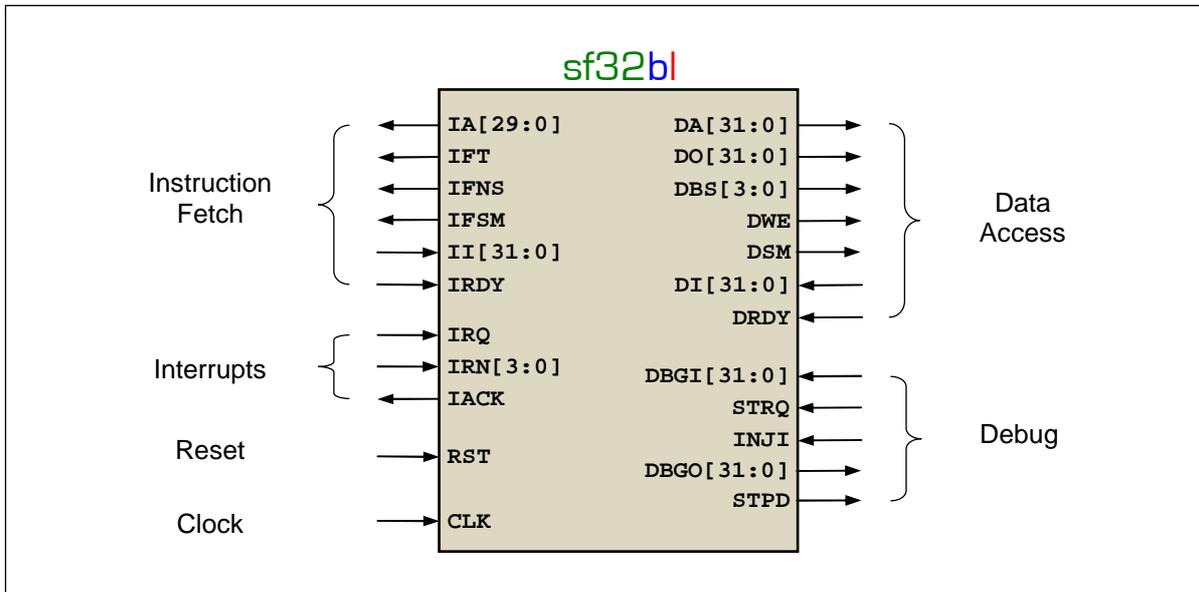
1.3 Scope of this manual

This sf32bl IMA reference manual contains the following detailed descriptions:

- **I/O Overview**, overview of interfaces and I/O signals
- **Interface Details**, detailed I/O signal descriptions and interface timing
- **Instruction Execution Timing**, effective execution time of instructions, data dependencies and stall conditions
- **Compatibility**, hardware and software compatibility, drop in replacement options

ISA specific details such as programming model and instruction set are not part of this IMA reference manual. This information can be found in the base (b) ISA (Instruction Set Architecture) reference manual.

2 I/O Overview



Signal	Direction	Width	Description
IA[29:0]	Output	30	Instruction Address
IFT	Output	1	Instruction Fetch
IFNS	Output	1	Instruction Fetch Non Sequential
IFSM	Output	1	Instruction Fetch System Mode
II[31:0]	Input	32	Instruction In
IRDY	Input	1	Instruction Ready
IRQ	Input	1	Interrupt Request
IRN[3:0]	Input	4	Interrupt Number
IACK	Output	1	Interrupt Acknowledge
RST	Input	1	Reset
CLK	Input	1	Clock
DA[31:0]	Output	32	Data Address
DO[31:0]	Output	32	Data Out
DBS[3:0]	Output	4	Data Byte Strobes
DWE	Output	1	Data Write Enable
DSM	Output	1	Data System Mode
DI[31:0]	Input	32	Data In
DRDY	Input	1	Data Ready
DBGI[31:0]	Input	32	Debug In
STRQ	Input	1	Stop Request
INJI	Input	1	Inject Instruction
DBG0[31:0]	Output	32	Debug Out
STPD	Output	1	Stopped

Clocking

The sf32bl is a fully synchronous design. All flip flops are triggered with the rising edge of the CLK input. All output changes occur after the rising edge of CLK. All inputs are sampled with the rising edge of CLK.

Control signals asserted state

All control signals are active high. The asserted state is '1' and the de-asserted state is '0'. The following signals are affected: **IFT**, **IFNS**, **IFSM**, **IRDY**, **IRQ**, **IACK**, **RST**, **DBS [3 : 0]**, **DWE**, **DSM**, **DRDY**, **STRQ**, **INJI**, **STPD**.

Debug Interface

If the debug interface is not used inputs **STRQ**, **INJI** and **DBGI [31 : 0]** should be connected to ground.

3 Interface Details

3.1 Instruction Fetch

Signals

IA[29:0]	Instruction Address (output); When IFT is asserted IA[29:0] is the address of the 32-bit instruction word to fetch. When IFT is de-asserted IA[29:0] is don't care. The sf32bl ISA defines a 4 GBytes instruction address space with 32-bit instruction addresses pointing to byte locations. Instruction addresses must be aligned on 32-bit boundaries which means the 2 LSBs of instruction byte addresses are always zero. IA[29:0] is a 32-bit word address and represents bits [31:2] of instruction byte addresses.
IFT	Instruction Fetch (output); IFT is the main control signal of the instruction fetch interface. When IFT is asserted outputs IA[29:0] , IFNS and IFSM are valid. When IFT is de-asserted these outputs are don't care.
IFNS	Instruction Fetch Non Sequential (output); When IFT is asserted IFNS indicates if the fetch is sequential (IA[29:0] = address of the preceding fetch + 1) or not (any address due to a change in program flow). When IFT is de-asserted IFNS is don't care.
IFSM	Instruction Fetch System Mode (output); When IFT is asserted IFSM indicates if the fetch is from the system mode (IFSM = 1) or application mode (IFSM = 0) address space. When IFT is de-asserted IFSM is don't care.
II[31:0]	Instruction In (input); When IRDY is asserted II[31:0] must be a valid instruction word. When IRDY is de-asserted II[31:0] is ignored.
IRDY	Instruction Ready (input); IRDY is the acknowledge handshake signal following IFT instruction fetch requests. IRDY must be asserted only as a response to an IFT request. For zero wait state instruction fetches IRDY must be asserted in the cycle following an IFT request. Wait states are inserted by delaying the assertion of IRDY by the required #of clock cycles.

General Rules

The sf32bl instruction fetch timing is designed for direct connection of synchronous memories. The following rules apply:

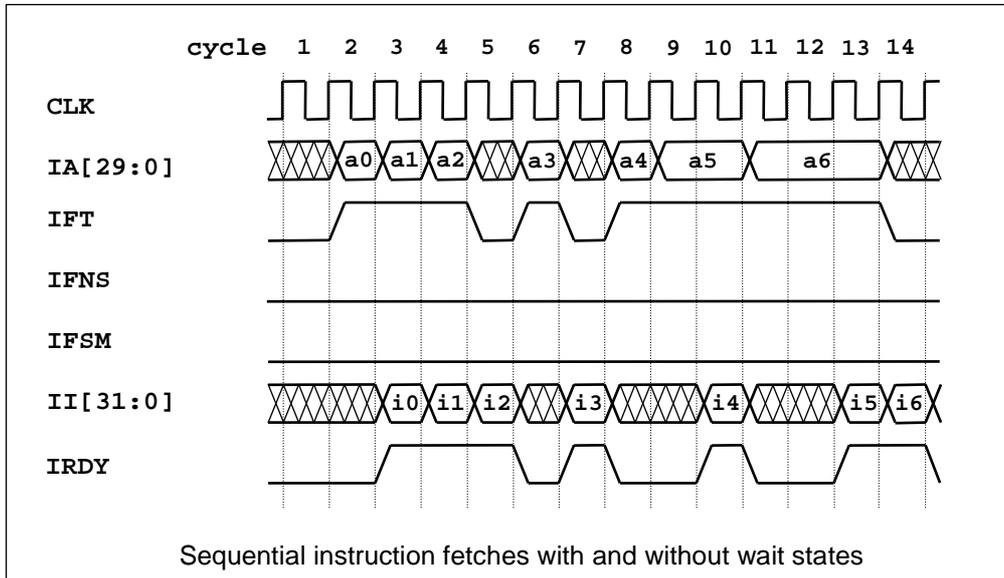
- Based on a handshake with **IFT** as request and **IRDY** as acknowledge
- For zero wait states fetches **II[31:0]** must be provided and **IRDY** must be asserted in the next cycle following an **IFT** request.
- If **II[31:0]** is not ready in the next cycle following **IFT** an arbitrary number of wait cycles can be inserted by delaying the assertion of **IRDY** until **II[31:0]** is ready.
- **IFT** asserted with **IFNS** de-asserted indicates sequential fetches. The address **IA[29:0]** is the address of the preceding fetch + 1.
- **IFT** and **IFNS** both asserted indicate non-sequential fetches. **IA[29:0]** can have any value with no relation to the preceding fetch. A preceding fetch not yet completed is aborted. The next **IRDY** and related **II[31:0]** are interpreted as response to the non-sequential fetch.
- When **IFT** is asserted **IFSM** indicates if the fetch is from the system mode address space (**IFSM** = 1) or from the application mode address space (**IFSM** = 0). Bus controllers can use **IFSM** to protect system code from execution in the application mode.
- **IFSM** is asserted for all instruction fetches of interrupt routines.
- If system mode protection is not required **IFSM** can be ignored.

Sequential Fetches

The figure below shows sequential instruction fetches in application mode with **IFNS** and **IFSM** de-asserted. There are gaps with no instruction fetches in cycles 1, 5 and 7. Although the sf32bl pipeline architecture is designed for one instruction per clock throughput instruction fetching gaps can occur as a result of pipeline stalls or execution of multi-cycle instructions, e.g. multiply or load/store multiple registers instructions. The processor fetches sequential instructions only when there is space available in its pre-fetch buffers. The pre-

fetch buffer concept makes sure that the processor never discards and re-reads sequential instruction words independent of instruction execution times and pipeline stalls.

The fetches in cycles 2, 3, 4, and 6 are done with zero wait states. Instruction words **i0**, **i1**, **i2** and **i3** read from addresses **a0**, **a1**, **a2** and **a3** are provided in the next cycle following the fetch and **IRDY** is asserted. Fetching of instruction words **i4** and **i5** from addresses **a4** and **a5** in cycles 8 and 9 is done with 1 and 2 wait states respectively. Fetching of **i6** from **a6** is done with zero wait states again.



The fetches with wait states show an important behavior of the sf32bl's pipelined instruction interface. With no pending fetch (waiting for **IRDY** of the preceding fetch) **IFT** is asserted with **IA[29:0]**, **IFNS** and **IFSM** valid for only one cycle. The fetch from **a4** in cycle 8 of the diagram illustrates the behavior. In cycle 9 the next fetch from **a5** is driven on the interface. Because the fetch from **a4** is not acknowledged yet in cycle 9 the fetch from **a5** remains stable on the interface. This means that bus logic that inserts wait states, e.g. to let another client access the instruction memory must latch the instruction address and control signals. E.g. if in the example shown below the bus controller grants access to the instruction memory to another client in cycle 8 and then reads from **a4** in cycle 9 to have **i4** ready in cycle 10 the address **a4** and corresponding control signals must be latched in registers because they are not available anymore at the interface in cycle 9.

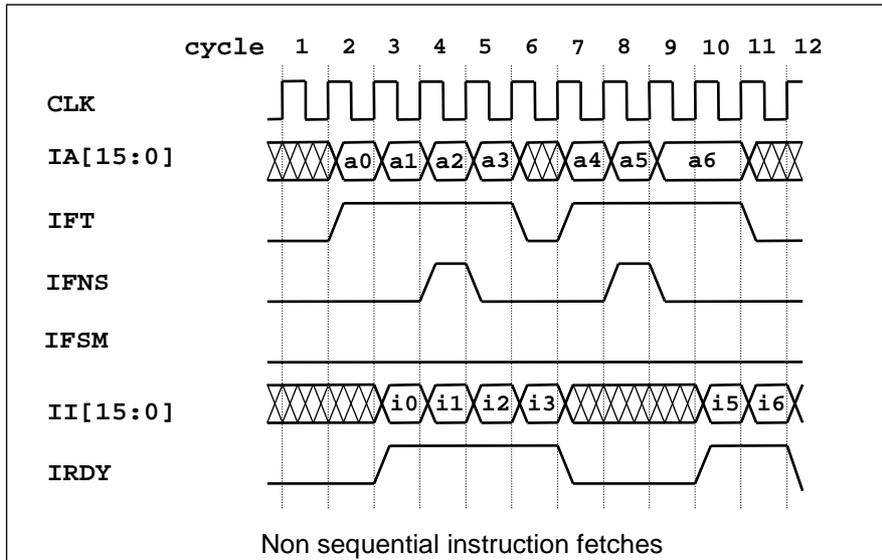
Non Sequential Fetches

Non sequential instruction fetches occur as a result of program flow changes (jump, branch, return or interrupt). The processor flushes the instruction pre-fetch buffer and does not wait for **IRDY** of a preceding fetch. If **IRDY** is asserted in the same cycle the corresponding instruction word **II[31:0]** is ignored. A pending instruction fetch that has not been acknowledged yet when a non-sequential fetch occurs is aborted. This means that the first **IRDY** following a cycle with **IFNS** asserted is always interpreted as acknowledge of the non-sequential fetch.

The next figure shows some example non-sequential fetch timings. The first example is in the middle of a fetch sequence with no wait states. In cycle 4 **IFNS** indicates a non-sequential fetch from **a2**. Instruction word **i1** from the fetch in cycle 3 is discarded. The second example shows a case where a preceding fetch is aborted. The non-sequential fetch from **a5** is started in cycle 8. Due to wait states the preceding fetch from **a4** is not completed yet. The instruction bus logic aborts this fetch and reads directly from **a5**. Instruction word **i5** is delivered (in the example with one wait state) in cycle 10.

In systems with no instruction fetch wait states, e.g. with a synchronous instruction memory directly connected output **IFNS** can be ignored. In systems with wait states e.g. with an instruction cache or with debug access to the instruction memory **IFNS** must be used to abort pending fetches.

IFSM is shown in the two instruction fetch timing diagram only for completeness of the interface signals. For the timing its state is not directly relevant. It can be viewed as an extra address bit.



3.2 Data Access

Signals

- DA[31:0]** Data Address (output); when **DBS[3:0]** is asserted ($\neq 0$) **DA[31:0]** is the byte address of the data access. When **DBS[3:0]** is de-asserted ($=0$) **DA[31:0]** is don't care.
- DO[31:0]** Data Out (output); for write accesses (**DBS[3:0]** $\neq 0$ and **DWE** asserted) **DO[31:0]** provides the data on four byte lanes. Only byte lanes with **DBS[n]** asserted have valid data. **DBS[0]** controls **DO[7:0]**, **DBS[1]** controls **DO[15:8]**, **DBS[2]** controls **DO[23:16]** and **DBS[3]** controls **DO[31:24]**. When **DWE** is de-asserted or when a **DBS[n]** is de-asserted the corresponding byte lane within **DO[31:0]** is don't care.
- DI[31:0]** Data In (input); when **DRDY** is asserted as response to a read access input data is expected at the active byte lanes of **DI[31:0]**. Active byte lanes are determined by **DBS[3:0]** in the cycle where the access was requested. **DBS[0]** controls **DI[7:0]**, **DBS[1]** controls **DI[15:8]**, **DBS[2]** controls **DI[23:16]** and **DBS[3]** controls **DI[31:24]**. When **DRDY** is de-asserted **DI[31:0]** is ignored.
- DBS[3:0]** Byte Strobes (output); **DBS[3:0]** is the main control signal of the data access interface. When **DBS[3:0]** is asserted ($\neq 0$) outputs **DA[31:0]**, **DWE** and **DSM** are valid. With **DWE** asserted **DO[31:0]** provides write data on the **DBS[3:0]** selected byte lanes. When **DBS[3:0]** is de-asserted these outputs are don't care. The following table summarizes the **DO[31:0]** and **DI[31:0]** byte lane use. Note that for reads the **DI[31:0]** byte lanes are active in the response cycle (**DRDY** asserted) following a **DBS[3:0]** request.

Access Type	DBS[3:0]	DA[1]	DA[0]	Valid Output	Expected Input
byte (8-bit)	0x1	0	0	DO[7:0]	DI[7:0]
	0x2	0	1	DO[15:8]	DI[15:8]
	0x4	1	0	DO[23:16]	DI[23:16]
	0x8	1	1	DO[31:24]	DI[31:24]
short (16-bit)	0x3	0	don't care	DO[15:0]	DI[15:0]
	0xC	1		DO[31:16]	DI[31:16]
long (32-bit)	0xF	don't care		DO[31:0]	DI[31:0]

- DWE** Data Write Enable (output); When **DBS[3:0]** is asserted ($\neq 0$) **DWE** indicates if the data access is a read (**DWE=0**) or write (**DWE=1**); When **DBS[3:0]** is de-asserted ($=0$) **DWE** is don't care

DSM	Data System Mode (output); When DSB[3:0] is asserted ($\neq 0$) DSM indicates if the data access is to the system address space (DSM =1) or application address space (DSM =0); When DSB[3:0] is de-asserted ($=0$) DSM is don't care
DRDY	Data Ready (input); DRDY is the acknowledge handshake signal following DBS[3:0] $\neq 0$ data access requests. DRDY must be asserted only as a response to a DBS[3:0] $\neq 0$ request. For zero wait state data accesses DRDY must be asserted in the cycle following a DBS[3:0] request. Wait states are inserted by delaying the assertion of DRDY by the required #of clock cycles. For read accesses input data must be provided at the active byte lanes of DI[31:0] in the cycle when DRDY is asserted.

General Rules

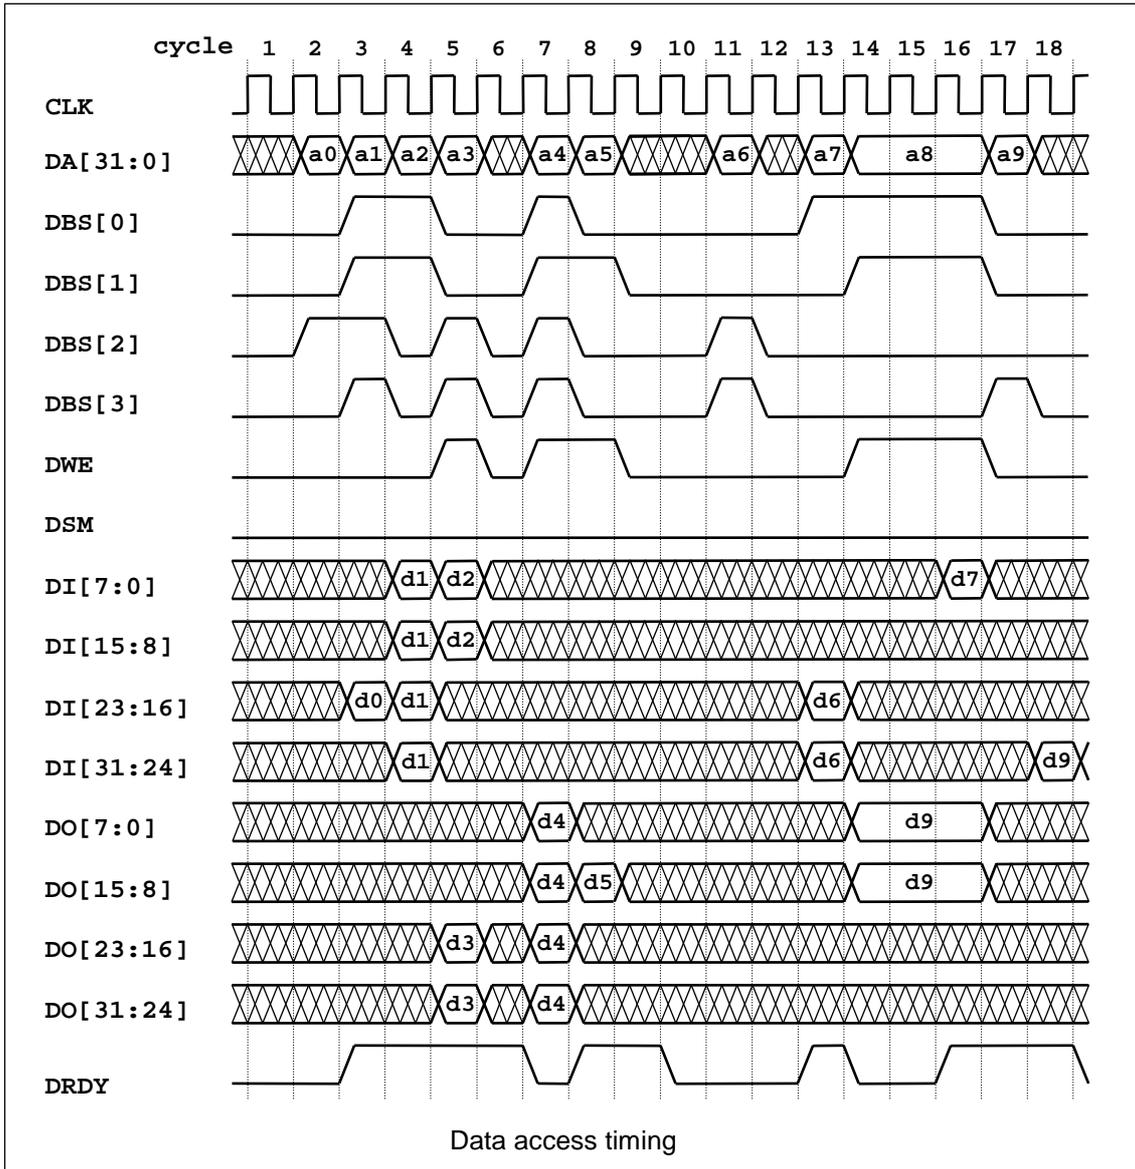
As with the instruction interface the sf32bl data interface is designed for direct connection of synchronous memories. The following rules apply:

- Based on a handshake with **DBS[3:0]** as request and **DRDY** as acknowledge
- For zero wait states access **DRDY** must be asserted in the next cycle following a **DBS[3:0]** request; for read accesses data must be provided at **DI[31:0]** in that cycle.
- If an access can't be serviced with zero wait states an arbitrary number of wait cycles can be inserted by delaying the assertion of **DRDY**.
- The maximum access rate is one access per cycle.
- To build secure systems the **DSM** output can be used to protect critical system resources from illegal access by application mode code. If this is not required the **DSM** output can be ignored.
- Read accesses to the interrupt vector table are done with **DSM** asserted.

Timing

The following diagram shows the sf32bl data access timing. Cycles 2 to 8 are zero wait state accesses of different types (size, read/write). Cycles 11 to 18 are data accesses with and without wait states and show an effect that are important to keep in mind when designing data bus logic for the sf32bl. As with the instruction fetch interface with no pending transaction output signals are driven for only one cycle. E.g. the 16-bit read from **a6** in cycle 11 is completed with one wait state in cycle 13, but the data bus output signals are de-asserted in cycle 12 (before the transaction is completed). Another example is the 8-bit read from address **a7** in cycle 13. The processor then starts a 16-bit write to **a8** in cycle 14. Because the read from cycle 13 is not completed yet the output signals (in this case **DA[31:0]**, **DBS[1:0]**, **DO[15:0]** and **DSM**) are extended until the preceding access is completed in cycle 16. In systems with data access wait states the bus logic must latch the output signals to be able to perform the access later when these signals are no longer valid.

DSM is shown only for completeness of the interface signals. For data access timing its state is not directly relevant. As with the instruction fetch interface it can be viewed as an extra address bit.



3.3 Interrupts

Signals

- IRQ** Interrupt Request (input); **IRQ** asserted signals an interrupt request with number **IRN[3:0]** to the processor
- IRN[3:0]** Interrupt Number (input); when **IRQ** is asserted **IRN[3:0]** is the number of the requested interrupt; when **IRQ** is de-asserted **IRN[3:0]** is ignored
- IACK** Interrupt Acknowledge (output); **IACK** is asserted for one cycle when the processor has latched **IRN[3:0]** and starts interrupt execution.

The interrupt concept of the sf32bl is designed for maximum performance. When a new interrupt request is acknowledged the processor does not flush the pipeline. Fetching of the interrupt vector from the data address space is done one the fly. As soon as the interrupt vector has been received instruction fetching is diverted to the interrupt start address. In ideal cases with no bus wait states and dependency stalls the only execution time overhead is the execution of the `rtir` instruction at the end of the interrupt service routine.

General Rules

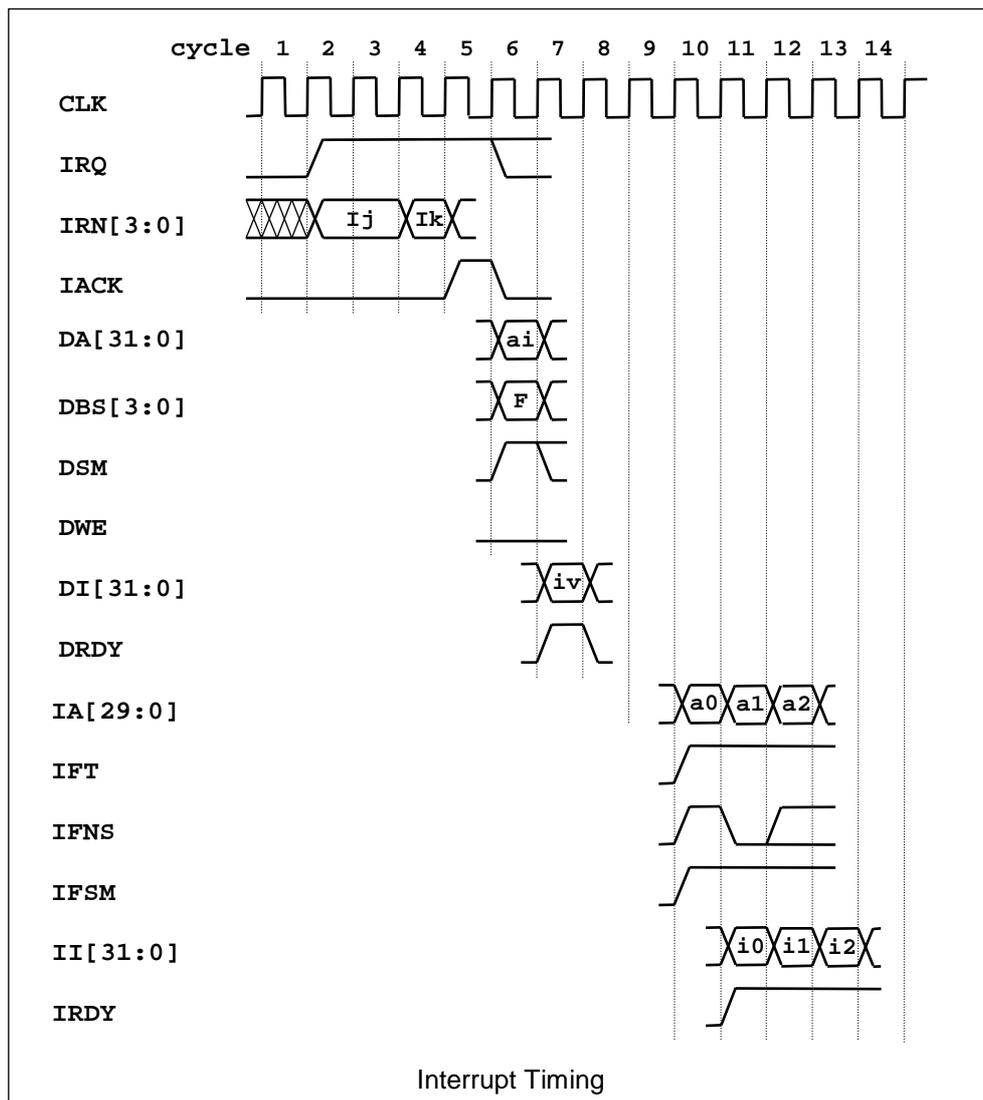
- Interrupts are acknowledged and executed only if enabled (see ISA reference manual)
- The interrupt number $IRN[3:0]$ may be changed from cycle to cycle at any time also while IRQ is asserted. When $IACK$ is asserted the $IRN[3:0]$ of the preceding cycle has been latched and the corresponding service routine will be executed.
- Simple interrupt controllers with no request queuing can ignore the $IACK$ signal

Timing

The following diagram is an example interrupt timing of a sf32bl system. Interrupt processing affects also signals of the instruction fetch and data access interfaces. To keep the diagram simple and clear only sections that are relevant for interrupt processing are shown for each signal and all instruction and data accesses are completed with zero wait states.

The sequence starts in cycle 2 with the assertion of IRQ and interrupt number Ij on $IRN[3:0]$. In most cases if the processor is not already executing another interrupt $IACK$ is asserted in the next cycle following IRQ . In the example $IACK$ is asserted later in cycle 5 to demonstrate that $IRN[3:0]$ is allowed to change while IRQ is asserted. $IACK$ asserted in cycle 5 indicates that the $IRN[3:0]$ value Ik of cycle 4 has been latched inside the processor and is the interrupt number that will be processed. Output $IFNS$ is not shown in cycles 1 to 6 because it is not relevant if the last instructions fetches before an interrupt is started are sequential or non-sequential. $IFSM$ is de-asserted because if not the processor would already be in an interrupt routine and the new request would not be acknowledged.

In the cycle following the $IACK$ pulse a 32-bit data read ($DBS[3:0] = 0xF$, $DWE = 0$) from address ai of the system data address space (DSM asserted) fetches the interrupt vector (start address).



With no wait states the data access is completed in cycle 7 with **DRDY** asserted and the instruction vector **iv** available at **DI [31 : 0]**. Three cycles later in cycle 10 fetching of instructions of the interrupt service routine starts. The first fetch from **a0** (**a0 = iv**) is non-sequential and **IFNS** is asserted. Because the pre-fetch buffers and execution pipeline are completely empty there are at least three consecutive instruction fetches as shown in the diagram. All these fetches have **IFSM** asserted while the processor is in the interrupt service routine.

3.4 Debug

Signals

DBGI [31 : 0]	Debug In (input); this port is used to inject instructions into the processor and to provide input data for the mfdp (move from debug port) and rspc (restore PC) instructions; when INJI is asserted DBGI [31 : 0] is interpreted as 32-bit opcode of the instruction to be injected; when a mfdp or rspc instruction is injected source data must be provided at DBGI [31 : 0] from the cycle following the assertion of INJI .
STRQ	Stop Request (input); the debug module asserts this signal to bring the processor into the debug state. The processor stops fetching new instructions and flushes its pipeline (executes all pending instructions and instructions in the pre-fetch buffer). As long as STRQ remains asserted the processor is held in the debug state; when STRQ is released the processor resumes normal operation.
INJI	Inject Instruction (input); when the processor is in the stopped state (STPD asserted) the debug module asserts INJI for one clock cycle to inject and execute individual instructions; in the cycle where INJI is asserted the opcode of the injected instruction must be provided at DBGI [31 : 0] ; when the processor is not in the stopped state INJI is ignored.
DBGO [31 : 0]	Debug Out (output); when in the stopped state a mtdp (move to debug port) or svpc (save PC) instruction is injected and executed destination data is provided at DBGO [31 : 0] .
STPD	Stopped (output); STPD asserted indicates that the processor is in the stopped state. The processor enters the stopped state after flushing its pipeline either when STRQ is asserted by the debug module or when a stop instruction is executed.

General Rules

- To use the **sf32bl** debug features a separate debug module is required that connects to the processor's debug interface and the debug Host PC. If debug functionality is not required the input signals of the debug port **DBGI [31 : 0]**, **STRQ** and **INJI** should be tied to GND.
- Injecting and executing instructions via the debug port is possible only when the processor is in the stopped state indicated by the **STPD** output signal.
- The stopped state is entered from normal operation either by asserting the **STRQ** input or by executing a **stop** instruction.
- To resume normal operation when the stopped state has been entered by **STRQ** assertion **STRQ** must be de-asserted.
- To resume normal operation when the stopped state has been entered by executing a **stop** instruction **STRQ** must be asserted and then de-asserted.

Timing

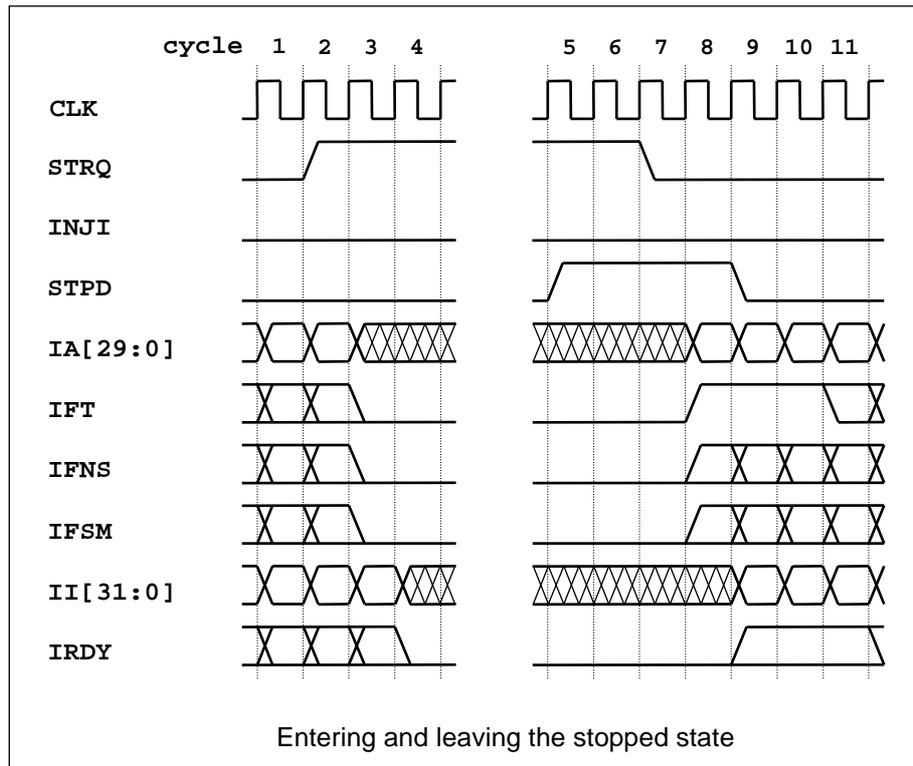
The first diagram following shows the interface timing at the beginning and end of the stopped state. The signals of the debug and instruction fetch interfaces are shown. In cycle 2 **STRQ** is asserted. Starting with the next cycle the processor stops fetching new instructions. Pending instructions and instructions in the pre-fetch buffer are executed until the pipeline is completely empty. When the pipeline is empty the **STPD** output is asserted indicating that the stopped state has been reached.

The stopped state can also be entered by executing a **stop** instruction during normal operation. When a **stop** instruction is executed remaining instructions in the pre-fetch buffer are discarded and the processor asserts the **STPD** output and enters the stopped state when the execution pipeline has been flushed.

In the diagram **STRQ** is de-asserted again in cycle 7 only 2 cycles after the stopped state has been entered. Normally this would not make much sense but the purpose of this diagram is to illustrate the timing only at

the beginning and end of the stopped state.

In the next cycle after **STRQ** has been de-asserted the processor starts fetching instructions again. One cycle later in cycle 9 the **STPD** output is de-asserted indicating that the processor has resumed normal operation.



The second diagram shows the timing of instruction injection and data I/O via the debug port while the processor is in the stopped state. The **STRQ** signal is not shown because it is not relevant if the stopped state has been entered due to **STRQ** assertion or after the execution of a **stop** instruction. Regarding interface timing there are three types of instruction injection:

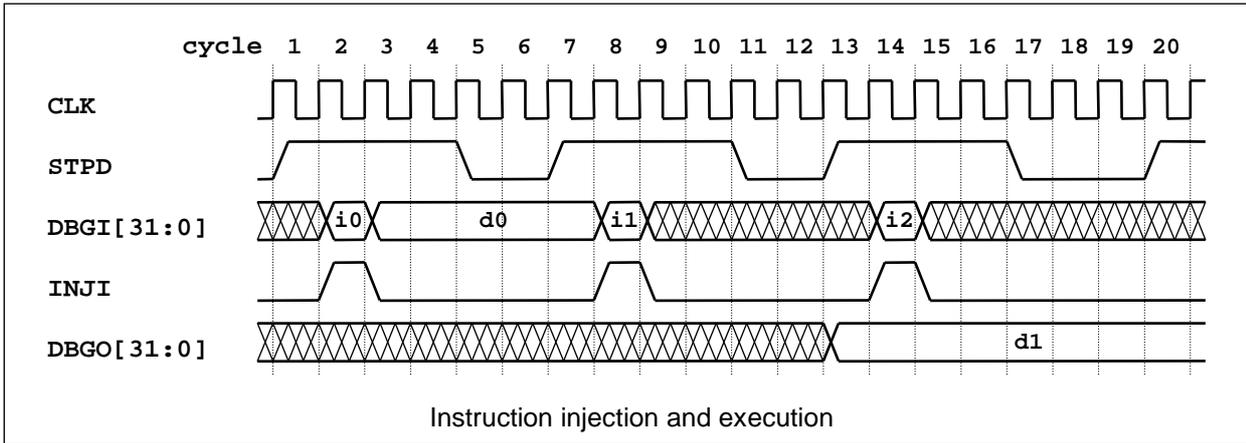
1. Injection with data input from **DGBI[31:0]**; only the dedicated debug instructions **mfdp** and **rspc** take a source operand from the debug port
2. Injection with data output to **DGBO[31:0]**; only the dedicated debug instructions **mtdp** and **svpc** output a destination operand to the debug port
3. Injection with no data I/O; all other instructions are of this type

The injection and execution behavior is common to all three types. The debug module asserts **INJI** for one cycle and drives the opcode of the instruction at **DBGI[31:0]**. Three cycles later the processor de-asserts **STPD** which indicates the execution of the injected instruction. **STPD** is asserted again when execution has finished and the pipeline is completely empty. The number of cycles **STPD** is de-asserted depends on the injected instruction. For some flow instructions like **svpc** or **rspc** **STPD** is de-asserted for only one cycle. For most computation instructions it is 2 or 3 cycles. For load/store instructions with zero wait states data access it is at least 5 cycles. Data access wait states add to the **STPD** de-asserted time. The debug module must wait for **STPD** being de-asserted and asserted again before it can inject the next instruction.

A type 1 example starts in cycle 2 where opcode **i0** is injected. In the following cycle when **INJI** is de-asserted the source operand **d0** is driven at **DGBI[31:0]**. It must be kept stable until **STPD** is re-asserted.

A type 2 example starts in cycle 8 where opcode **i1** is injected. When **STPD** is re-asserted in cycle 13 the destination operand **d1** is available at **DGBO[31:0]**. The example shows the behavior of an **mtdp** instruction. The second type 2 instruction **svpc** has different timing. **STPD** is de-asserted for only one cycle and the destination operand appears at **DGBO[31:0]** already in that cycle. The debug module should read the destination operands of type 2 instructions when **STPD** has been re-asserted. **DGBO[31:0]** is always valid then and remains stable until the next type 2 instruction is injected.

A type 3 example starts in cycle 14 where opcode **i2** is injected. **STPD** is de-asserted for three cycles which is a typical value for computation instructions



3.5 Reset

Signals

RST Reset (input); this is a synchronous reset; when **RST** is asserted the processor is reset with the next rising edge of **CLK**.

General Rules

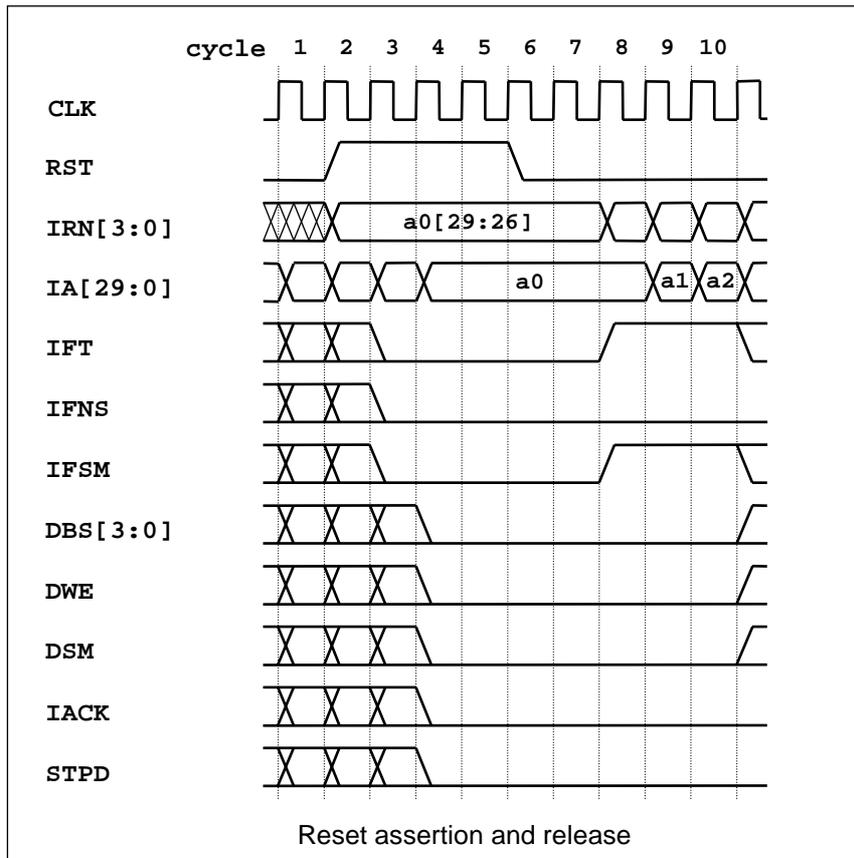
- **RST** can be asserted at any time. The processor does not wait for any pending interface transactions or instructions.
- **RST** needs to be asserted for only one active edge of **CLK** to fully reset the processor.

Timing

The diagram shows the **RST** assertion and release timing. Only signals that are directly affected are shown. Output signals not shown are either undefined or keep their state. Beside **RST** the only input signal relevant for reset is **IRN[3:0]**.

In cycle 2 **RST** is asserted. In the following cycle (3) instruction fetching stops, **IFT**, **IFNS** and **IFSM** are de-asserted. One cycle later (4) all control outputs of the processor are de-asserted. **IA[29:0]** takes the value of the reset start address **a0**: **IA[29:26] = IRN[3:0]** and **IA[25:0] = 0**. This state remains unchanged as long as **RST** remains asserted.

In cycle 6 **RST** is de-asserted. **IRN[3:0]** must continue to provide the upper 4 bits of the reset start address for the following 2 cycles. In cycle 8 instruction fetching starts from **a0** with **IFSM** asserted (system mode).



IRN[3:0] can take any value from this point.

4 Instruction Execution Timing

4.1 Effective Execution Times

The following table provides effective execution times for all sf32bl instructions except for the dedicated debug instructions `mtdp`, `mfdp`, `svpc`, `rspc` and `stop` which are not for use in normal program code sequences.

The numbers provided in the **Cycles** column are best case numbers assuming no stalls caused by operand dependencies or data access wait states (load/store instructions). The **Stalls** column contains abbreviations of stall conditions that are further explained in the “**Stall Conditions**” section later in this chapter.

Instructions are grouped by addressing modes and common execution time properties. Instructions with multiple addressing modes may appear in different non-consecutive places.

Instructions	Addressing Mode	Cycles	Stalls	Comment(s)
<code>move</code>	$C17_s, Rd, CND$	1	-	-
<code>comp cmpc</code>	$C17_s, Rs1$	1	D1	-
<code>addt subf</code> <code>andb iorb xorb</code>	$C16_U, Rs1, Rd$	1	D1	-
<code>mlcu mlcs</code>		2		
<code>addh</code>	$C32_U, Rs1, Rd$	1	D1	-
<code>addt subf addc subc</code>	$C12_U, Rs1, Rd, CND$	1	D1	-
<code>move negt abs1 inv</code> <code>clzr ibos ibol</code>	Rs, Rd, CND	1	D1	-
<code>btst btcl bttg</code>	$BTI5_U, Rs1, Rd, CND$	1	D1	-
<code>btts</code>	$BTI5_U, Rs1$	1	D1	-
<code>shlz shlf shru shrs</code>	$SHC5_U, Rs1, Rd, CND$	1	D1	-
<code>addt subf addc subc</code> <code>andb iorb xorb</code>	$Rs0, Rs1, Rd, CND$	1	D1	-
<code>btst btcl bttg</code> <code>shlz shlf shru shrs</code>			D1, D2	
<code>mult</code>		3	D1	iterative
<code>mlhu mlhs</code>		4		
<code>btts</code>	$Rs0, Rs1, CND$	1	D1	-
<code>comp cmpc</code>	$Rs0, Rs1$			
<code>stbt</code> <code>stsh</code> <code>stlg</code>	$Rs, DA16_s, CND$	1	D1	data access wait states add to the effective execution time
	$Rs, (DO12_s, An), CND$			
	$Rs, (An, Au12_s) *, CND$			
	$Rs, (An, Ru) *, CND$			
	$Rs, (Rx, An), CND$			
	$RGS, (An) +$ $RGS, -(An)$	n	D1	n = #of register in RGS
<code>ldbz ldbs</code> <code>ldsz ldss</code> <code>ldlg</code>	$DA16_s, Rd, CND$	1	-	Data access wait states add to the effective execution time
	$(DO12_s, An), Rd, CND$			
	$(An, Au12_s) *, Rd, CND$			
	$(An, Ru) *, Rd, CND$			
	$(Rx, An), Rd, CND$			
	$(An) +, RGS$			
	$-(An), RGS$	n	D1	n = #of registers in RGS

stie clie scie rsie	implied	1	-	-
jump jpsr	IA29 _U	2	-	-
jump jpsr	implied (TA)		D3	TA dependency
rtsr	implied	2	D4	SP dependency
rtir			-	-
brlc	IO16 _S	2	D5	branch taken
		1		branch not taken
		0		loop cache hit
brxx (conditional)	IO16 _S	2	D6	branch taken
		1		branch not taken

4.2 Stall Conditions

Extra cycles add to best case execution times if stall conditions caused by operand dependencies occur during instruction execution. Instructions have to wait if one or more of their source or destination operands are scheduled to be updated by a preceding instruction that has not finished execution yet. These conditions are instruction and addressing mode specific. Seven conditions exist named **D1** to **D6**. Affected instruction/addressing-mode combinations have the relevant conditions listed in the **Stalls** column of the execution times table.

The following paragraphs are more detailed descriptions of individual stall conditions with hints how they can be avoided.

Computation Latency

With the sf32bl pipeline structure register destination operands of computation instructions are updated only one cycle after a directly following instruction reads its source operands. If the following instruction has the same register as source operand this instruction would have to be stalled by one cycle to wait until the source operand is ready. For most cases a forwarding mechanism is implemented that uses the ALU output as source directly and bypasses the register file to avoid stalls. Exception is the **D3** stall condition.

D1 (Source operand pending update)

The **D1** stall occurs if instructions use the destination register of a directly preceding load from memory instruction as source operand. Load from memory instructions update their destination register three cycles after an immediately following instruction reads its source operands. As with destination operands of computation instruction there is a forwarding mechanism that bypasses the register file and uses the load from memory destination operand as source one cycle before it is written to the register file. But because of three cycles latencies there is still a two-cycles stall if an instruction uses the destination register of a directly preceding load from memory instruction as source operand.

In many cases such stalls can be avoided by instruction re-ordering so that there are at least two other instructions between the load from memory and the instruction that uses the load destination as source.

D2 (Computation destination not forwarded)

As described under **Computation Latency** a forwarding mechanism bypasses the register file to avoid stalls if instructions try to read the destination register of a directly preceding computation instruction as source operand. But there are some exceptions where forwarding is not done because it would create critical timing paths and decrease the processor's maximum clock rate. The following types of register source operands are affected and cause a one-cycle stall if the register is the destination operand of the directly preceding instruction:

- Indirect shift counts of shift instructions
- Indirect bit index of bit manipulation instructions
- Index of load/store instructions with indirect + index addressing mode

To avoid **D2** stalls instructions must be re-ordered such that the instruction that generates the register destination operand is not the directly preceding instruction.

D3 (TA pending update)

In principle this is the same as **D1** but for register **TA** (Target Address) and only for the **jump** and **jpsr** instructions with the implied addressing mode that use **TA** as source operand. There is no forwarding when **TA** is used as indirect jump address. A preceding instruction that updates **TA** with a latency > the cycle distance to the **jump/jpsr** with **TA** as source causes stall cycles. This is the case for load instructions with **TA** as destination operand and a cycle distance < 3 and also for directly preceding **move** instructions with **TA** as destination operand.

D4 (rtsr with SP pending update)

This is similar to **D3** but for register **SP** (Stack Pointer) and only for the **rtsr** instruction that uses **SP** as source operand. There is no forwarding when **SP** is used as return address. A preceding instruction that updates **SP** with a latency > the cycle distance to the **rtsr** causes stall cycles. This is the case for load instructions with **SP** as destination operand and with a cycle distance < 3 and also for directly preceding **move** instructions with **SP** as destination operand.

D5 (br1c with LC pending update)

Again this is similar to **D3-D4** but for register **LC** (Loop Counter) and for **br1c** instructions which use **LC** as source operand. There is no forwarding when **LC** is used as source operand of loop counter branches. A preceding instruction that updates **LC** with a latency > the cycle distance to the **br1c** causes stall cycles. This is the case for load instructions with **LC** as destination operand and with a cycle distance < 2 and also for directly preceding **move** instructions with **LC** as destination operand.

D6 (brxx with unresolved speculation)

The sf32bl implements a speculation scheme for conditional branches. Conditional branch instructions do not wait until pending **CC** updates of preceding instructions are completed. However a stall condition occurs if a new conditional branch is decoded and the speculation of a preceding branch is not yet resolved. In this case the new branch has to wait until all **CC** update operations of instructions that have been issued before the preceding branch are completed. **D6** stalls typically occur with instruction sequences where two or more conditional branches are placed close together.

4.3 Loop Cache

The sf32bl implements a single entry loop cache to accelerate the performance of simple loops. When a **br1c** (loop counter branch) instruction with backward branch direction is executed and the loop cache is not active yet the processor latches start-address, end-address and the opcode of the first instruction of the loop in internal registers.

From the second iteration on the last instruction (loop branch) and the first instruction of the loop are not fetched from instruction memory to eliminate the 2-cycle latency that the non-sequential instruction fetch of the loop branch would cause. Instead the loop branch is implicitly executed and the opcode of the first instruction is retrieved from the register where it has been stored when the cache became active. As a result the effective execution time of the loop branch is zero cycles.

If beside the implicit loop branch any other non-sequential instruction execution is encountered the loop cache is de-activated immediately. E.g. if a conditional branch inside the loop is taken the cache is de-activated and the next execution of the loop branch is done normally as in the first iteration but with this normal **br1c** execution the cache is activated again. If in the next iteration the conditional branch is not taken the cache concept will take effect and the next loop back branch will be done with zero cycles effective.

5 Compatibility

5.1 Software

The sf32bl is fully compatible with the sf32 (b) (base) ISA.

5.2 Hardware

The sf32bl interface signals and timing are the same as those of the sf32bu and sf32dl. These three processors can replace each other without changing any of the surrounding hardware.

5.3 Replacement Options

Sf32bl and sf32bu are software compatible and can replace each other regarding both hardware and software.

The sf32dl can replace sf32bu and sf32bl regarding both hardware and software because it supports the sf32 (b) (base) ISA.