

# An Innovative, Segmented High Performance FPGA Family with Variable-Grain-Architecture and Wide-gating Functions

Om Agrawal, Herman Chang, Brad Sharpe-Geisler, Nick Schmitz, Bai Nguyen, Jack Wong,  
Giap Tran, Fabiano Fontana and Bill Harding

Vantis Corporation  
995 Stewart Drive  
Sunnyvale, CA 94088

## 1. ABSTRACT

This paper describes the Vantis VF1 FPGA architecture, an innovative architecture based on 0.25u (drawn) (0.18u Leff)/4-metal technology. It was designed from scratch for high performance, routability and ease-of-use. It supports system level functions (including wide gating functions, dual-port SRAMs, high speed carry chains, and high speed IO blocks) with a symmetrical structure. Additionally, the architecture of each of the critical elements including: variable-grain logic blocks, variable-length-interconnects, dual-port embedded SRAM blocks, I/O blocks and on-chip PLL functions will be described.

## 2. GOALS, CHALLENGES AND APPROACH FOR THE VF1 FAMILY

In developing Vantis' first FPGA family from a clean slate, our goal was to focus initially on the mainstream market needs of 12-36K+ logic gate range with an architecture that was scalable to higher densities. It had to be a high-performance, re-programmable FPGA family using SRAM technology; with a mainstream look-up-table (LUT) based architecture. It had to be easy-to-use with a competitive cost structure. The combination of high performance and ease-of-use was our highest priority.

One particular challenge for us was to develop a mainstream look-up-table based architecture that was innovative, and differentiated with a superior value proposition. We believed from the beginning, despite FPGAs being around for more than a decade, that there was significant room for innovation, differentiation, and creativity. To differentiate, we focused on each and every element: the logic block, memory, interconnect, I/O block, clocking and miscellaneous functions such as PLL. While the basic building blocks (such as logic, interconnect, memory, and I/Os) were conceptually the same as other mainstream FPGAs, the fundamental differences had to be in the basic architecture and structure of each and every component.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

FPGA 99 Monterey CA USA

Copyright ACM 1999 1-58113-088-0/99/02...\$5.00

We were very aware that for FPGAs using deep-sub-micron technology, gate delays tend to be quite fast, but interconnect and wiring delays are the dominant portion of total delay. Programmable interconnect switches, distributed resistance and capacitive loading associated with FPGAs make it worse.

Our key technical challenges for developing the architecture included optimal balancing of logic and interconnect resources with maximum performance, ease-of-use and competitive cost structures. Additionally, we wanted to develop an architecture that exploited silicon features with mapping and place and route tools as well as an FPGA family that maximized the application spectrum for both control and data path applications.

We knew fixed grain look-up-table logic block structures, whether narrow or wide, resulted in potentially inefficient solutions for both area and speed. To achieve a high-performance FPGA family we focused on three things: a) an innovative logic block designed with variable-logic-granularity, b) an intelligent, segmented variable-length-interconnect structure designed for performance, and c) innovative mapping and place and route algorithms optimized for our architecture. In order to address broader application spectrum needs for both data path and control applications, we also needed an FPGA that was not only register rich and interconnect rich, but fan-in rich.

## 3. THE VANTIS VF1 FPGA FAMILY OVERVIEW

Designed with an advanced 0.25u (drawn), (0.18u Leff) four-metal, CMOS SRAM technology, the VF1 family is designed to offer density spanning from 12K to 36K logic gates (85,000 system gates when on-chip embedded memory is included). It has a scaleable architecture allowing density up to 250K+ logic gates (1 million system gates with on-chip embedded memory). With 7 to 7.5ns fast pin-to-pin delays, internal data path pipe-lined frequency up to 250+ MHz and external bus speed up to >133 MHz, the VF1 FPGA family is designed to meet the needs of high performance datacom and telecom system designers.

Vantis FPGAs are intended to address mainstream design needs for both register and I/O intensive applications with packages ranging from 144-pin TQFP at the low-end to 160/208/240-PQFP at the mid-range, and 256/352 BGA

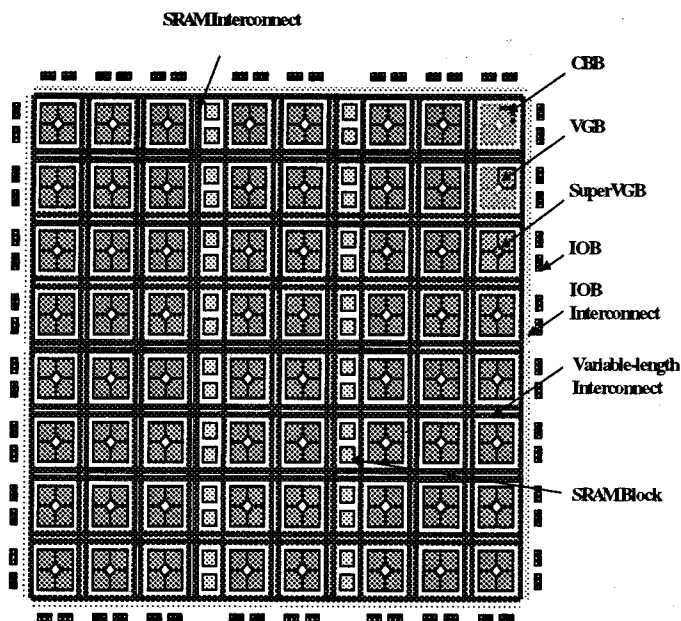
packages at the high-end. This includes up to 292 I/Os, and logic FFs ranging from 784 at the low-end to >2,300 at the high-end. The VF1 family also features 2.5-volt internal operation, with 3.3V VCC and 5V tolerant I/Os allowing optimal performance and compatibility with existing voltage standards. These devices are infinitely re-programmable with an in-system-programmability (ISP) feature. They also offer 100% guaranteed testability with a focus on system oriented features such as embedded dual-port RAMs, 3-state buses, carry chain, high speed IO blocks, slew rate control, phase-lock-loops (PLL) etc. [1].

### 3.1 Vantis VF1 FPGA Variable-Grain-Architecture™ with Logic Hierarchy

Figure 1 shows the block diagram of a Vantis VF1 FPGA device. A Variable-Grain-Architecture is at the heart of the VF1 family. The Variable-Grain-Architecture allows the device to vary the logic grain size of an FPGA to suit a broad range of applications, as well as allowing it to be adaptable to the synthesis and mapping tools to optimize the silicon efficiency with the best possible speed.


The VF1 family architecture consists of three distinct levels of logic hierarchy: Configurable-Building-Block™ (CBB™), Variable-Grain-Block (VGB™), and Super-Variable-Grain-Block (SuperVGB™). (Figure 2).

Figure 1. VF1 FPGA Architecture



The Super Variable-Grain Block (Super VGB) is the highest level building block in the VF1 architecture. It is physically laid out in a symmetrical manner and consists of four VGBs arranged in a 2x2 mirrored image manner, that can be


Figure 2. VF1 Super VGB Architecture and Logic Hierarchy



combined to create complex high performance functions with logic width up to 32 or 48 inputs, using local building blocks and local interconnect resources. Supporting SuperVGBs at the top level are dual-port embedded SRAM and input/output blocks.

Variable Grain Block (VGB) is the second level (Figure 3). The VGB includes four CBBs arranged symmetrically in an L-shaped manner and contains logic to combine two or more CBBs to implement wide logic functions. It also has built-in wide-gating logic to support complex logic functions with up to sixteen parallel inputs within a single VGB with 2 levels of LUT delays. The VGB also includes high-speed carry logic to support high performance arithmetic functions and common control logic for all CBBs.

Figure 3. Variable Grain Block (VGB)



**Configurable Building Block (CBB):** The CBB is the lowest level of logic hierarchy. Each CBB consists of two independently operable 8-bit look-up tables (3-LUTs), each with its separate three inputs and one output, and a flexible storage element. Two outputs of these LUTs are multiplexed to provide a 16-bit LUT (4-LUT) capability for a single 4-input function. Thus, each CBB can also implement one 4:1 Mux or two 2:1 mux.

A CBB consists of two parts: a configurable combinational element (CCE) (Figure 4) and a configurable sequential element (CSE) (Figure 5). The CCE receives all logic inputs and generates combinational outputs. The CSE stores and routes the outputs.

The CCE contains two separate 8-bit, three-input look-up tables (3-LUTs). The CCE receives inputs via the VF1 variable-length-interconnect routing resources, direct connections from adjacent VGBs, and local feedback within the VGB. An Input switch matrix routes the inputs to the inputs into the LUTs. The two 3-LUTs may generate individual outputs or they may be combined into a 16-bit 4-LUT that decodes four inputs. If the 3-LUTs operate independently, one output follows the fast-path route to the CBB output. The CSE has a flexible structure. It can receive inputs from the CCE, the carry logic or the wide gating function via a mux. The output of this mux may be stored in the CSE register or it may bypass the register and go directly to an output via a separate mux. The output of this mux drives the direct connect output of the CSE to provide connectivity to adjacent VGBs.

Figure 4. Configurable Combinational Element (CCE)

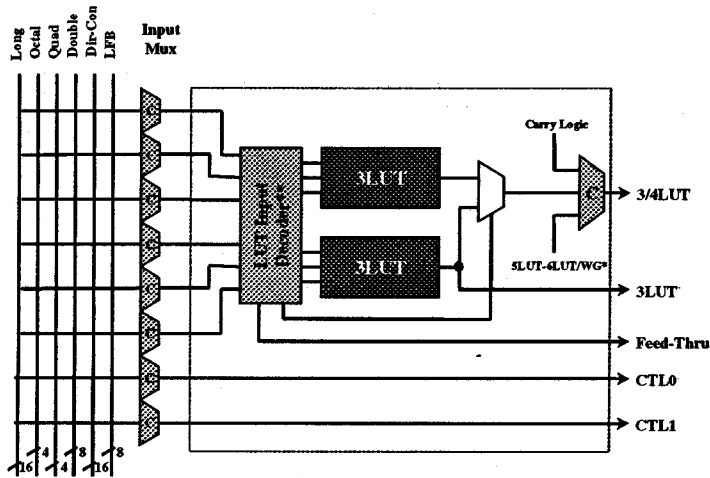
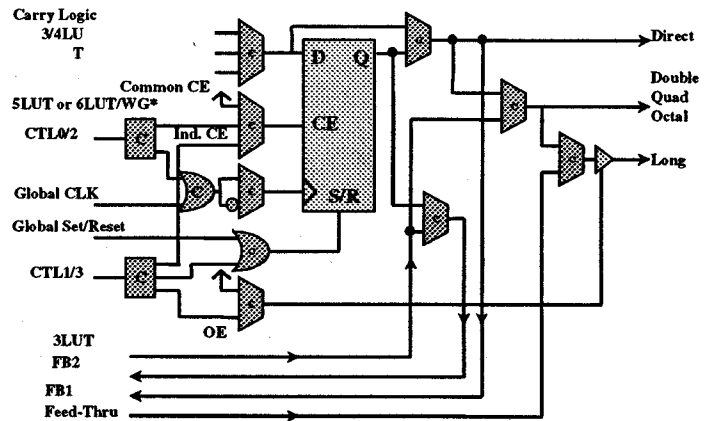



Fig. 5 Configurable Sequential Element (CSE)




Each VGB can be combined in a wide variety of ways to create either very simple or very complex logic structures. Each VGB may be configured as a full CBB quad providing up to 16- or 24-input functions, as one or more twin CBBs providing up to 8- or 12-input functions, or as four single CBBs providing up to 4- or 6-input functions each. A VGB can be viewed as a fine-grained structure when each CBB is used to implement a separate simple logic function, or it can become a coarse-grained architecture when the entire VGB is dedicated to a single, complex function. For maximum logic packing, each VGB can implement 8 functions of 3-inputs, four functions of 4-inputs (Figure 6), two functions of 5-inputs (Figure 7), one function of 6-inputs, or one partial function of 16-inputs. It can also implement either two 8:1 muxes or two 4:1 muxes, and can implement a maximum of 13:1 mux function.

Figure 6. Four-Input CCE Configurations (Four per VGB)



CBBs within a VGB can be combined or operated independently. Two CBBs can be combined per side, while the other two can operate independently. The combined CBBs implement any five-input function while the independent CBBs implement two three- or one four- input functions. All four CBBs can be combined to provide any six input function.

Figure 7. Five-Input Function Using Two CBBs (Two per VGB)



Prior research had shown that for a look-up-table based architecture, the best number of inputs was between three and four, and that it was beneficial to include a D flip-flop in the logic block [2].

We chose two independently operable and composable 3-LUTs with 1 D flip-flop as our basic CBB for the following reasons:

1. Our desire was to expand the granularity to 3-input functions at the low end. We believed that fixed 4-LUT as base building block was expensive and wasteful.
2. Composable 3-LUT was efficient for handling a single 2:1 mux (with one 3-LUT) and a 4:1 Mux (with 2 3-LUTs).

This required each 3-LUT to have its independent 3-inputs with the ability to be shared with the adjacent 3-LUT to operate as 4-LUT. Our desire for Variable-logic-granularity required some overhead within the VGB for configuring logic as 3-LUT, 4-LUT, 5-LUT and 6-LUT. However, the logic overhead was minimal and felt to be worth it for achieving faster speed. Our goal was to maximize the flexibility of each logic block [3, 4, 5, 6, and 7].

### 3.2 Common Control Functions

Each VGB is organized in a nibble-wide fashion with 4 CBBs and 4 flip-flops. All four CBBs within a VGB share flexible common control functions for clock, clock enable, set/reset and output enable. Control functions were shared per VGB basis to minimize die size and achieve balance. Each VGB

derives a common clock signal from 4 global and four local clock sources. However, each CSE has individual polarity control for the clock. This allows certain flip-flops within the VGB to be triggered by different edges of the same clock. A common clock enable signal is also derived from 4 local control signals. Besides a common clock enable, each CSE has the ability to receive an individual clock enable signal. A common local Set/Reset signal is also derived from 4 local control signals. Each CSE receives this common Set/Reset or Global Set/Reset signal. Each CSE has the ability to either Set or Reset its individual storage element on an individual storage element basis. Also, a common output enable control is derived per VGB. This OE control is used to control the shared drivers between VGBs (for the Super VGB) for driving 3-statable long lines.

### 3.3 Special Wide Gating Logic

Special wide-gating logic that is part of the VGB architecture is used to implement configurations up to 48 inputs (partial functions) in only three logic levels. The wide gating logic includes a dedicated 16-bit LUT within each VGB that is used to combine CBBs into functions with up to sixteen inputs using all four CBBs in one VGB (Figure 8). In this example, each CBB within a VGB is configured to fully decode four inputs. Each of the four CBBs generates an output that becomes an input to the 16-bit LUT in the wide gating logic. This LUT fully decodes the four inputs from the CBBs and provides wide-gating function (AND, NAND, NOR) or flexible XOR, X-NOR functions up to 16 or 24 inputs. Output of each CBB is either a full function of 4-inputs or a partial function of 6-inputs. The built-in wide-gating LUT allows signals to stay within a VGB obviating the need for routing signals between VGBs and achieves faster speed. Each VGB can also implement a high-speed 4-bit adder, 4-bit subtractor, 4-bit up/down counter, 4-bit shifter and an 8-bit comparator, using built-in carry logic.

This configuration does not decode all 65,536 possible combinations of sixteen inputs. Instead, it decodes sixteen combinations of four inputs in each CBB for a total of 64 possible combinations. The wide gating 16-bit LUT (4-LUT) decodes sixteen possible combinations. The circuit therefore, decodes 1,024 possible combinations (16\*64). For most logic functions, this is quite adequate, and is accomplished using only the high-speed, short-intra-connect logic contained within a single VGB.

Variable granularity structure does impose slight variation of delay with logic width. Three input functions are faster than 4-input functions. Four-input functions are slightly faster than 6-input functions. Table 1 below summarizes speed for different granularity logic functions.

Table 1. Speed for variable-Grain-Logic Functions

Parameter	Combinatorial CCB Delays	Unit
3-LUT	CBB input -> LUT -> CBB Output	1.8 ns
4-LUT	CBB Input -> LUT -> CBB output	2.6ns
5-LUT	CBB input -> LUT -> CBB output	3.0ns
6-LUT	CBB input -> LUT -> CBB output	4.8ns
Wide gating	CBB input -> LUT -> CBB output	3.4ns
Fast Path	CBB 3-LUT Fast path	1.7ns

It is important to note a couple of things: a) speed for partial wide-gating functions is faster than the full 6-LUT function, and b) there is a special fast path to allow the lower 3-LUT to get its output to other VGBs. This allows the mapping and place/route tools to intelligently map, pack and place/route appropriate width functions within a VGB to achieve best speed or area utilization.




Figure 8. Special Wide Gating Logic -

### 3.4 High Speed CarryLogic

Each VGB includes built-in high-speed carry logic that facilitates the implementation of arithmetic circuits such as adders, subtractors, bit shifters, up/down counters and comparators. To improve the arithmetic speed, the carry chain within a VGB is placed between the CCEs and the CSEs within each CBB.

A VGB receives a carry input from a preceding VGB in the

arithmetic chain, and generates a carry for the following VGB. The carry chain between VGBs starts with the bottom VGB in a column and proceeds vertically through the column. Each column of VGBs has its own carry chain (Figure 9).

Figure 9. Carry Routing between VGBs



### 3.5 Super VGB and Shared Drivers

The Super VGB is arranged in a symmetrical manner and consists of four mirrored VGBs with four sets of shared long connect multiplexers/drivers, arranged in a very symmetrical manner. Each side has a set of 4 3-stateable drivers, to provide expanded wide gating and wide multiplexing and high-speed data connectivity to 3-stateable lines on all four sides. The shared drivers allow functions up to 32 or 48-inputs or up to 26:1 multiplexer function in both horizontal and vertical direction. These shared drivers also allow feed-through of signals for providing high-speed connectivity between parallel and perpendicular lines. Each VGB can drive two sets of 4 3-stateable shared drivers. This allows each VGB to connect to 3-stateable horizontal or vertical long lines in a symmetrical manner. The symmetrical arrangement of the Super VGB improves logic density and minimizes interconnect length for implementing complex functions. Inputs can come from any direction on the chip and outputs can go in any directions. Compared to architectures that force logic paths to flow in one particular direction, this Super VGB symmetry shortens signal paths and thus improves performance and density.

### 3.6 Routing Resources And Hierarchical Interconnect Structure

Two key innovations of Vantis' interconnect architecture are a) embodiments of an orthogonal, symmetrical and homogeneous routing structure, and b) high performance interconnect hierarchy. Orthogonal, symmetrical routing resources offer symmetrical amounts of routing resources in both horizontal and vertical directions. This removes any directional

constraints on the place and route software for any signal flow, and significantly eases the tasks for developing powerful, architecture specific place and route tools.

The VF1 FPGA architecture also incorporates an optimized hierarchical Variable-Length Interconnect™ (VLI™) Hierarchy. The richness of the VLI resources provides many benefits to the user. First and foremost, VLI provides highly predictable performance with minimal variability. Secondly, the hierarchical VLI resources provide optimal length resources for every net and enable silicon efficient results that enhance fitting designs first time. This results in reduced capacitive wire loads and significant power savings for designers.

The VF1 architecture provides four levels of high-performance interconnect resources:

High-speed local feedback provides intra-VGB connectivity and allows CBB outputs to feedback to the inputs of CBBs within the same VGB. Inter-VGB direct connect routes the outputs of every CBB in every VGB to the inputs eight nearby VGBs and to IOBs. Variable-length Interconnect resources provide programmable interconnects that may span two VGBs, four VGBs, eight VGBs, and the entire FPGA.

*Intra-VGB Connectivity* is the lowest level of interconnect hierarchy, and is achieved by eight high-speed dedicated feedback lines provided within a VGB for achieving fast intra-VGB connectivity (two from every CBB) (Figure 10). *Inter-VGB Direct Connect* is the second level of interconnect hierarchy. This rich, symmetrical, orthogonal and extended direct connect structure is a unique innovation of the VF1 FPGA architecture. It offers significant benefits to the users. First, it minimizes connection lengths for maximum performance resulting in faster speed, better speed predictability and improved routability. This alleviates routing congestion and allows a VGB to connect to its 8 adjacent VGBs - 2 in each direction. Second, the orthogonal symmetrical direct connect structure removes the directional constraints on place and route software. (Figure 11). *Inter-VGB VariableLength Resources* serve as the 3<sup>rd</sup> interconnect hierarchy. Twin, Quad, Octal and Long-Connect orthogonal, symmetrical routing resources are available for providing high-speed short, medium and long net connections between VGBs. Twin lines span two VGBs. Quad-lines span four VGBs. Octal lines span eight VGBs. Long lines span from edge-to-edge uninterrupted.

Short connections deliver better performance than long connections and provide better routability. Interconnect resources for the VF1 family in order of performance are: local feedback within a single VGB, direct-connect lines between VGBs and from VGBs to IOBs, double-length lines spanning two VGBs, Quad lines spanning four VGBs, Octal lines spanning eight VGBs; and long lines spanning the entire FPGA. Based on the characteristics of these interconnect lines, the timing driven place and route tool selects appropriate length interconnect wire to meet the timing specified by various timing constraints.

Figure 10. Intra-VGB Connect (Local Feedback)






Figure 11. Inter-VGB Connect



*Symmetrical Interconnect structure* - The interconnect structure for the VF1 family is designed to be symmetrical in both the horizontal and vertical dimension. Each channel includes 16 long lines, 4 Octal lines, 4 Quad lines, 8 double lines, 16 direct connect lines and 8 feedback lines. The long lines are also tri-statable in both vertical and horizontal directions in a symmetrical manner. (Figure 12).

Figure 12. Symmetrical Variable-Length Interconnect Resources



Prior research for segmented FPGAs showed the merits of wire segments of lengths 1, 2 or 4. [8]. We chose wire segments of lengths 2, 4, 8 and longlines and dedicated high speed local connectivity to achieve high performance. Intra-VGB feedback and a rich set of direct connects between VGBs eliminated the need for single-length wire segments. This further allowed us to better optimize our switch boxes [9]. With rich local connectivity (within the VGB) and rich inter-VGB direct connects, there was no reason for single-length wire segments between VGBs

IO and global resources (clock, set/reset) interconnect act as the top or the 4<sup>th</sup> level interconnect hierarchy. The global clock network provides low-skew clock signals for high fan-out nets. There are four dedicated global high-speed, low-skew clock signals with optional PLL (phase-lock-loop) control. While typically used to minimize skew, the PLLs can also be programmed as a clock multiplier to provide on chip frequency that is 2x - 3x the frequency of the external system clock, up to 200 MHz.

There are also dedicated interconnect resources for I/O Connectivity to facilitate routing to and from the I/O to support pin retention, incremental design changes and density migration: I/O direct connect resources, I/O shift connect and I/O long connect resources. The I/O direct connect resources provide high speed connectivity to logic blocks. The I/O-shift connect resources enable density migration up or down while maintaining pre-specified pin-outs in a given package. This additional routing flexibility is key to effective pin-retention for design changes and better timing retention for incremental changes. The I/O-Long Connect resources are capable of three-state operations and connect the I/Os to the adjacent Long-Connects on the core.

### 3.7 Other System Functions: High-Speed, Embedded Dual-Port SRAM Blocks, IO Blocks AND JTAG

The Vantis VF1 FPGAs consist of two-columns of dedicated, high-speed, cascable true dual-port embedded SRAM blocks. Each row of VGBs is associated with two embedded SRAM blocks. Each RAM block is configured to provide a 32-x 4 true dual-port RAM capability with completely independent READ and WRITE ports and separate READ and WRITE clocks. The READ port supports both high speed asynchronous and synchronous operation, while the WRITE port supports very high speed, synchronous operation only. With 5-ns access time, these high-speed dual-port memory blocks allow implementation of on-chip FIFOs running up to 200MHz. The embedded SRAM blocks have a dedicated address/control bus to provide easy expansion for either depth or width.

The dual-port configuration (one read/write port and one read port) allows an application to read from the read port while it is reading from or writing to the read/write port. This allows applications such as FIFOs and register stacks to run much faster and requires half as many memory bits as a single-port RAM would require to implement.

Figure 13 shows the organization of a 32x4 embedded SRAM block. One port of each SRAM block is a read/write port and the other is a read-only port. The read/write port consists of a write/read address input that may be stored in Read Address registers, or may bypass the register and go directly to the Read/Write port. For write operation, the write address is stored in the Read/Write port and write data is stored in the Write Data registers.

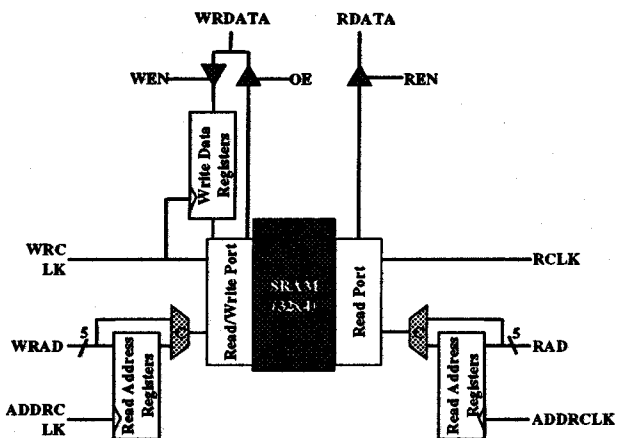
Memory read and write addresses come from dedicated SRAM address buses. There are five read address lines, five write address lines, and six control lines (including global clocks) connected to each 32x4 memory block. The SRAM address bus connects to VLI quad and long lines. Read/Write data from the read/write port connects to VLI long lines. Read data for the read output port and output enable lines connect to any VLI resources.

The VF1 embedded memory supports six single- and dual-port synchronous and asynchronous read and synchronous write operations. All single-port operations use the read/write port. The read-only port is used for dual-port operations. All write operations are synchronous. Read operations may be synchronous or asynchronous.

In dual-port operations, it is possible to read from the read port at the same time the read/write port is performing a read or a write. It is also possible to access the same address simultaneously. If the read/write port writes to an address at the same time the read port reads the address, the read port will read the old contents of the address until the next clock cycle, at which time the contents of the address will change to the new data. Modes supported include: Single port synchronous read/write; single-port synchronous write/asynchronous read; single port synchronous

write/asynchronous read, registered read address; dual-port synchronous read/write; dual-port synchronous write/asynchronous read; dual-port synchronous write/asynchronous read, registered read address.

Figure 13. VF1 Dual-Port SRAM Blocks



We chose this particular RAM block structure for three reasons:

- a) fast access speed
- b) potentially less waste, if not used
- c) scalability to higher densities without impacting the architecture

IOB regions lie on all four sides of the FPGA. Each programmable IOB includes a pad, input logic, and output logic. The input and output sections function separately from each other, sharing only the I/O pad and common Set/Reset logic. The Common Set/Reset logic is either the VF1 Global Set/Reset, or a local set/reset.

Figure 14 shows an I/O block. Each I/O Block is designed to provide high-speed on-chip/off-chip accessibility. Each I/O block has two separate and independently operable flip-flops - one dedicated for input and one for output. The one dedicated for input can also act as a transparent latch. The input flip-flop also has a delay element to provide Zero Hold Time. The VF1 FPGA devices also have IEEE standard 1149.1 compatible JTAG boundary scan capability with a separate port to be configured in-system via serial port from external configuration memories. Pins for JTAG are dedicated and not multiplexed for dual functions.

#### 4.SILICON/SOFTWARE DEVELOPMENT IN HARMONY

The VF1 FPGA architecture was developed in conjunction with the software tools and intended to be software centric from the very beginning. The architecture was driven strongly by the tools and features were implemented in the silicon to a) achieve an optimal balancing of logic and interconnect resources, b) make the tasks of mapping, place and route tools

easy, and c) to ensure that all the features were usable by the tools. A significant amount of energy was spent in architecting to balance the logic/interconnect resources to achieve speed, die size and ease the task of mapping and place/route tools. The architecture was driven with the goal of making it easy-to-use with a) an abundant and diverse style of routing resources,

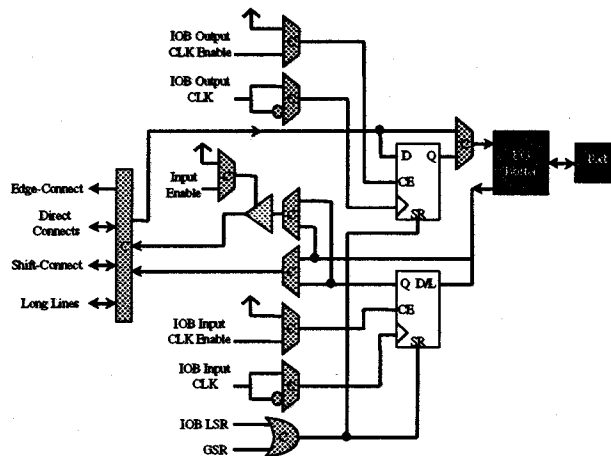


Figure 14. VF1 Input/Output Blocks

b) supporting incremental design, c) an ability to retain pins for design changes, and d) achieving a fast place and route time. The VF1 FPGA architecture was iterated many times during the development to ensure we had the optimal balancing of logic and interconnect structure to achieve our desired speed, cost and ease-of-use goals. Architecture iterations included the following:

1. Layout of the basic VGB (from 4-sided to L-shaped)
2. Optimization of Interconnect resources
3. Elimination of single length resources
4. Optimization of double, quad, octal and long lines
5. Implementation of shared drivers and connections to long lines
6. Implementation of Input Switch Matrices for VGBs (some iterations)
7. Implementation of switch boxes (some iterations)
8. Dual-port SRAM blocks interconnect (some iterations)
9. Connectivity of dual-port SRAM blocks to interconnect resources (some iterations)
10. Interconnect of IO lines for pin-locking and density shifting
11. Control functions inside VGBs
12. Interconnect of IOs to long lines
13. Implementation of CCEs and CSEs

#### 5. VF1 PERFORMANCE OBJECTIVES

High performance was our number one priority. Users measure system performance in many ways, but there are four factors that are commonly held as significant - combinatorial time on & off chip for control signals, data path access to I/O registers & memory blocks and arithmetic performance of the carry



chain. Each of these was further broken down into a total of 9 separate cases for performance specifications. Table 2.

Table 2. Nine Cases for Performance Objectives

#	Description	Objective	Actual	Units
1	Fast Comb Path	7.5	7.3	ns
2	Diagonal Comb path	15	12.7	ns
3	3 level Comb logic	15	14.7	ns
4	3 level, IO reg Fmax	75	79	MHz
5	Shift Reg – Long Line	175	177	MHz
6	16-bit Adder	10	9.2	ns
7	Fmax External * (System Speed)	125	139	MHz
8	SRAM Read	3.0	2.8	ns
9	SRAM Write & Read	4.0	3.9	Ns

\* (1/TSU + TCO of IO register)

## 5.1 Mapping Optimization For Area and Performance

Logic mapping can be tuned to give best speed at the expense of circuit area or the reverse. Users want to be able to control the depth of critical paths through the logic and optimize area for those paths not on the critical paths. Our optimizer examines the user logic for paths that have significant depth and fanout, marks those for special processing and applies speed or area optimization to the circuit as directed by the menu profiles. The critical paths are evaluated to estimate their compliance with the performance objectives from the timing constraint GUI and processed accordingly. Longer paths are rearranged where feasible to reduce the depth. Fifteen customer benchmark circuits were processed to examine the result of tuning for either speed or area and the results are presented in Table 3 below.

Table 3. Optimized for AREA

Bench Mark #	Total CBBs	3-LUTs	4-LUTs	Registers (FFs)
BM-1	819	683	130	520
BM-2	544	259	284	220
BM-3	1045	588	366	608
BM-4	745	362	270	465
BM-5	636	478	211	228
BM-6	1518	774	474	878
BM-7	1388	1431	145	1289
BM-8	707	500	184	459
BM-9	1058	1010	283	703
BM-10	780	430	352	341
BM-11	643	394	254	252

BM-12	284	99	32	98
BM-13	959	646	159	575
BM-14	860	301	577	198
BM-15	574	446	223	239

Table 4 shows the data optimized for speed.

Table 4. Optimized for speed

Bench Mark #	Total CBBs	3-LUTs	4-LUTs	Registers (FFs)
BM-1	825	745	166	520
BM-2	580	274	268	220
BM-3	1087	574	386	608
BM-4	739	405	315	465
BM-5	651	473	222	228
BM-6	1566	763	434	878
BM-7	1454	1495	145	1289
BM-8	727	504	180	459
BM-9	1075	1017	259	703
BM-10	843	430	319	341
BM-11	705	379	327	252
BM-12	284	99	32	98
BM-13	950	687	197	575
BM-14	876	292	608	198
BM-15	638	386	211	239

Table 5: Mapper, Placer, Router run times

Design	Mapper Time	Placer Run Time	Router Run Time
BM-1	0:02:40	0:07:58	0:02:16
BM-2	0:03:32	0:08:47	0:01:33
BM-3	0:01:52	0:13:54	0:11:01
BM-4	0:01:21	0:05:34	0:00:23
BM-5	0:01:04	0:09:29	0:04:17
BM-6	0:03:01	0:16:56	0:01:17
BM-7	0:01:30	0:10:56	0:10:22
BM-8	0:02:07	0:10:55	0:13:55
BM-9	0:01:31	0:09:30	0:34:41
BM-10	0:01:28	0:15:39	0:14:00
BM-11	0:02:06	0:15:10	0:01:46
BM-12	0:03:33	0:13:58	0:01:57
BM-13	0:13:20	0:18:22	0:06:23
BM-14	0:02:11	0:13:52	0:10:52

## 6. CONCLUSION

The Vantis VF1 architecture has been driven by a significant number of architectural innovations. At a higher level, the fundamental and unique innovations in the VF1 family of FPGAs are in the structure of each of the basic components: Variable-Grain-Block, hierarchical Variable-Length-Interconnect and flexible, high-speed dual-port embedded SRAMs. Variable-Grain-Architecture and hierarchical, Variable-Length-Interconnect is key for achieving optimal silicon efficiency (high utilization), speed, low cost and wide

gating functions. These architectural concepts were developed in close conjunction with the mapping and place/route tools to ease development of timing driven synthesis, optimization, mapping and timing driven place and route tools. The tools fully exploit the capabilities of the architecture and provide the system designers with optimal speed and high quality-of-results. With a simple, symmetrical, and homogeneous Variable-Grain-Architecture and Variable-Length-Interconnect hierarchy, Vantis' new VF1 FPGA family is a novel attempt to raise mainstream FPGA architectures to new heights.

## 7. ACKNOWLEDGEMENTS

We are thankful to many Vantis employees and independent contractors who helped us with the development of this unique architecture. Though we cannot name all of them here, we would like to explicitly acknowledge the contributions made by our key folks from the software development, engineering and marketing teams.

## References

- [1] Vantis VF1 Data Sheet 1998
- [2] Rose, J., Francis, R.J., Lewis, D., and Chow, P. Architecture of Programmable Gate Arrays: The Effect of Logic Block Functionality on Area Efficiency. IEEE journal of Solid State Circuits, Vol. 25, No. 5, October 1990, pp. 1217-1225
- [3] Brown, S., Khella, M., and Varnesic, Z. Minimizing FPGA Interconnect Delays. IEEE Design and Test of Computers Winter (1996) 16-23.
- [4] Carter, W., Duong, K., Freman, R., Hsieh, H., Ja, J. Y., Mahoney, J. E., Ngo, N. T., and Sac, S. L. A User Programmable Reconfigurable Gate Array. Proc 1986 Custom Integrated Circuits Conference, May 1986, pp. 233-235.
- [5] Hill, D., and Woo, N. S. The Benefits of Flexibility in Look-up Table FPGAs. In FPGAs, W. Moore and W. Luk Eds., Abingdon 1991. Edited from the Oxford 1991 International Workshop on Field Programmable Logic and Applications, pp. 127-136.
- [6] Rose, J.S., Francis, R.J., Chow, P and Lewis, D. The Effect of Logic block Complexity on Area of Programmable Gate Arrays. Proc 1989 Custom Integrated Circuits Conference, May 1989, pp. 5.3.1 - 5.3.5
- [7] Rose, J., and Brown, S. The Effect of Switch box Flexibility on Routability of Field Programmable Gate Arrays. Proc 1990 Custom Integrated Circuits Conference, pp. 27.5.1 - 27.5.4, May 1990.
- [8] Betz, V., and Rose, J. Directional Bias and Non-uniformity in FPGA global Routing Architectures. Proc. ICCD'96 (1996) 652-659.
- [9] Rose, J., and Brown, S. Flexibility of Interconnection Structures in Field Programmable Gate Arrays. IEEE journal of Solid State Circuits, Vol 26, No. 3, pp. 277-282, March 1991.

---

Authors' e-mail addresses:

[Om.agrawal@vantis.com](mailto:Om.agrawal@vantis.com)  
[herman.chang@vantis.com](mailto:herman.chang@vantis.com)  
[brad.sharpe-geisler@vantis.com](mailto:brad.sharpe-geisler@vantis.com)  
[nick.schmitz@vantis.com](mailto:nick.schmitz@vantis.com)  
[bai.nguyen@vantis.com](mailto:bai.nguyen@vantis.com)  
[jack.wong@vantis.com](mailto:jack.wong@vantis.com)  
[giap.tran@vantis.com](mailto:giap.tran@vantis.com)  
[fabiano.fontana@vantis.com](mailto:fabiano.fontana@vantis.com)  
[bill.harding@vantis.com](mailto:bill.harding@vantis.com)