

TFT Panel Support in the Vybrid Microcontroller Family

How to Choose a Suitable Panel and Configure Vybrid

by: Steve McAslan

Contents

1 Introduction

Each Animation and Compositing Engine (2D-ACE) module on the Vybrid MCU family is capable of driving a TFT LCD panel of different sizes. The upper limit on the size of each panel depends on a range of factors including 2D-ACE clock speed, platform clock speed, memory interface speed, graphic data encoding, and number of graphic layers active. This application note explains how each of these factors contributes to the limit and provides guidelines for the maximum panel size for various MCU configurations.

2 The 2D-ACE module

The 2D-ACE performs two primary tasks on the Vybrid architecture. Firstly it fetches, composites, and blends graphics dynamically from on- or off-chip memory and secondly, it provides the final graphic content to a TFT LCD panel connected to the Vybrid GPIO pads.

The 2D-ACE itself provides fundamental and practical upper limits to the panel size. There are no practical lower limits because the 2D-ACE supports all known available small panel sizes.

Since the 2D-ACE does not rely on a frame buffer, its performance is dependent on the availability and speed of the memory in or attached to the Vybrid. A graphic is fetched

1	Introduction.....	1
2	The 2D-ACE module.....	1
2.1	Hardware upper limits to panel size.....	2
2.2	Impact of memory bandwidth.....	2
2.3	Memory optimisation techniques.....	3
2.4	Other memory types.....	5
3	Suggested panel sizes for different use cases.....	6
3.1	How to interpret the calculations.....	6
3.2	Single panel applications.....	7
3.3	Dual panel applications.....	9
4	Conclusion.....	10

from memory when it is associated with a 2D-ACE layer. Each 2D-ACE allows any combination of up to 64 layers to be displayed at a time and up to 6 of these layers to be blended together at each individual output pixel value. Therefore, the peak demand from the 2D-ACE is 6 graphics at the highest quality encoding of 32 bits-per-pixel (BPP). Each of those source graphics can come from any memory location; so the peak demand on a specific memory interface would also be 6 x 32 BPP.

2.1 Hardware upper limits to panel size

Configuration of the 2D-ACE panel interface is described in detail in AN4647: Configuring and Using the 2D-ACE on Vybrid Microcontrollers, available on freescale.com. This configuration allows a number of upper values which define the specification limits on panel size. However, in practice, most panels will be smaller than the size that the specification allows.

The absolute maximum size of a panel that can be connected to the Vybrid is 2048 x 2048 pixels as defined by the 2D-ACE DISP_SIZE register, but this is not a practical size because for a normal 60 Hz refresh rate, the panel would require a pixel clock value of greater than 300 MHz which is higher than the clock architecture can supply.

The pixel clock in the 2D-ACE is derived from either the PLL1 PFD2 or PLL3 clock sources on the MCU, which gives a usual upper limit of 480 MHz. The 2D-ACE divides this by 2 at a minimum to give a theoretical pixel clock limit of 240 MHz, or a panel size of approximately 1920 x 1440 pixels.

2.1.1 Pixel clock rates and content of panel

This pixel output clock rate is limited in usability by the challenge of high speed circuit design (especially EMC). It is also limited by the rate at which pixels can be fetched since the 2D-ACE must fetch and blend each of the pixels before outputting them to the panel.

All incoming pixels must pass across the internal NIC301 crossbar which typically operates at 133 MHz and allows 64-bit transfers. Therefore the 2D-ACE can fetch each 32 BPP layer at a peak rate of 266 Mpixels/s (two pixels per clock cycle). This data rate must be shared among all layers that are visible on the panel so the peak rate is actually 266 Mpixels per second per layer.

This means that the rate at which layers can be fetched is reduced by the number of overlapping layers in the design. For example a frame which uses 4 x 32 BPP overlapping layers means that each of the layers can have one pixel fetched at an average incoming rate of 66.5 Mpixels/s. Since all incoming pixels must be fetched before a single output pixel can be displayed this in turn means that the maximum rate of the output pixel clock is also 66.5 MHz.

In practice, a pixel clock rate closer to 60 MHz is more reasonable because it is more straightforward to create from the source clock frequency. This pixel clock value would support a panel of 1024 x 768 but although this may be a practical hardware limit, it is likely that memory and display specifications will reduce this further.

2.2 Impact of memory bandwidth

The Vybrid MCU supports various memory types of which the highest bandwidth performance is found on SDRAM, on-chip RAM (OCRAM), and serial flash (QuadSPI). Other memory sources are also available but do not have the speed or the memory-mapping required by the 2D-ACE (see [Other memory types](#)).

It is necessary to understand the effect of memory bandwidth on the 2D-ACE configuration to see how it affects usable panel sizes.

2.2.1 Impact of memory availability

Some members of the Vybrid family include an SDRAM interface which allows up to 2 GB to be connected to the system while others have a limit of 1.5 MB of on-chip RAM (OCRAM). It is likely that the size of available RAM will be an upper limit on the graphic content available and indirectly on the size of the panel that can be connected. This is because the 2D-ACE fetches graphics stored in memory to display on the panel and the less memory there is, the less sophisticated or (more likely) the smaller the panel may be.

For this reason, a Vybrid application with no external SDRAM limits the maximum panel size because there is simply not enough memory to store the graphics required for the application and the external memory available is too slow.

2.2.2 Influence of memory speed

As previously discussed in [Pixel clock rates and content of panel](#), the 2D-ACE can fetch pixels at a peak rate of 266 Mpixels/s for a 32 BPP graphic. This is equivalent to a raw transfer rate of just over 1 GB/s. This means that the memory system must also be able to provide a combined operating rate as fast to allow the 2D-ACE to operate at its maximum performance and so support the largest panel possible.

- **SDRAM:** The maximum SDRAM data rate is 1600 MB/s which is based on a 400 MHz bus, DDR3 memory, and a 16-bit interface. At this rate, it would be possible to supply approximately 6 x 32 BPP layers each frame to the 2D-ACE for a 60 MHz pixel clock where only graphics are stored in the SDRAM and layer access is at its most efficient level. However, the maximum transit bandwidth is only 1 GB/s, so a full 6-layer 32 BPP blend is not possible at a 60 MHz pixel clock on a single panel. Reducing the pixel clock to 40 MHz allows the six layers to be fetched and blended (40 x 6 x 4 = 960 MB/s). The peak SDRAM rate allows a reasonable assumption that the 1 GB/s peak bandwidth could be comfortably achieved while allowing for random data access into the memory. Alternatively, the memory could supply graphic content to more than one 2D-ACE simultaneously.
- **OCRAM:** The OCRAM also operates at the platform clock of 133 MHz and provides a 64-bit interface. Therefore its bandwidth is also approximately 1 GB/s or 6 x 32 BPP layers each frame (at 40 MHz). The limitation with this memory is that it cannot be increased in size.
- **Serial flash:** The slowest and most cost-effective of these memories is the serial flash attached to the QuadSPI modules. Two independent interfaces are available with both allowing 4-bit single and double data rate operation. QuadSPI0 also allows two serial flashes to operate in parallel giving an effective 8-bit transfer at single and double data rate. A typical operating specification for the fastest option (parallel DDR) is 66 MHz which yields a peak memory bandwidth of 132 MB/s. In practice, the effective bandwidth is less because of the command overhead of each memory read but a sustained rate of more than 100 MB/s is considered a realistic expectation. This supports around 1 x 16 BPP layer each frame for a 40 MHz pixel clock. Reducing the pixel clock speed obviously increases the number of layers that can be fetched.

The serial flash is also the memory most likely to be used when the application executes “in-place” (XiP). This is the configuration where the CPU fetches its program directly from the serial flash and relies on its instruction cache to provide a suitable level performance. In this case, the serial flash bandwidth is shared between the application and the 2D-ACE and the use of the flash needs to be carefully planned. In this case, it may be possible that best performance is achieved when graphics are copied from the flash into OCRAM before use.

Given the above calculations, it is apparent that even with a fast DDR3 SDRAM, the largest panel supported with full color encoding is WVGA (800 x 480, c. 33 MHz). The use of external flash will further reduce this capacity if direct fetch from memory is required.

However, there are many techniques that can be employed to minimise the memory bandwidth and capacity required and so allow more layers to be blended on a larger panel.

2.3 Memory optimisation techniques

There are three benefits to optimising the requirements for graphic memory:

- Graphics which use fewer BPP for each pixel can make more efficient use of bandwidth.

the 2D-ACE module

- Graphics which take less space in memory save cost by reducing the overall memory footprint.
- Buffering graphics from slower and cheaper memory into RAM allows better performance.

It is possible to reduce the footprint of the graphics by using more advanced features of the Vybrid. The memory space saving that is achievable depends heavily on the specific graphics and may come at the expense of time to display the image.

The Vybrid includes a number of features that make it possible to reduce the memory footprint required for a given application. These include:

- Optimising the original graphic encoding and using tile mode of the 2D-ACE
- Compressing the graphic data using Run Length Encoding (RLE) and using either the 2D-ACE or independent RLE module to unpack the image when required
- Generate graphics as required using the OpenVG Graphics Processing Unit (GPU)
- Using the DMA module to rotate and mirror graphics instead of storing multiple images

Each of these features is explained in the following subsections.

2.3.1 Optimising graphic encoding

The 2D-ACE can display graphics in 13 different encodings without loss of quality in the blend.

NOTE

The quality of the image is always determined by the source graphic format.

In most cases, the use of 32 BPP ARGB images is excessive. There are various factors which make this the case:

- Most Vybrid applications will be on relatively small panels with limited resolution and limited brightness.
- The range of colors and gradients in most images can be easily represented by a smaller color gamut.
- Many applications only require a 16- or 18-bit panel and so, full 24-bit RGB color is wasteful unless the temporal dithering function is enabled.

Given these factors, it is normal for many applications to use 16-bit formats as a standard. This immediately doubles the effective number of layers which can be stored in memory and fetched with a given bandwidth. For some images, a look-up-table format is suitable; available formats include 8 BPP which gives a 256 color palette. More likely is that applications can use the 8 BPP Transparency format which gives full-quality alpha anti-aliasing for a given image color. Either of these two latter options give 4x bandwidth and footprint benefits. All of these options can be used in combination, so it is possible to use a single high-quality 32 BPP image along with 8 BPP or 16 BPP formats on a single frame.

When analysing panel options, a combination of format usages is assumed.

2.3.2 Compressing the graphics

The memory footprint may be reduced by compressing the source graphic using RLE. There are two options to decompressing the image.

- Standalone module: The standalone module does not allow any net saving of memory bandwidth since the decompressed image must ultimately be stored into memory although the decoding can be done during frame refresh porches.
- 2D-ACE module: This module decompresses the RLE data internally and so it does save some bandwidth as well as memory space; however, there are more restrictions on this approach (limited format support, partial extraction not possible and limited to one layer at a time)

The usefulness of RLE compression is entirely dependent on the source image but since it is lossless, it is a simple exercise to examine the benefit of the approach by compressing the image.

2.3.3 Using the OpenVG GPU

The OpenVG GPU allows the creation and manipulation of images and so removes the need to store them in external flash and subsequently copy them to RAM. There is no direct saving of memory bandwidth since the images must always be stored in RAM before use; however, only the images required for the current or upcoming frames need to be stored in RAM.

The OpenVG engine can draw and translate images using vector data or source raster images and in both cases, significant memory savings can be made by manipulating the images as they are needed for the display.

2.3.4 Manipulate images using the DMA

The Vybrid includes an advanced Direct Memory Access module called the eDMA which allows manipulation of data ordering while a copy is taking place. This DMA uses independent inner and outer loops to calculate how to choose the next source and destination address and this allows operations such as rotations and mirroring and integer re-sizing. As a result, this module is commonly used to provide memory saving for objects that rotate or have a consistent pattern around a centre. In these cases, it is possible to store only a small portion of the image or pattern— typically 45°—and reconstruct the remainder of the graphic by using the eDMA translations.

2.4 Other memory types

The Vybrid supports other memories beyond those described here but they have more limitations than the memories already described. They may be useful and cost-effective interfaces in the system, but do not directly influence the configuration of the panel for most applications.

Some of the other memory types are described briefly as follows.

2.4.1 NAND flash controller

The NAND Flash Controller (NFC) provides an industry-standard interface to NAND flash modules. It accesses the memory contents by a series of commands which copy the the flash contents into RAM and therefore, is not available for direct access by the 2D-ACE.

In effect, the NFC requires the use of SDRAM or RAM as a buffer for graphics and is therefore of no direct influence on the performance of the 2D-ACE in the system.

2.4.2 Flexbus controller

The Flexbus controller is a multi-function module that supports connection to several different types of memories and other peripherals. It operates as a memory-mapped module, so it is possible for the 2D-ACE to directly read from external memories such as NOR flash.

The Flexbus operation is limited to a maximum of 57 MHz and requires a minimum of an 8-bit bus and around four other control signals. A multiplexed but parallel address bus would normally add at least another 10 to 12 bits. The use of a large number of pins with a limited operating speed means that that this interface is not the most cost-effective solution for providing graphics for the 2D-ACE.

2.4.3 USB controller

The USB controller implements a full high-speed USB interface so that it can stream data at a peak rate of 60 MB/s. It can access external memory via a series of commands and data transfers into RAM and therefore is not available for direct access by the 2D-ACE.

The USB requires the use of SDRAM or RAM as a buffer for graphics and is therefore of no direct influence on the performance of the 2D-ACE in the system.

2.4.4 SDHC Controller

The Secure Digital Host Controller (SDHC) provides an industry-standard interface to Secure Digital (SD) and Multimedia Cards (MMC). It accesses the memory contents by a series of commands which copy the the flash contents into RAM and therefore is not available for direct access by the 2D-ACE.

The SDHC requires the use of SDRAM or RAM as a buffer for graphics and is therefore of no direct influence on the performance of the 2D-ACE in the system.

3 Suggested panel sizes for different use cases

In this section, a number of different application use cases are considered and the Vybrid configuration for each is suggested.

NOTE

Since the recommendations rely on typical requirements, it is possible that the experienced designer or expert software engineer may be able to extract a higher performance from the system than shown here.

3.1 How to interpret the calculations

All calculations are made using graphics that are of the same size as the panel. Although this is not a typical use case, it gives an insight into the capability of the system and the users can use this as a reference point for their own graphics. For example, if the calculation shows that a single 16 BPP graphic may be stored in OCRAM, then it will be possible to store the sum of the same surface area in multiple smaller graphics and graphics with smaller encoding size.

Consider the case where 2.5 16 BPP graphics can be fetched from SRAM. In this case, it would be possible to fetch 5 full-sized 8 BPP graphics or 5 16 BPP graphics where each is one half of the panel size.

Note that in all cases, the 2D-ACE can never use more than six layers simultaneously and therefore any calculation which supports more than this at 32 BPP indicates that any combination of graphics is possible for that panel size.

The tables shown in the following subsections contain a bandwidth calculation and a memory size calculation. For internal OCRAM, the lack of storage can be a greater limit than the memory bandwidth.

Each of the tables contains up to 9 entries from the values described below:

- Width—the width of the panel chosen
- Height—the height of the panel chosen
- Nominal clock—the typical pixel clock frequency of a panel of this size
- 32 BPP OCRAM capacity—the maximum number of panel-sized 32 BPP layers that can be stored using 1 MB of OCRAM
- 32 BPP OCRAM % load—the fraction of memory bandwidth used if the maximum number of 32 BPP graphics (6) are fetched from OCRAM to blend a single output pixel. In other words, the memory bandwidth that the 2D-ACE consumes when fetching 6 layers from OCCRAM for this panel (irrespective of the size of the graphics)
- 32 BPP QuadSPI capacity – the maximum number of panel-sized 32 BPP layers that can be fetched from an external serial flash

- 32 BPP SDRAM % load – the fraction of memory bandwidth used if the maximum number of 32 BPP graphics (6) are fetched from SDRAM to blend a single output pixel. In other words the memory bandwidth that the 2D-ACE consumes when fetching 6 layers from SDRAM for this panel (irrespective of the size of the graphics)
- 16 BPP OCRAM capacity—the maximum number of panel-sized 32 BPP layers that can be stored using 1 MB of OCRAM
- 16 BPP OCRAM % load—the fraction of memory bandwidth used if the maximum number of 16 BPP graphics (6) are fetched from OCRAM to blend a single output pixel. In other words, the memory bandwidth that the 2D-ACE consumes when fetching 6 layers from OCCRAM for this panel (irrespective of the size of the graphics)
- 16 BPP QuadSPI capacity—the maximum number of panel-sized 16 BPP layers that can be fetched from an external serial flash
- 16 BPP SDRAM % load—the fraction of memory bandwidth used if the maximum number of 16 BPP graphics (6) are fetched from SDRAM to blend a single output pixel. In other words, the memory bandwidth that the 2D-ACE consumes when fetching 6 layers from SDRAM for this panel (irrespective of the size of the graphics)

An entry of “-“ in any table means that the value is unobtainable by the Vybrid. Recommended screen sizes are highlighted.

Note that in all cases, the peak bandwidth of the memory is considered as 100% usage. In practice, the average memory bandwidth is lower than this; however the panel does not consume data at a constant rate due to the horizontal and vertical porches. In addition, there are several buffers in the system, panel-sized graphics are rarely used, and the 2D-ACE performs a pre-emptive fetch; so in practical applications, the true usable bandwidth is very hard to predict. For these reasons, the usable bandwidth is always stated against peak performance. A user should note that it is unlikely that bandwidth above c. 90% should be considered achievable in most applications.

3.2 Single panel applications

The Vybrid has the ability to use either of the two 2D-ACE modules in a single panel application, but in most cases, the 2D-ACE 0 (DCU0) is the obvious choice. This is because it is designed to use GPIO pins that are close together and which share fewest other functions that are not available on other pins. These examples assume the use of DCU0.

3.2.1 Small (144-pin) LQFP applications

The LQFP packages do not provide support for either SDRAM or Flexbus, so memory for application and graphics is limited to the OCRAM and one or two QuadSPI interfaces. For this use case, assume that only a single QuadSPI serial flash is available but that it can be run as DDR at 66 MHz. Also, assume a platform clock of 133 MHz.

Width	Height	Nominal clock (MHz)	32 BPP			16 BPP		
			OCRAM max	OCRAM % load	QuadSPI capacity	OCRAM max	OCRAM % load	QuadSPI capacity
320	240	5	3.25	12	3.32	6.51	6	6.5
480	272	9	1.91	21.6	1.84	3.82	10.8	3.61
640	480	25	0.81	60	-	1.62	30	1.3
800	480	33	0.65	79.2	-	1.3	39.6	-
800	600	40	0.52	96	-	1.04	48	-
1024	768	65	0.31	-	-	0.63	78	-

144-pin LQFP Applications:

Suggested panel sizes for different use cases

From the table given above, it is clear that panels above 480 x 272 (WQVGA) place an unreasonable demand on the storage capability of OGRAM (c. 1 layer) and the bandwidth capability of the serial flash (max 1 16 BPP graphic). A WQVGA panel can be considered to be supported well by this system since it allows storage of large 32 BPP graphics and multiple 16 BPP graphics in both OGRAM and serial flash.

3.2.2 Large (176-pin) LQFP applications

The LQFP packages do not provide support for either SDRAM or Flexbus, so memory for application and graphics is limited to the OGRAM and one or two QuadSPI interfaces. For this use case, assume that two QuadSPI serial flashes are available and configured in parallel and that both can be run as DDR at 66 MHz. Also, assume a platform clock of 133 MHz.

Width	Height	Nominal clock (MHz)	32 BPP			16 BPP		
			OCRAM max	OCRAM % load	QuadSPI capacity	OCRAM max	OCRAM % load	QuadSPI capacity
320	240	5	3.25	12	6	6.51	6	13
480	272	9	1.91	21.6	3	3.82	10.8	7
640	480	25	0.81	60	1	1.62	30	2
800	480	33	0.65	79.2	1	1.3	39.6	2
800	600	40	0.52	96	-	1.04	48	1
1024	768	65	0.31	-	-	0.63	78	1

176-pin LQFP applications

From the table given above, it is clear that panels above 480 x 272 (WQVGA) place an unreasonable demand on the storage capability of OGRAM (c. 1 layer) and the bandwidth capability of the serial flash (max 1 16 BPP graphic). A WQVGA panel can be considered to be supported well by this system since it allows storage of large 32 BPP graphics and multiple 16 BPP graphics in both OGRAM and serial flash.

3.2.3 BGA applications

The BGA package supports SDRAM for application and graphics and one or two QuadSPI interfaces. For this use case, assume that there is a large DDR3 RAM at 400 MHz and two QuadSPI serial flashes configured in parallel and running as DDR at 66 MHz. Also, assume a platform clock of 133 MHz. This example assumes that the SDRAM is the primary storage memory from graphics.

Width	Height	Nominal clock (MHz)	32 BPP			16 BPP		
			SDRAM max	SDRAM % load	QuadSPI capacity	SDRAM max	SDRAM % load	QuadSPI capacity
320	240	5	Unlimited	7.5	6	Unlimited	3.75	13
480	272	9	Unlimited	13.5	3	Unlimited	6.75	7
640	480	25	Unlimited	37.5	1	Unlimited	18.75	2
800	480	33	Unlimited	49.5	1	Unlimited	24.75	2
800	600	40	Unlimited	60	-	Unlimited	30	1
1024	768	65	Unlimited	97.5	-	Unlimited	48.75	1

364-pin BGA applications

From the table given above, it is clear that panels up to 800 x 600 (SVGA) are supported in this system. In practice, the more common WVGA (800 x 480) is a more common form factor. A WVGA panel can be considered to be supported well by this system since the bandwidth of the SDRAM allows several graphics. Larger panels may be supported depending on the acceptable maximum loading of the external memory.

3.3 Dual panel applications

In these examples, assume an equal graphic demand on two identical panels. This may not be true in all cases, but gives an insight into the overall performance limits of the systems.

NOTE

The memory bandwidth of each panel is independent. This means that each panel can consume up to 1 GB/s independently from the memory sources.

3.3.1 Large (176-pin) LQFP applications

The LQFP packages do not provide support for either SDRAM or Flexbus, so memory for application and graphics is limited to the OCRAM and one or two QuadSPI interfaces. For this use case, assume that two QuadSPI serial flashes are available and configured in parallel and that both can be run as DDR at 66 MHz. Also, assume a platform clock of 133 MHz.

Width	Height	Nominal clock (MHz)	32 BPP			16 BPP		
			OCRAM max	OCRAM % load	QuadSPI capacity	OCRAM max	OCRAM % load	QuadSPI capacity
320	240	5	3.25	24	3	6.51	12	6
480	272	9	1.91	43.2	1	3.82	21.6	3
640	480	25	0.81	-	-	1.62	60	1
800	480	33	0.65	-	-	1.3	79.2	1
800	600	40	0.52	-	-	1.04	96	-
1024	768	65	0.31	-	-	0.63	-	-

176-pin LQFP applications

From the table given above, it is clear that panels above 480 x 272 (WQVGA) place an unreasonable demand on the bandwidth capability of the serial flash (max 1 16 BPP graphic). WQVGA panels can be considered to be supported reasonably well by this system but 16 BPP graphics must be considered as a maximum data size per pixel.

3.3.2 BGA applications

The BGA package supports SDRAM for application and graphics and one or two QuadSPI interfaces. For this use case, assume that there is a large DDR3 RAM at 400 MHz and two QuadSPI serial flashes configured in parallel and running as DDR at 66 MHz. Also, assume a platform clock of 133 MHz. This example assumes that the SDRAM is the primary storage memory from graphics.

Conclusion

Width	Height	Nominal clock (MHz)	32 BPP			16 BPP		
			SDRAM max	SDRAM % load	QuadSPI capacity	SDRAM max	SDRAM % load	QuadSPI capacity
320	240	5	Unlimited	15	3	Unlimited	7.5	6
480	272	9	Unlimited	27	1	Unlimited	13.5	3
640	480	25	Unlimited	75	-	Unlimited	37.5	1
800	480	33	Unlimited	99	-	Unlimited	49.5	1
800	600	40	Unlimited	-	-	Unlimited	60	-
1024	768	65	Unlimited	-	-	Unlimited	97.5	-

364-pin BGA applications

From the table given above, it is clear that panels up to 800 x 480 (WVGA) are supported in this system. A WVGA panel can be considered to be supported well by this system since the bandwidth of the SDRAM allows up to 6 32 BPP graphics on each panel. Larger panels may be supported depending on the acceptable maximum loading of the external memory.

NOTE

The usability of the serial flash in this case is likely very restricted.

4 Conclusion

The Vybrid platform comfortably supports panels of various sizes; however the memory capacity and bandwidth of the system impose upper limits on the practical size of these panels.

For the LQFP package, panels of no larger than WQVGA (480 x 272) are comfortably supported. For the BGA package, up to two WVGA (800 x 480) panels can be connected assuming a suitable DDR3 memory is also connected.

How to Reach Us:

Home Page:

www.freescale.com

Web Support:

<http://www.freescale.com/support>

USA/Europe or Locations Not Listed:

Freescale Semiconductor
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
+1-800-521-6274 or +1-480-768-2130
www.freescale.com/support

Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
www.freescale.com/support

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku,
Tokyo 153-0064
Japan
0120 191014 or +81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor China Ltd.
Exchange Building 23F
No. 118 Jianguo Road
Chaoyang District
Beijing 100022
China
+86 10 5879 8000
support.asia@freescale.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductors products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation, or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any liability, including without limitation consequential or incidental damages. "Typical" parameters that may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals", must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claims alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

RoHS-compliant and/or Pb-free versions of Freescale products have the functionality and electrical characteristics as their non-RoHS-complaint and/or non-Pb-free counterparts. For further information, see <http://www.freescale.com> or contact your Freescale sales representative.

For information on Freescale's Environmental Products program, go to <http://www.freescale.com/epp>.

Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners.

© 2013 Freescale Semiconductor, Inc.