New Monolithic A/D Technology

- Smaller geometry, lower core voltages and power dissipation
- Much higher sample rates and bit accuracy
- Wideband input circuitry optimized for direct IF sampling
  - IF (intermediate frequency) signals are usually greater than Fs
  - Differential, transformer coupled inputs minimize noise
- High Performance Integrated S&H (sample-and-hold)
  - Higher immunity to clock waveform symmetry and level
- Improved multi-stage flash conversion techniques
- Digital sample code generation and error correction
  - Devices can be calibrated and trimmed during production
- Improved thermal tracking of DC offset, gain, and linearity
- Improved power supply noise rejection and immunity

Because of all of these market segments, wideband A/D converters have made some tremendous advances in the last five years.

This is partly due to silicon process improvements, but that's not all.

Because many applications require direct sampling of IF signals well above 100 MHz, new wideband input stages were developed.

One of the most important advances is the sample-and-hold (or track-and-hold) at the front end.

Just as important, are new sample clock interfaces and drivers.

At these speeds, you need state-of-the-art flash and multi-stage flash conversion techniques.

New techniques in digital error code correction and thermal compensation circuitry help eliminate errors in bit accuracy, linearity and gain.

Lastly, these new devices are more immune to power supply and system noise.

Monolithic A/D Converters for Fs ≥ 200 MHz and bits ≥ 8

Manufacturer	Part No.	Sample Rate	Chans	No. Bits	Input BW
Atmel	AT84AS008	2200 MHz	1	10	3300 MHz
Atmel	TS83102G0B	2000 MHz	1	10	3300 MHz
Maxim	MAX108	1500 MHz	1	8	2200 MHz
National	ADC081000	1000 MHz	1	8	1700 MHz
National	ADC08D1000	1000 MHz	2	8	1700 MHz
Atmel	JTS8388B	1000 MHz	1	8	2000 MHz
Maxim	MAX104	1000 MHz	1	8	2200 MHz
Atmel	AT84AD001B	1000 MHz	2	8	1500 MHz
Maxim	MAX106	600 MHz	1	8	2200 MHz
Atmel	AT84AD004B	500 MHz	2	8	1000 MHz
Maxim	MAX101A	500 MHz	1	8	1200 MHz
Atmel	TS8308500	500 MHz	1	8	1300 MHz
Maxim	MAX1121	250 MHz	1	8	600 MHz
Analog Dev	AD9480	250 MHz	1	8	750 MHz
TelASIC	TS1411	250 MHz	1	14	1000 MHz
Analog Dev	AD9430	215 MHz	1	12	700 MHz
Analog Dev	AD9410	210 MHz	1	10	500 MHz
Analog Dev	AD9054	200 MHz	1	8	350 MHz

Here's a long list of some of the commercially available, monolithic A/D converters with sampling rates of at least 200 MHz and resolution of at least 8 bits.

These are all candidates for board level products for embedded systems.

Notice I've listed the input bandwidth at the right, just to highlight the importance of direct IF sampling applications, also call undersampling.

In the next section, we'll discuss more about the principles and rules of sampling.

Section 2. Sampling and Filtering Techniques

Direct Baseband RF Signal Acquisition

- Antenna signals are usually in the microvolt range
- RF amplifier boosts signal to full scale input voltage of the A/D - usually 0 to +10 dBm
- RF amplifier often includes a tuned bandpass filter centered on the signal of interest
- No analog frequency translation before the A/D
- Appropriate for HF signal frequencies (3 - 30 MHz)

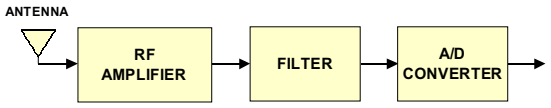

Most receiver systems start with a signal originating from an antenna, and that signal is often at the microvolt level, so it must first be amplified by an RF amplifier stage.

The amplifier is usually a tuned RF circuit which only passes the frequency band of interest, providing signal gain **within** that band and rejecting noise and unwanted signals in adjacent frequency bands.

If the RF input signal is at a low enough frequency, it can be digitized directly by an A/D converter, and no analog translation is necessary.

For example, you can usually perform direct baseband sampling on HF signals with no translation required, since the frequency content is below 30 MHz.

Analog RF Frequency Translation

- Analog Translation to Baseband

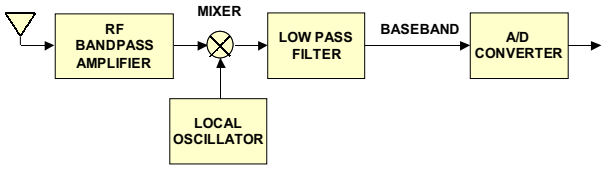

- Analog Translation to IF (Intermediate Frequency)

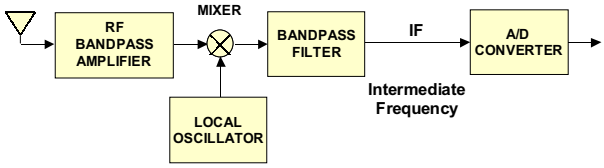

In case the antenna signal frequency is too high to be digitized directly by the A/D converter, it has to be translated down using an analog mixer and local oscillator.

The top diagram shows a simplified representation of this analog translation to baseband with a low pass filter following the mixer.

The bottom diagram shows the translation to an intermediate frequency or IF — this is quite common. In this case, the filter is a bandpass filter centered at the IF frequency.

So far, we've discussed three types of front end circuitry:

1) Direct sampling with no translation (on the previous slide)

2) Analog translation to baseband

3) Analog translation to IF

But how do we design the filters in each case?

Let's go back to review some fundamental sampling theory.

Pentek, Inc. • One Park Way, Upper Saddle River, NJ 07458 • Tel: (201) 818-5900 • Fax: (201) 818-5904 • Email: info@pentek.com • Web Site: www.pentek.com

Filtering Helps Avoid Noise and Aliasing

- In all systems, the A/D input must be filtered for two important reasons:
  - Eliminate out of band noise
  - Eliminate aliasing
- Nyquist sampling theorem requires the input signal bandwidth must be less than one-half the sampling rate of the A/D converter
- Some systems (like an IF stage) provide inherent bandlimiting before the A/D
- Fundamental Sampling Modes
  - Baseband Wideband Sampling
  - Baseband Pre-select Sampling
  - Undersampling

Filters ahead of the A/D are needed primarily for two reasons: to eliminate out of band noise and to eliminate out of band signals that can cause aliasing.

Nyquist tells us that whenever you sample a signal with an A/D, the bandwidth of that signal must be less than half the sampling rate of the A/D.

Filters help us guarantee that this rule is met. Sometimes the bandwidth is already limited by the signal source, like the output of an IF stage that takes advantage of the IF filter bandwidth. But each case has to be analyzed individually.

The design of the filter is also critically linked to the sampling mode. Here we've listed three fundamental sampling modes:

1) Baseband **wideband** sampling

2) Baseband **pre-select** sampling

3) Undersampling, which is also sometimes called sub-sampling

To help you get a feel for the filter requirements of each mode, we present a convenient tool for analyzing the effects of sampling in the frequency domain.

Fan Fold Paper Model to Visualize Sampling

- Plot the spectrum of the input signal on transparent fan-fold printer paper scaled so the frequency axis is aligned with multiples of Fs on the backward folds

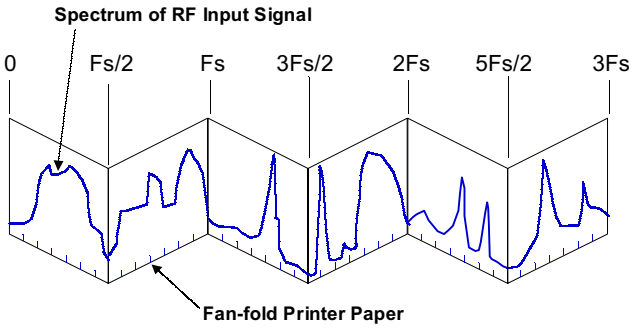

This simple technique has been very useful to our customers and our own applications engineers to help them understand what happens during sampling.

Imagine that we have a stack of fan-fold computer printer paper with transparent sheets.

Now, we assign the frequency axis along the bottom edge of this paper, scaled so that multiples of the sampling rate line up with the backward folds of the paper, as shown.

Now, using that frequency scale, we plot out the spectrum of the signal we want to sample, with amplitude plotted on the vertical axis.

Fan Fold Paper Model to Visualize Sampling

- Now collapse the stack of transparent fan-fold paper and look through all the sheets
- This represents how sampling "folds" the entire RF input spectrum into a single page from 0 to Fs/2
- Once aliasing occurs, there is no way to undo it

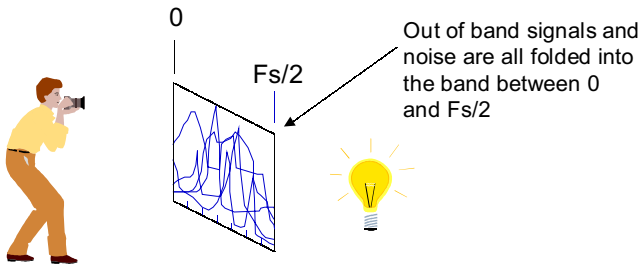

Now let's collapse the stack of transparent paper flat together and hold the stack up to a light so we can see through all the sheets.

We are now looking at the frequency plot of the sampled signal at the output of the A/D converter.

Notice that we've lost a lot of information because we can't tell which sheet a particular signal is on. And, unfortunately, after sampling, that information is lost forever.

We've also contaminated any particular signal with signals from other sheets which have folded on top of it.

Not only that, we've also folded the noise from all the sheets so they pile up in the region between DC and the half sampling rate, potentially ruining your signal to noise ratio.

How do we avoid this mess in each of the three sampling modes?

Baseband Sampling of Wideband Signals

- For baseband signals over a wide frequency range, use a low pass filter with cutoff frequency, Fc, less than Fs/2, where Fs is the A/D sample rate
- After sampling, only the baseband signal is captured, eliminating folding of aliased signals and noise

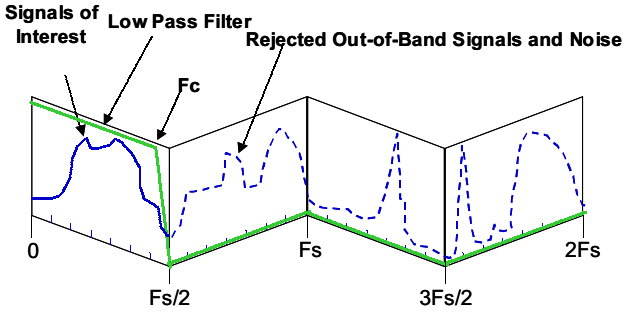

For the baseband **wideband** sampling mode, where we want to look at everything from DC up to a frequency below the half sampling rate, we can install a low pass filter with a cutoff frequency, Fc, located below Fs/2.

The frequency response of the filter is shown in green.

Now all of the out-of-band signals and noise on the pages above Fs/2 are eliminated so that when the folding occurs, it doesn't corrupt the baseband signal.

Baseband Sampling of Pre-Select Signals

- For narrowband signals at baseband, using a pre-select bandpass filter can optimize the dynamic range of the A/D converter by rejecting strong adjacent signals and out-of-band signals and noise
- Pre-select filter is a bandpass filter whose passband is centered on the signal of interest

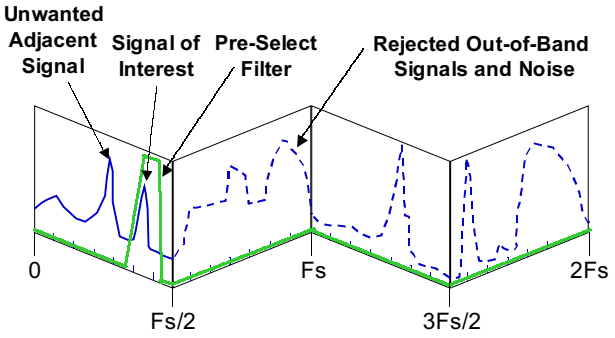

For the baseband **pre-select** sampling mode, we need to use a bandpass filter with the frequency response shown in green.

We get the same benefits as the previous case for out-of-band signals and noise above Fs/2, but more importantly, we can keep large adjacent signals like the one shown, from getting to the A/D converter.

The reason for this is that if the large unwanted signal gets through to the A/D converter, it uses up the dynamic range of the A/D.

For applications where there are known, strong unwanted signals, this technique can be extremely useful in improving the signal-to-noise ratio of the smaller signal of interest.

Principles of Undersampling

- For narrowband signals above Fs/2, undersampling can be used to intentionally "alias" the input signal
- Very useful for IF outputs of UHF/VHF receivers
- Successful undersampling needs careful selection of:
  - Signal Frequency
  - Signal Bandwidth
  - Bandpass Filter
  - Sampling Frequency

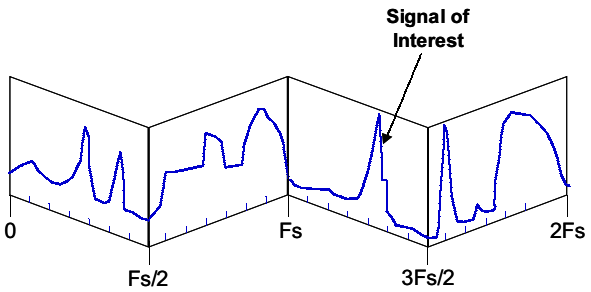

The third sampling mode, called undersampling or sub-sampling, is ideal for many systems that use an analog RF translator front end. These receivers usually deliver IF outputs, often at 21.4 or 70 MHz, with bandwidths ranging from a few kilohertz to tens of MHz, depending on the receiver.

If we wanted to perform baseband sampling on a 70 MHz signal, we would have to choose a sampling rate of well over 140 MHz. This may require an A/D that adds cost and power to the system.

However, because the IF signal is inherently bandlimited, we can take advantage of the folding caused by sampling and use a lower frequency A/D.

This is a little tricky since you have to carefully choose the sampling frequency and filtering according to the signal frequency and bandwidth.

Let's see how.

Principles of Undersampling Design: Step 1

- Step 1: Design a bandpass filter or IF filter to pass the band of interest and reject all other signals to meet spurious and S/N requirements
- Tradeoffs
  - Sharper filter adds complexity, expense, calibration, space, etc.
  - Sharper filter allows lower A/D sample rate

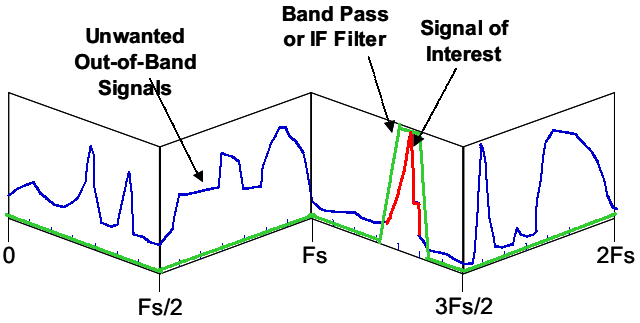

The fan fold paper really comes in handy here.

First, design a bandpass filter that rejects unwanted signals and noise.

This is often fully satisfied by the standard IF filter in the RF translator, but you do have to check this.

Sharper filters add cost and maintenance but they do let you get away with a lower sampling rate as we'll see on the next slide.

Second (top of next column), choose a sampling frequency so that the passband of the filter, along with its skirts, falls entirely on a single page of fan fold paper.

There are many possible solutions to each case, so you have to pick the one that works best. You may have to go back and forth a few times to readjust the filter and sampling rate to get the best scheme.

Principles of Undersampling Design: Step 2

- Step 2: Choose a sampling frequency so that the filter pass band and skirts fall entirely within one page of the fan-fold paper
- Tradeoffs
  - Higher sampling rate allows broader bandwidth & simpler filter
  - A/D's with lower sampling rates are more accurate & less expensive

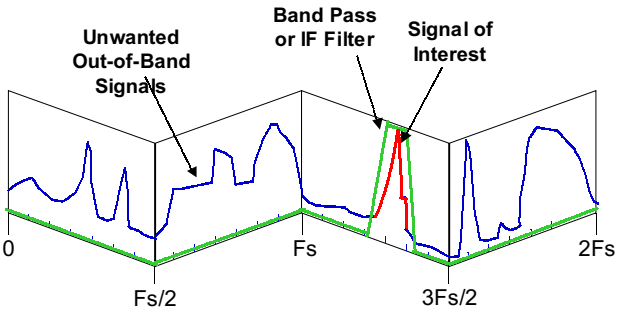

Here are some trade-offs to consider.

With a higher sampling rate, the pages are wider and the filter becomes less complex. Also, there is a lower noise density folded into the 0 to Fs/2 band after sampling.

At higher sampling rates, however, the A/D is more expensive and the number of bits of accuracy drops off.

You also need to be sure that the A/D has a good wideband input stage to handle the IF signal with minimum distortion.

Equally important is the aperture uncertainty or phase jitter of the sample-and-hold amplifier, which is usually part of the A/D.

To make this job easier, many A/D converters are now specifically characterized to operate in undersampling applications.

Undersampling Performs Frequency Translation

- Signal of interest folds into the 0 to Fs/2 region
- Undersampling performs an automatic frequency translation
- Translated image may be reversed in frequency depending on which side of the "fold" the input falls

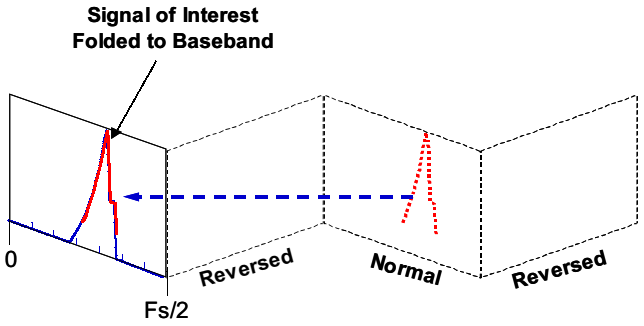

The effect of undersampling, as you probably expected, is that the IF signal is folded down to the first page. This is really an automatic frequency translation, performed for free by the sampling process.

For the signals on every odd numbered sheet, the effect is a frequency translation by a multiple of Fs. For the signals on even numbered sheets, there is a reversal of the frequency axis on that sheet, followed by a translation by an odd multiple of Fs/2. Again, this is much easier to follow by visualizing the fan-fold model.

This undersampling technique is extremely popular in software radio systems which almost always follow the A/D converter with a digital down converter (DDC).

Regardless of where the undersampling folding process translated the signal of interest, the DDC can translate it down to 0 Hz as a complex baseband signal. Once the complex signal is at baseband, the reversal of the frequency axis is easily undone by simply changing the sign of the Q component.

Guidelines for Sampling and Undersampling

- Use the fan fold paper to validate your sampling plan for the characteristics of your input signal
- Carefully evaluate A/D specifications for operation in the undersampling mode
- Ensure low-noise, wideband circuitry in the front end ahead of the A/D
- Transforming coupling often is superior to an amplifier for IF or RF input signals
- Eliminate as many out-of-band signals and noise as possible, since they will fold
- Ensure the the sample clock is clean with low phase noise and jitter

There are usually several different sample clock frequencies that will work for undersampling. While the fanfold paper model can show all of the correct frequency plans, the best choice will usually be determined by several other important practical considerations shown above.

Some A/D converters are specifically characterized for undersampling applications, while others are designed only for baseband sampling, so be sure to verify the specifications.

Noise and distortion on the input signal must be minimized so these components don't fold into the sampled signal. Special care must be taken to preserve the purity of the sample clock signal.

Undersampling can be an extremely valuable tool for software radio applications, since it can eliminate at least one additional stage of analog frequency translation and simplify system design.

Undersampling allows you to use an A/D converter with a lower sampling rate, which usually means more bits of resolution and better dynamic range. This lower sample rate also reduces the cost and complexity of the next stage of digital signal processing, recording, storage, or transmission.

Section 3. FPGA Technology

FPGAs: The Essential Companion for High Speed A/Ds

- On-chip processor cores
- Internal clock rates up to 600 MHz
- Reduced power with core voltages near 1 volt
- Dedicated on-chip hardware multipliers
- Memory densities of over 10 million bits
- High-density BGA and flip-chip packaging
- Flexible memory structures
- Logic densities of over 10 million gates
- Silicon geometries near 0.1 microns
- On-board gigabit serial interfaces
- Over 1200 user I/O pins
- Configurable interface standards

FPGAs, or Field Programmable Gate Arrays, are commonly coupled to high speed A/Ds for several key reasons: they can perform real-time digital signal processing faster than general purpose programmable processors and they offer extremely high speed interfaces to other system components, including built-in interfaces to the new high-speed serial switched fabrics.

BGA and flip chip packages provide plenty of I/O pins to support these on-board gigabit serial transceivers and other user-configurable system interfaces.

Other important features are on-chip processor cores, computation clocks of up to 600 MHz, and lower core voltages to keep power and heat down.

In the late 1990s, dedicated hardware multipliers started appearing and now you'll find literally hundreds of them on chip as part of the DSP initiative launched by virtually all FPGA vendors.

High memory densities coupled with very flexible memory structures meet a wide range of data flow strategies. Logic slices with the equivalent of over 10 million gates result from silicon geometries shrinking down to 0.1 microns.

FPGAs: New Development Tools

- High Level Design Tools
  - Block Diagram System Generators
  - Schematic Processors
  - High-level language compilers for VHDL & Verilog
  - Advanced simulation tools for modeling speed, propagation delays, skew and board layout
  - Faster compilers and simulators save time
  - Graphically-oriented debugging tools
- IP (Intellectual Property) Cores
  - FPGA vendors offer both free and licensed cores
  - FPGA vendors promote third party core vendors
  - Wide range of IP cores available

To support such powerful devices, a whole new world of design tools are appearing that now open up FPGAs to both hardware and software engineers. Instead of just accepting logic equations and schematics, these new tools accept entire block diagrams as well as VHDL and Verilog definitions.

Choosing the best FPGA vendor often hinges heavily on the quality of the design tools available to support the parts.

To minimize some of the tricky timing work for hardware engineers, excellent simulation and modeling tools help you quickly analyze worst case propagation delays and suggest alternate routing strategies to minimize them within the part. This can really save you hours of tedious trouble shooting, not only during design verification but also for production testing.

In the last few years, a new industry of third party IP or intellectual property core vendors now offer thousands of application-specific algorithms, ready to drop into the FPGA design process to help beat the time-to-market crunch and to minimize risk.

FPGAs: Key Resources for DSP

- Parallel Processing
- Hardware Multipliers for DSP
  - FPGAs can now have over 500 hardware multipliers
- Flexible Memory Structures
  - Dual port RAM, FIFOs, shift registers, look up tables, etc.
- Parallel and Pipelined Data Flow
  - Systolic simultaneous data movement
- Flexible I/O
  - Supports a variety of devices, buses and interface standards
- High Speed
- Available IP cores optimized for special functions

Like ASICs, all of the logic elements in FPGAs can execute in parallel. This includes the hardware multipliers, and you can now get over 500 of them on a single FPGA.

This is in sharp contrast to a programmable DSP, which normally have just a handful of multipliers that must be operated sequentially, instead.

FPGA memory can now be configured in the design tool to implement just the right structure for your task including dual port RAM, FIFOs, shift registers and other popular memory types.

These memories can be distributed along the signal path, interspersed with the multipliers and math blocks, so that the whole signal processing task operates in parallel, in a systolic pipelined fashion.

Again, this is dramatically different from sequential execution and data fetches from external memory in a programmable DSP.

As we've said, FPGAs now have specialized serial and parallel interfaces to match requirements for high speed peripherals and buses.

FPGAs Bridge the Software Radio Application Task Space

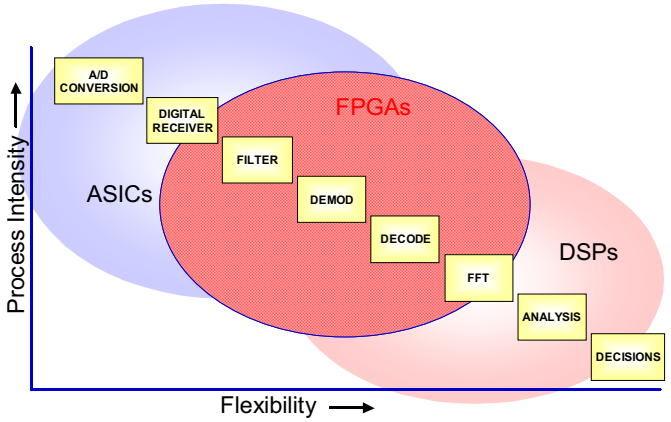

As a result, FPGAs have significantly invaded the application task space as shown by the center bubble in the task diagram above.

They offer the advantages of parallel hardware to handle some of the high process intensity functions like digital receivers, and the benefit of programmability to accommodate some of the decoding and analysis functions of DSPs.

These advantages may come at the expense of increased power dissipation and increased product costs, but these considerations are often secondary to the performance and capabilities of these remarkable devices.