Opposite Page:

Circuit board including IC in 0.18 µm CMOS process. Courtesy of A. Chen (A.I. Akinwande and H.-S. Lee)

Sponsor: 3M and MARCO Focused Research Center on Integrated Circuits and Systems (C2S2) (MARCO/DARPA)

Integrated Circuits and Systems



Integrated Circuits and Systems

- µAMPS-2: A Fully-Integrated Energy-Agile Wireless Sensor Node
- Energy and Quality Scalable Wireless Communication
- Energy Efficient Multitarget-Multisensor Tracking on the uAMPS Platform
- Energy-Scalable 1024 Real-Valued FFT for Wireless Sensor Networks
- Novel Techniques Addressing Delay and Power Issues in Deep Sub-Micron Interconnect
- Design Methodology for Fine-Grained Leakage Control in MTCMOS
- PLL-Based Optical Clock Distribution
- Power-Aware Reconfigurable Hardware for Digital Baseband Processing
- Radio Design for a Power-Aware Microsensor Node
- Exploration of Sub-Threshold Circuit Design Topologies
- Ultra-Wideband Radio Antenna Design
- Ultra-Wideband Radio Transceivers
- Ultra-Wide Band (UWB) Short-Haul Data Communication
- Wireless Gigabit Local Area Network
- 5.8 GHz Wideband Receiver for Wireless Gigabit LAN
- Analog Base-band Processor for Wireless Gigabit LAN
- Effect of Circuit Nonlinearity on System Performance Metrics, BER and Spectral Efficiency
- On-Chip Cross-Talk Analysis for an Array of Transceivers on a Single Chip
- Linear Power Amplifier Design for WiGLAN
- Smart Active-Matrix Display Drivers For Organic Light Emitting Devices
- A CMOS-Compatible Compact Display
- A Programmable, Wide Dynamic Range CMOS Imager with On-Chip Automatic Exposure Control
- Mixed-Signal Design in Deeply Scaled CMOS Technology
- A CMOS Bandgap Current and Voltage References
- Flicker Noise in Scaled CMOS Devices
- Device Level Optimization of Phase Noise in Integrated LC VCOs
- Radio Frequency Digital-to-Analog Converter

continued Integrated Circuits and Systems

- Low Power RF Front-End for Wireless Microsensor Systems
- Substrate Noise Coupling and Reduction Techniques in Mixed-Signal Systems
- An Advanced Model Based Vision System for Intelligent Transportation Systems
- Efficient Traffic Monitoring
- Sensor Fusion for Automobile Applications
- Superconducting Bandpass Delta-Sigma A/D Converter
- Circuit and System Level Tools for Thermo-Aware Reliability Assessments of IC Designs
- Intelligent Transportation Systems
- The Low-Power Bionic Ear Project
- The Visual Motion and Inertial Sensing Project
- Spike-Based Hybrid Computers Project

µAMPS-2: A Fully-Integrated Energy-Agile Wireless Sensor Node

Personnel N. Ickes and C. Schurgers (A. P. Chandrakasan)

Sponsorship

DARPA

In recent years, the idea of wireless microsensor networks has garnered a great deal of attention and interest. A distributed wireless microsensor network consists of hundreds to several thousands of small sensor nodes scattered throughout an area of interest. Each node individually monitors the environment and collects data as directed by the user, but the network collaborates as a whole to deliver high-quality observations to a central base station. The large number of nodes in a microsensor network enables highresolution, multi-dimensional observations and faulttolerance that are superior to more traditional sensing systems. With these advantages in mind, microsensor networks hold great promise for applications such as warehouse inventory tracking, location-sensing, machine-mounted sensing, patient monitoring, and building climate control.

In the μ AMPS-2 project, our aim is to build a highly integrated, yet versatile sensor system with an extreme focus on energy efficiency. This is achieved through custom ASIC design, while supporting the flexibility that is crucial in the dynamic operating environments that are typical for sensor networks. To this end, our hardware exhibits energy-agility: adapting the internal settings and circuit parameters on-the-fly to track the most energy optimal operating point under those varying conditions.

The μ AMPS-2 architecture consists of a micropower DSP, surrounded by dedicated accelerator blocks for functions performed frequently by each sensor node. Accelerators for FFTs, FIR filtering, error correction coding/decoding, and data encryption are currently being designed. Each of these blocks is optimized for its specific task, while incorporating support for aggressive energy-agility. Furthermore, fine-grain shutdown modes are provided such that power is only consumed when strictly needed. Unused accelerators are dormant when not actively processing information. This architecture of highly optimized, on-demand hardware support for energy intensive tasks allows for ultra low-power data manipulation and lowers the processing burden on the DSP core.

This DSP core is at the heart of the µAMPS-2 node and employs a custom instruction set architecture optimized for microsensor applications. It has a 16bit datapath, and will run at up to 10MHz, which is sufficient for typical microsensor applications. Clock speed can be dynamically reduced to tens of kilohertz or less for energy savings when less computational power is required. The on-chip data (~32kB) memory and instruction cache minimize the energy expended storing sensor measurements and fetching program instructions. Because many microsensor algorithms operate on blocks of sensor measurements, a sophisticated DMA (direct memory access) engine is employed to move data around the chip. The DMA engine can move blocks of measurements between accelerators and memory without intervention from the DSP.

Το μAMPS-2

Energy and Quality Scalable Wireless Communication

µAMPS-2 Digital ASIC Architecture DSP Extra bus provides





Personnel

R. Min (A. P. Chandrakasan)

Sponsorship

DARPA, ARL Collaborative Technology Alliance, Hewlett-Packard under the MIT Alliance, and NDSEG Fellowship

Next-generation wireless devices will be characterized by shrinking size and increasing density, which together place an increasing emphasis on energy efficient operation. Moreover, emerging wireless applications such as microsensor networks will exhibit high operational diversity, reflected by time-varying environmental conditions, user requirements, and the nodes' own role in the network. Effective energy management strategies must therefore foster *energy scalability* through graceful energy vs. quality trade-offs in response to changing operational conditions.

Graceful energy vs. quality scalability for wireless communication requires two prerequisites. First, the notion of communication "quality" must be defined. Hence, we define communication quality by four of its fundamental metrics: range, delay, reliability, and energy. We then introduce a basic API that allows an application to specify these metrics. Various combinations of delay, reliability, and energy can be chosen by direct specification. The communication range desired by an application can be expressed in a variety of ways, such as the distance to a specified node, a group of nodes, or the *n* nearest nodes. A second prerequisite is hardware that enables graceful trade-offs of energy and performance. The µAMPS-1 node provides flexible coding and output power settings.

Figure 2 then illustrates the energy required by the µAMPS-1 node for communication of variable reliability (bit error rate) and range. Note that, as the reliability or range requirements of communication increase, the energy required for communication increases monotonically. Each "step" in the graph corresponds to the minimum-energy selection of radiated power and coding policy that will meet the required quality constraints. The range and reliability of communication increase as more power is radiated from the transmitter or stronger error-correcting codes are used, resulting in longer transmit and receive times, and higher decoding energy.

continued

Energy Efficient Multitarget-Multisensor Tracking on the uAMPS Platform

Personnel

K. Atkinson (A. P. Chandrakasan and C. Rohrs)

Sponsorship

ARL Collaborative Technology Alliance

For the uAMPS project, we have a network of acoustic sensing nodes producing a line-of-bearing to a sound source. The question then becomes how to integrate the data to produce the most accurate estimate of the target position. The line-of-bearing measurements are quite noisy in practice and simple triangulation would yield poor tracking performance. To achieve superior estimation accuracy we have implemented algorithms based on the Kalman filter, a fundamental algorithm of estimation theory.

The Kalman filter has two main advantages. First, it incorporates information regarding the statistical properties of both the target's motion and the measurements available, allowing for improved performance in the presence of noise. Second, it produces not only the estimate of the target's state, but also probabilistic information regarding the accuracy of the estimate (covariance matrix).

The covariance matrix and other probabilistic information computed by the Kalman filter serve a variety of purposes. For example, they can be used to validate measurements and reject those that are clearly spurious by computing the probability a particular measurement came from the target. (See Figure 3) Monitoring the covariance matrix can also indicate which sensors are providing useful information to localize the target, and allow for unneeded sensors to be shut down to save power.

Work is currently underway to extend the system to simultaneously track multiple targets. In this scenario the main challenge involves determining the measurement to target associations. The technique used is known as joint probabilistic data association, and computes the probability that each measurement validated for a target actually originated from that target. The state estimate is then updated with an appropriately weighted combination of all validated measurements.



Fig 2: Energy-scalable communication on µAMPS-1 node.

continued

Energy-Scalable 1024 Real-Valued FFT for Wireless Sensor Networks

Personnel A. Wang (A.P. Chandrakasan)

Sponsorship

Energy efficient Digital Signal Processors (DSP's) is an important component of wireless sensor networks, where tens to thousands of battery-operated microsensors are deployed remotely and used to relay sensing data to the end-user. Given the constantly changing environments of sensor devices and the extreme constraints on battery lifetimes, system level energy-aware design considerations should be taken into account.

Energy-aware design is in contrast to low power design, which targets the worst case scenario and may not be globally optimal for systems with varying conditions. The energy-awareness of a system can be increased by adding additional hardware to cover functionality over many scenarios of interest and to tune the hardware such that over a range of scenarios, the system is energy-efficient. One algorithm that is widely used in sensor and wireless applications is the Fast Fourier Transform (FFT). In the area of sensor signal processing, the FFT is used in frequency domain beamforming, source tracking, harmonic line association and classification.

Energy-quality scalability for an system is needed if the environment of the device changes constantly. An energy-aware FFT will be able to adapt energy consumption as energy resources of the system diminish or as performance requirements change. Therefore, it is advantageous to design the FFT with energy scalability hooks such as variable memory size and variable bit precision, so that it can be used for a variety of scenarios. Our design focuses on a Real-Valued FFT (RVFFT) which can scale between 128-512-point FFT lengths and can operate at both 8 and 16-bit precision computation.

In this work two energy-aware architecture designs are developed. These architectures are evaluated in the context of a variable bit precision and variable

7



Fig. 3

We currently have a real-time Kalman filter based tracking application running on the uAMPS nodes for a single target. (Figure 3 shows a screenshot.) Work in is progress to implement the multiple target tracking algorithms in addition to power-saving enhancements to shut down unnecessary sensors.

Intel Fellowship and DARPA

Novel Techniques Addressing Delay and Power Issues in Deep Sub-Micron Interconnect

Personnel

T. Konstantakopoulos (A. P. Chandrakasan)

Sponsorship

MARCO Focused Research Center on Interconnect (MARCO/DARPA)

In deep-submicron technologies the primary component of delay is shifting from logic gates to the interconnect network. Buses can no longer be considered as a set of independent lines that don't interact. A more appropriate model would treat the bus as a distributed system where a transition on a line would affect adjacent lines as well. However, the transitions on a bus can be grouped into delay classes depending on the effective capacitance that the driver circuit needs to charge.

An effective approach to reduce delay in interconnect is by eliminating the transitions that are relatively time consuming. We are using coding schemes to accomplish that by increasing the number of lines in the bus, thus imposing some redundancy. In our implementation, which is shown in Figure 5, we are mapping a 4-line bus to a 6-line bus. A test chip to verify the proposed coding scheme has been fabricated and is being tested.



Fig. 5

FFT length RVFFT. The scalability of the RVFFT was implemented in a 0.18-micron process for energyawareness measurements and hardware verification. Figure 4 shows a die photo of the scalable RVFFT implementation.



Fig. 4: Die photograph of the Energy-Scalable RVFFT.

Design Methodology for Fine-Grained Leakage Control in MTCMOS

Personnel B. Calhoun (A. P. Chandrakasan)

Sponsorship Texas Instruments and DARPA

Texas Instruments and DARPA

Multi-threshold CMOS is a popular technique for reducing standby leakage power with low delay overhead. Most MTCMOS designs use large sleep devices to reduce standby leakage at the block level. The use of sleep devices at a local, gate level remains largely unexamined. One reason for choosing large sleep FETs is the design complexity associated with placing them locally. Segmenting the sleep transistor into many devices at the gate level creates a greater possibility for sneak leakage paths. Generally, sneak leakage paths are high leakage paths between power and ground that remain during sleep mode. Analysis of sneak leakage paths at the local level can provide insight that eliminates much of this complexity. Locally placed sleep devices offer several advantages over the large sleep FET approach, such as guaranteed circuit functionality at high speed, standard-cell MTCMOS design, and improved noise margins.

We developed a formal examination of sneak leakage paths and a design methodology that enables gatelevel insertion of sleep devices for sequential and combinational circuits. A fabricated 0.13μ m, dual VT testchip employs this methodology to implement a lowpower FPGA architecture with gate-level sleep FETs and over 8X measured standby current reduction. The methodology also allows local sleep regions that reduce leakage in active CLBs by up to 2.2X (measured) for some CLB configurations.

An MTCMOS circuit uses a high VT sleep device between a low VT circuit and one rail, usually the ground rail. Whether the sleep device is one large device or many small ones, the basic structure of MTCMOS circuits suggests that sneak leakage paths must occur only where the sleep device(s) can be bypassed. Thus, the focal point for preventing sneak leakage paths is the interface between MTCMOS and CMOS circuits. The rules isolate cases where sneak leakage occurs at this interface. The testchip confirms that gate-level sleep devices can provide standby leakage savings. Placing the entire chip in sleep mode provides a measured reduction in leakage current by from 7.0X to 8.6X. We also propose fine-grained sleep regions and implement them on the testchip. The chip measurements match closely with simulation, and they show the benefit of the sleep region technique. The use of gate-level sleep devices allows inactive circuit regions to enter sleep at a fine grain. The other circuit components remain active and with unaffected performance. The total steady-state power (clock-gated) for an active CLB reduces by from 10% to 2.2X for different configurations. Figure 6 shows an annotated die photo of the chip.



Fig. 6: Annotated Die Photo

PLL-Based Optical Clock Distribution

Personnel

A. Kern (A.P. Chandrakasan)

Sponsorship

MARCO Focused Research Center on Interconnect (MARCO/DARPA)

Because optical signals do not experience the timevariant propagation delays associated with the distribution of electrical signals, optical synchronization could potentially be used to generate extremely precise clocks. Many existing optical clock distribution systems, however, fail to fully utilize this potential because they are limited by the skew and jitter introduced in the transimpedance amplifier stages used for the optical to electrical signal conversion.

Removing the explicit optical to electrical signal conversion step would result in a reduction of skew and jitter. This direct conversion step may be eliminated by generating the clock with an optically-locked PLL that contains a phase detector capable of directly comparing the phases of an optical and an electrical signal. (See Figure 7) The proposed phase detector uses the electrical feedback signal to steer the photocurrent generated by the optical signal.

The PLL is based on the traditional architecture and contains the phase detector, loop filter, VCO, and feedback divider. A LC VCO was chosen to minimize oscillator jitter and frequency range. Because a phase detector is used for comparison, the VCO range must be within a factor of two of the reference frequency or the loop could lock to a harmonic. Due to its superior jitter performance, a synchronous divider was used to generate the 1.6 GHz output frequency from an optical input frequency of 200 MHz.

A first-generation test chip will be fabricated in the TSMC 0.18 um process and tested. Future work will examine using an additional electrical feedback loop to aid frequency acquisition and investigate circuits for direct phase-frequency comparison of optical and electrical signals.



Fig. 7

Power-Aware Reconfigurable Hardware for Digital Baseband Processing

Personnel

F. Honoré (A.P. Chandrakasan)

Sponsorship

MARCO Focused Research Center on Interconnect (MARCO/DARPA)

This project continues previous work on building energy-scalable solutions for wireless systems. A Field Programmable Gate Array (FPGA) architecture is the basis for exploring a fine-grain hardware approach to creating a platform for signal processing by introducing low-level power control. A novel Configurable Logic Block (CLB) with enhancements for distributed arithmetic computation and power control regions was validated through a testchip using a 0.13-um dual threshold voltage process.

The new CLB design yields much more efficient algorithm mapping over existing implementations by reducing logic block utilization by 50% or better. For more efficient power usage regions within the logic block have the ability to automatically power down when not in use. The results show that this fine-grain MTCMOS approach achieves significant power savings in the range of 1.2x to 2.7x by reducing active mode subthreshold leakage.

Interconnect overhead for FPGA's can be a significant fraction of the power budget for large designs. Programmable switch elements introduce large delays in long paths. Additionally, deep submicron buses have large parasitic interwire coupling capacitances that further increase delay of long wires. A new programmable switch architecture is being explored that will allow long paths to be pipelined to meet critical path timing or alternatively allow reduced voltage operation with minimal performance impact. Bus coding techniques are applied on groups of long routes to mitigate the impact of the interwire capacitance effect. Additionally, routing area overhead is reduced by more efficient placement of configuration storage elements and utilizing the available layers of metal.

By introducing these fine-grain hardware controls at several levels, this low power FPGA forms the basis for a platform that allows for effective in-system energydelay tradeoffs for energy-constrained systems. .92 mm



Radio Design for a Power-Aware Microsensor Node

Personnel

D. Wentzloff (A. P. Chandrakasan)

Sponsorship

DARPA

The goal of the MIT μ AMPS project is to develop a wireless network of small, low-power, general purpose sensor nodes that can collectively gather information about their surroundings and wirelessly relay it to a base station. The nodes are general purpose because they have a generic A/D interface that can connect to a variety of sensors, and the DSP and hardware accelerators can be selectively used or disabled depending on the complexity of the processing required by the application. The MIT μ AMPS project is now in phase II: the development of a two-chip, power-aware node. One chip will include all of the digital hardware and sensor A/D converters. The second chip will include all of the RF hardware for the radio.

Since the first prototype radio for the μ AMPS project was built, a need has arisen for long range wireless communication between nodes. To accommodate this, a new radio is being designed with the addition of a 1 Watt power amplifier that will improve the transmission distance by a factor of 10. Other improvements to the radio are lower idle and receive power consumption and the use of a radio chip with an integrated 8-bit microcontroller. The 8-bit microcontroller can be used as a protocol processor for the radio, and will greatly simplify the interface from the radio to the digital circuits.

Using this prototype radio as a benchmark, a chip will be fabricated with a complete, ultra low power radio having multiple knobs for varying its performance. This will allow the power consumption of the radio to be fine tuned based on the node's environment and proximity to other nodes. To accomplish this, RF components must be re-designed with the appropriate adjustments and feedback for regulating the performance. A test chip has been fabricated in collaboration with Nisha Checka that has seven Voltage-Controller Oscillators (VCO) operating at different frequencies, and digital circuits that intentionally inject noise into the substrate. By adjusting the power consumed by a VCO, the signal to noise ratio at the output can be controlled. This chip is currently in the testing phase. A picture of the die is shown in the figure. By adding functionality like this to a radio, the protocol processor will be able to vary parameters such as linearity and bandwidth in addition to transmit power. This has the potential of lowering the overall energy consumption per bit of the radio.



Fig. 9 Die photo of the VCO chip and digital circuits

Exploration of Sub-Threshold Circuit Design Topologies

Ultra-Wideband Radio Antenna Design

Personnel J. Cline (A. P. Chandrakasan)

Sponsorship SRC Fellowship and DARPA

Sub-threshold logic design is a circuit technique that entails lowering the voltage supply below the threshold voltage of transistors to reduce the dynamic power dissipation. At such low voltage levels, sub-threshold currents are used to charge and discharge the logic.

My research aims to compare and contrast different logic design methodologies for the sub-threshold regime. The logic studied was evaluated in full-adder designs and implemented in a 16x16 bit Baugh-Wooley Multiplier.

An important observation obtained from the implementation above relies on the fact that subthreshold current depends exponentially on the gate voltage. In strong inversion, this exponential relationship gives an 'on' current to 'off' current ratio of approximately a few thousand to one. In contrast, subthreshold designs have voltages around a few hundred millivolts and thus, the ratio of the currents can be as small as 10:1. When implementing CMOS logic in the sub-threshold regime, designs must account for this small current ratio and the process variations must be tightly monitored. At the worst process corners, the already small ratio of currents can affect the speed and even the circuit functionality. Thus, the worst case process corners dictate the lowest functional voltages.

In conclusion, this research compared different logic families, such as transmission gate, static CMOS, and dynamic logic, implemented in the sub-threshold regime. The performance and energy tradeoffs were analyzed for an array of process variations and variable supply voltages, as low as 100mV.

Personnel

J. Powell (A. P. Chandrakasan)

Sponsorship

Presidential Fellowship, DARPA and HP under the MIT Alliance

The recent allocation of the 3.1-10.6 GHz spectrum by the Federal Communications Commission for Ultra Wideband (UWB) radio applications has presented a myriad of exciting opportunities and challenges for design in the communications arena, including antenna design. Ultra Wideband Radio requires power spectral densities of -43.1dBm/MHz and bandwidths greater than 50% of the center frequency. Due to the power constraints, many UWB designs occupy the entire bandwidth. Successful transmission and reception of an Ultra Wideband pulse that occupies the 3.1-10.6 GHz spectrum require an antenna that has linear phase and VSWR \leq 2 throughout the entire band. Linear phase ensures constant group delay which is imperative for transmitting and receiving a pulse with minimal distortion. VSWR ≤ 2 is required for proper impedance matching throughout the band, ensuring at least 90% total power radiation. This corresponds to a return loss of greater than 10 dB throughout the band. Compatibility with an integrated circuit also requires an unobtrusive, electrically small design.

One method for achieving broadband characteristics uses Babinet's Equivalence Principle of duality and complementarity. The principle states that the product of the input impedances of two planar complementary antennas is such that $Z_1Z_2=\eta^2/4$. This is illustrated in the spiral slot antenna design in Figure 10, which incorporates two complementary spirals in one antenna element. This design shows that the metal spiral is the exact complement of the free-space spiral, which then requires that $Z_1=Z_2=\eta/2$. This impedance can be adjusted based on the choice of dielectric constant value.

Preliminary simulations of this design have shown promising results with regard to VSWR bandwidth and linear phase, at a size of 4.5 cm x 4.5 cm. Multiple implementations will be designed on various PC boards to determine an optimum dielectric constant value and dielectric thickness based on bandwidth and beamwidth

continued

Ultra-Wideband Radio Transceivers

Personnel

R. Blazquez, F. S. Lee, P. Newaskar, J. Powell, D. D. Wentzloff (A. P. Chandrakasan)

Sponsorship

HP-MIT Alliance, Presidential Fellowship, DARPA, Air Force Research Laboratory

The recent approval of Ultra-Wideband (UWB) wireless technology by the Federal Communications Commission has presented a myriad of exciting opportunities for circuit design, system design and antenna design. Depending on the application, UWB signals utilize bandwidths from DC to 960 MHz or 3.1 GHz to 10.6 GHz. Contrary to traditional narrowband, single-tone radio signals, a UWB signal is typically composed of a pulse train of sub-nanosecond pulses modulated either in polarity or position as shown in Figure 11. The narrowness of the pulse in the time domain corresponds to the wideness of the band in the frequency domain. Since the total power is spread over such a wide swath of frequencies, its power spectral density is extremely low. This minimizes the interference caused to existing services that already use the same spectrum. On account of the large bandwidth used, UWB links are capable of transmitting data over tens and hundreds of megabits per second.



Fig. 11 : Transmitting Information in a UWB System

To date, we have implemented and tested a base-band (DC to 960MHz) front end UWB chip shown in Figure 12 (our future research will target communication in the 3.1-10.6 GHz band). Currently we are designing a full base-band transceiver using BPSK modulation, with a symbol rate of 20 megabits per second.

Our research approach for UWB radio transceivers is divided into these five sections: transmitter, antenna, analog front-end receiver, analog-to-digital conversion/ mixed-signal processing, and digital backend.

measurements. Antenna radiation patterns, gains and efficiencies will be measured in an anechoic chamber. Pulse transmission and reception will be tested on a UWB discrete system to qualitatively determine the effects of pulse transmission and reception in the time domain. Lastly, this antenna will be tested with a UWB IC transceiver.

In summary, this research involves the design, implementation and characterization of an Ultra Wideband antenna for integration with a UWB IC transceiver.





TRANSMITTER

The transmitter for our first architecture consists of a three to one switch. Every 20 nanoseconds, the transmitter output switchs from the idle state of 0.9 volts, to either the power rail, 1.8 volts or ground rail, 0 volts, to provide positive or negative pulses. The output from the transmitter is ac-coupled to the transmit antenna. The power supply is scaled to scale the transmit power output.

A block diagram of a general UWB transceiver is shown in Figure 13.



Fig. 12 : Layout of Analog front-end for base-band UWB

ANTENNA

UWB requires power spectral densities of −43.1dBm/ MHz and bandwidths greater than 50% of the center frequency. Due to the power constraints, many UWB designs occupy the entire bandwidth. Successful transmission and reception of an Ultra Wideband pulse that occupies the entire FCC allocated spectrum require an antenna that has linear phase and VSWR ≤ 2 throughout the entire band. Linear phase



ensures constant group delay, which is imperative for transmitting and receiving a pulse with minimal distortion. VSWR ≤ 2 is required for proper impedance matching throughout the band, ensuring at least 90% total power radiation. This corresponds to a return loss of greater than 10 dB throughout the band. Compatibility with an integrated circuit also requires an unobtrusive, electrically small design. Preliminary simulations of a spiral slot antenna have shown promising results with regard to VSWR bandwidth and linear phase, at a mere size of 4.5 cm x 4.5 cm. Multiple implementations will be designed on various PC boards to determine an optimum dielectric constant value and dielectric thickness based on bandwidth and beamwidth measurements.

ANALOG FRONT-END RECEIVER

An analog front-end has been designed for the baseband, zero to 960MHz UWB band. The LNA of the front end was chosen to be single-ended, and utilizes known noise-canceling techniques to achieve a broadband low noise figure. The noise figure of the entire front-end is under 3.8dB between the frequencies of 100MHz to 1GHz. Between 10MHz and 100MHz, the noise figure is degraded by 1/f noise, and is as high as 8dB. The gain of the entire front end is 60dB, differentially. A singleended to differential circuit converts the single-ended

continued

LNA signal to a differential signal, and feeds into a cascade of differential stages to obtain the overall 60dB gain.

ANALOG-TO-DIGITAL CONVERSION

Digitizing a large-bandwidth RF signal near the antenna introduces its own set of challenges and has traditionally been considered infeasible. A high-speed, high-resolution analog-digital converter (ADC) is difficult to design, and is extremely power-hungry. However, due to the unique spectral characteristics of UWB signals and their noise environment, it can be shown that reliable detection of a UWB signal is achievable with very few bits of resolution in the ADC. A theoretical analysis of the problem has been carried out that validates the above hypothesis. It has been determined that 4 bits of resolution are sufficient for reliable detection of a UWB signal. Subsequently, design issues for a high-speed, low-resolution ADC in a fine line-width process have been explored, based on the design and implementation of a 4-bit, 4 giga-samples time-interleaved FLASH converter in 0.18mm CMOS.

DIGITAL BACKEND RECEIVER

For the first implementation of a base-band UWB receiver, a mostly digital implementation was chosen. Detection and all necessary signal processing is performed in the digital domain. As the ADC is

providing 4 giga-samples per second, each of them with 4 bits, this is the total data rate that it is necessary to process in the digital back end.

The wireless channel will imply several echoes plus the usual attenuation and interferences that degrade the reception. The premise of the design is to provide fair protection against these problems.

An important part of the receiver, at least during the beginning of the communication is the synchronization. Apart from being necessary for the detection of the signals, it allows to drastically reduce the power consumption. In order to do this an architecture based in parallel correlators are proposed and implemented.

Ultra-Wide Band (UWB) Short-Haul Data Communication

Personnel A. Chow (A.I. Akinwande and H.-S. Lee)

Sponsorship

MARCO Focused Research Center on Circuits and Systems (C2S2) (MARCO/DARPA)

High-level computation can be integrated on-chip to perform image processing, or data compression/decompression, or intelligent power management. High-resolution displays require large input data bandwidth; for example, computer monitors typically require over 2GHz bandwidth and interface circuits dissipate high power. As an example, The Silicon Image Sil 161B digital video interface receiver dissipates 800 mW. Interface circuits using compression and/or circuit techniques such as low-swing signaling can reduce the interface power dramatically, lowering the overall system power.

We are investigating a RF wireless link between the display and the host. For high resolution displays, even with on-chip data compression, the I/O data rate will still be very high. For this reason, the traditional narrow-band wireless link is not a suitable technology. We propose ultra-wideband data communication technology for host-to-display data communication. This technology can potentially be extended to chipto-chip and back plane data communication as well. The ultra-wideband communication, which has been in limited use for medium-to-long range (~mile), lowdata rate communication, employs a train of impulses rather than a single frequency RF carrier. The impulse train has a very wide frequency spectrum, typically DC- GHz range. Since the energy is spread in such a wide frequency range, there is negligible interference with traditional narrowband RF systems. Unlike narrowband transceivers, highly frequency selective circuits are unnecessary to facilitate the integration of the entire transceiver. Also, the effect of the multipath can be mitigated, and even exploited by measuring the arrival time and the phase of the multipath signals. For this reason, the untra-wideband technology is more suitable for short-range, fixed environment communication than the application that has been in use. The host-to-display, chip-to-chip, and backplane communication can benefit from the ultra-wideband communication because they are typically short-range, fixed environment communication. The short-range nature of the host-to-display,

chip-to-chip, and backplane communication could provide a reasonable signal-to-noise ratio, which combined with ultra-wideband, would provide potentially very high data rate required in such data communication. Also, the host-to-display wireless link has an added possibility of broadcasting to multiple displays.

Our focus is an UWB receiver in which most signal processing is performed in the digital domain. This requires an A/D converter with extremely high sampling rate (> 10 GHz) and moderate (6 bit) resolution. There are several options to achieve such high sampling rates, one of which is time-interleaved analog to digital converters. The speed increase is achieved by placing several converters in parallel. This method requires that the individual converters, which make up the parallel combination, be matched. Mismatches in non-idealities, such as gain error, timing error, and voltage offset, greatly degrade the performance of such systems. Calibration is often used to reduce these mismatches.

This research will focus on using digital signal processing techniques to perform background calibration on the individual converters. Many converters use several input references ($V_{max'} V_{gnd'}$ and a ramp function) to measure the gain error, offset error, and timing error. We are investigating an alternative calibration waveform to measure all of these errors that can be derived in one pass. For example, a single frequency sine wave calibration tone when digitized by individual channels and Fourier transformed will produce all necessary measurements. These non-idealities can be directly taken from the amplitude, offset, and phase of the digitized sin wave. Furthermore, by using spread spectrum modulation techniques, the calibration can be performed in the background.

Wireless Gigabit Local Area Network

Personnel A. Chandrakasan, H.-S. Lee, and C.G. Sodini

Sponsorship

Center for Integrated Circuits and Systems (CICS)

The exploding number of electronic devices or "appliances" requiring high bandwidth communication will continue to drive the need for higher speed (Gigabitper-second, Gb/s) networking. We assume that the Next Generation Internet (NGI) will carry high-speed data to and from the home or office. However, a Local Area Network (LAN) within these structures is necessary to continue high-speed data transmission to and from end-use devices, such as cameras, displays, printers, high resolution video, mobile communicators, and novel devices. The enabling technology for this rich set of applications is a *wireless Gb/s LAN*, (WiGLAN), connected to the NGI.

The WiGLAN offers several research challenges. First, there is a wide range of data rates, quality of service, and need for real time transmission to and from the appliances. For example, voice transmission over the network will not require high data rates, but may require low power dissipation for portability. Interactive video transmission requires real time transmission and very high data rates especially as high resolution video and 3D graphics become available. System resources will need to be adaptive in order to support this wide range of appliances. Second, since many of the appliances will require portability, low power design techniques at the circuit, chip architecture, and overall system level will be required. Third, this research requires synergy between a variety of disciplines including, communication system design at the physical layer, low power circuit and system design, digital signal processing algorithm and IC design, mixed signal IC design, and RFIC design. It also lends itself to a number of demonstration projects using some of the technology which results from this research. Besides the educational component of the PhD researchers directly involved, this program will generate a number of IC's and algorithms which can be demonstrated by Masters student design projects.

A block diagram of the Wireless Gigabit Local Area Network, WiGLAN, is shown in Figure 14. We envision a network server being the gateway between the NGI and the local area network. Each appliance is attached to the network through a WiGLAN adapter which is capable of providing a wireless connection to the network. This adapter should be physically small, implying a high degree of integration of the electronic functions required to interface digital data from the appliance to and from the network. The quality of service, QoS, which is a function of data rate and bit error rate, should be scaleable with power dissipation to permit battery operation of many appliances.

The network requirements of high bandwidth efficiency and real time transfer led to our choice of a multi-carrier modulation, such as Orthogonal Frequency Division Multiplexing, (OFDM) using M-Quadrature Amplitude Modulation, (MQAM) signal constellations. We plan to digitize the entire signal bandwidth (150 MHz) available at the 5.8 GHz ISM band and adapt the bit rate (change M) within sub-bands according to the available Signalto-Noise Ratio (SNR) and interference in the sub-band. A programmable digital signal processor will perform this adaptive modulation.

The adaptive bit rate processor located in the network server will estimate the channel capacity by measuring the SNR and interference within sub-bands across the entire 150 MHz signal band. The channel estimation algorithm is a subject of this research. Depending on the SNR and interference, data modulation will range from simple Phase Shift Keying (PSK) up to 256 level QAM with intermediate levels of QAM, (i.e. 4-QAM, 16-QAM, etc.) allowing for transmission of approximately 1b/Hz for PSK up to 8b/Hz for 256-QAM.

In order to provide the capacity enhancements required to support the target data rates, the system to be developed will make extensive use of multiple-element antenna arrays for both transmission and reception. A key component of the proposed research will therefore be the development of computationally and power efficient space-time coding and space-time processing algorithms that exploit the substantial diversity benefit inherent in the use of such antenna arrays. At the implementation level, multiple-element antenna arrays require a separate receive and transmit channel for each antenna element. To efficiently meet this requirement, we propose to build a system of parallel radios divided into three distinct Integrated circuits, namely RF, Mixed signal, and DSP.

The WiGLAN network adapter consists of three functions: digital signal processing for multi-carrier adaptive bit rate QAM, a baseband analog processor performing data conversion and filtering, and an RF transceiver function which interfaces the modulated baseband data to a 5.8 GHz carrier. We will design and characterize integrated circuits to perform these functions.



Fig. 14: Wireless Gigabit Local Area Network

5.8 GHz Wideband Receiver for Wireless Gigabit LAN

Personnel L. Khuon (C. G. Sodini)

Sponsorship

Center for Integrated Circuits and Systems (CICS) and SRC

To take advantage of space-time diversity algorithms, multiple receiver front ends are needed on a single chip. Direct conversion does not require an image reject filter and simplifies the Radio Frequency (RF) filtering requirements. The nature of multiple receivers on chip, however, implies that the homodyne's local oscillator radiation would significantly interfere with nearby receivers since its frequency is in-band to the desired RF signal. In addition, a direct conversion receiver performs In-phase and Quadrature (I/Q) demodulation in the analog domain. This results in an I/Q phase imbalance, directly impacting the bit-error-rate performance. With a heterodyne architecture, the received signal could be digitized at a low IF and the functions of I/Q demodulation along with channel selection could be performed in the digital domain.

The receiver for the WiGLAN performs amplification, filtering, and downconversion of the 150 MHz signal centered at 5.8 GHz. The receiver downconverts the Radio Frequency (RF) signal to a low Intermediate Frequency (IF) that is fed to the analog baseband processor where it is equalized and digitized. The design approach for the receiver is based upon block level analyses that consider the gain, noise, and linearity tradeoffs necessary for the WiGLAN's adaptive modulation scheme.

The focus of this research is the design of on-chip filters within the framework of the Wireless Gigabit LAN (WiGLAN) receiver design. To reduce the effect of image frequencies for the heterodyne receiver, dual conversion architecture is selected to allow for optimized frequency planning. As such, filters are needed for the band selection and image rejection at the RF, band selection at the first IF, and anti-aliasing at the low IF. The primary challenges are to obtain the necessary band filtering and image rejection with an integrated approach without severely degrading the system's noise and linearity performances. An initial integrated image reject filter was fabricated on IBM BiCMOS 7HP process technology (See Figure 15). The filter incorporates an on-chip inductor that has its quality factor enhanced through the use of a negative resistance circuit. The filter's center frequency is tunable externally with a DC voltage. In addition, by controlling the DC current of the negative resistance circuit, the rejection response is also adjustable. This simple notch filter circuit only performs a rejection at the image frequency. To perform rejection of a band of frequencies, higher order filters are necessary, and this initial circuit may serve as the building blocks for more complex responses for both the image band rejection and the signal band selection at the RF and IF stages. Besides the impact on the receiver's noise and linearity performances, the design of integrated filters must also consider issues of stability and possible automatic frequency response adjustments to account for device tolerances.



Fig. 15: Die photo of an LNA with notch filter for image rejection on IBM BiCMOS 7HP process.

Analog Base-band Processor for Wireless Gigabit LAN

Personnel M. Spaeth (H.-S. Lee)

Sponsorship

SRC and Center for Integrated Circuits and Systems (CICS)

The base-band analog processor performs necessary signal processing on the 150 MHz base-band signal in the Transmit (Tx) and Receive (Rx) signal paths of a wide-band wireless local area network. There are tremendous technical challenges in the development of this base-band analog processor due to the high data rates and complex modulation schemes employed in a wireless network. The analog circuits in both the transmit and receive sections of the processor must handle 150MHz of signal bandwidth with signal-to-noise ratio in excess of 75dB (12 bits). In the receive section, these circuits include a low-noise wide-band anmplifier, a programmable gain-amplifier, an anti-alias filter, a channel equalization filter (if required), and finally an A/D converter.

This work focuses on the implementation of the extremely high speed, high resolution, and widebandwidth A/D converter in the Rx section of the base-band analog processor. In order to digitize the 150 MHz-wide signal band, the A/D converter must have an effective sampling rate above the Nyquist frequency of 300 MHz. To ease the anti-alias and digital filtering requirements, a sampling frequency above twice Nyquist will be used. The preliminary estimate of the A/D converter resolution needed to handle the wide dynamic range of the received signal is 12 bits. Additionally, any harmonic and intermodulation distortion in the signal path produces spurious signals in other sub-bands, so the A/D converter must exhibit very high spurious-free dynamic range (SFDR) in addition to wide bandwidth. At present, such high performance is beyond the capability of monolithic silicon integrated circuits.

To achieve high performance operation, some degree of parallelism is often employed. In a parallel timeinterleaved converter, any mismatch in the gain, offset, or timing of the constituent channels results in undesirable harmonics in the output spectrum. Therefore, the time interleaving schemes commonly used today employ a small degree of parallelism, so that the harmonics lie either out of the signal band of interest or below the quantization noise floor. Our approach is to use large-scale parallelism (128 active channels) in a time-interleaved pipeline A/D converter. Back-end digital calibration is applied to account for static gain, offset, and timing mismatch errors between channels, so that the resulting calibrated output has sufficiently low spurious harmonics.

Measurement and calibration techniques for gain and offset errors are performed using standard calibration techniques. By digitizing a fast ramp using one converter as a fixed timing reference for the remaining converters, the relative timing skew between channels can be discerned. The calculated timing offsets are then used to re-time the output data stream using polynomial interpolation in the DSP in the back-end. Thus, all of the calibration is performed using simple algebraic operations with minimal latency. To allow all of the calibration operations to be performed in the background, a small fraction of the available channels are systematically pulled out for calibration, while a novel token-passing control scheme selects which of the 'active' converters will sample the incoming signal.

Figure 16 shows a top-level block diagram of the proposed A/D converter. 129 identical pipeline A/D channels are organized into 16 banks of 8 converters, with one additional converter used only as a skew timing reference. In this scheme 2 banks are pulled out at a time for calibration, so the remaining 112 converters operate at about 5.5 MHz to achieve the desired 600MHz aggregate sampling rate. 14 bit pipelines are used to generate 12 bit digitally error-corrected outputs. The converter bank that is actively digitizing the input signal receives the output of the front-end anti-aliasing filter. The converter banks that are under calibration may digitize DC values for gain

continued

and offset measurements or the fast ramp for timing skew measurements. The converter has two sets out outputs so that digitized signal samples and calibration data may be output simultaneously. The back-end DSP averages the calibration data, and generates the algebraic coefficients needed to correct the gain, offset, and timing mismatch errors.

While the infrastructure generate calibration data is on-chip, the sample output from the IC is raw and uncalibrated, with the calibration occurring offchip. This split architecture allows for other novel schemes to be employed, taking advantage of the massive amount of data being output to resolve higher resolution, but lower data rate samples. In addition to straight forward calibration, non-uniform sampling methodologies and oversampling techniques are also being explored, so that the core IC design can be more intelligently applied to a broader range of applications.



Effect of Circuit Nonlinearity on System Performance Metrics, BER and Spectral Efficiency

Personnel

F. Edalat (C.G. Sodini)

Sponsorship

Center for Integrated Circuits and Systems (CICS) and The National Defense Science and Engineering Graduate Fellowship

As the demand for RF spectrum increases, high-speed data transmission over radio channel is likely to benefit from high bandwidth efficiency obtained from transmission of multi-carrier signals (OFDM or TDM) with Multilevel Quadrature Amplitude Modulation (M-QAM). In addition, from the circuit point of view, power amplifiers are more desirable to operate at high power levels than at low power levels for improved power efficiency. Higher efficiency means that a larger percentage of the dc (e.g. battery) power is delivered to the load. However, at high power levels, power amplifiers become nonlinear. Since the envelope of a QAM-modulated signal is not constant, it requires a high level of linearity of the transceiver components to achieve an acceptable performance.

Traditionally, to operate in the linear region, a large back off from the 1-dB compression point of the power amplifiers was suitable. This leads to low power efficiency and high power consumption, which cannot be tolerated in portable wireless systems. A higher power efficiency of a power amplifier can be obtained at the expense of more nonlinear distortion. In other words, as the input power enters the nonlinear region of power amplifier, power efficiency increases while the performance degrades. In addition, nonlinearity introduces a degradation of spectral efficiency. As a consequence, for a given nonlinear power amplifier, we would like to find the maximum input power - in the nonlinear region of the amplifier - that achieves BER below the maximum target BER and spectral efficiency above the minimum allowable spectral efficiency. Such information is extremely valuable to circuit and system designers, facilitating the collaboration between the two, and hence leading to more efficient practical design solutions in a shorter amount of time. However, obtaining such information requires knowledge of how circuit nonlinearity parameters affect system performance metrics, BER and spectral efficiency, which is the focus of this research.

To investigate the effect of circuit nonlinearity on the system performance, a communication system with a nonlinear power amplifier model is simulated. The transmitter of the communication system consists of an M-QAM random source, creating a sequence of symbols chosen randomly from the signal constellation. The symbol sequence is converted to a continuous-time waveform by a square-root-raised-cosine filter, and then added to an Additional White Gaussian Noise (AWGN) characterizing the circuit noise introduced by other transmitter components, such as the VCO and mixer. The resulting signal is the input to the nonlinear power amplifier model, and its output goes to the AWGN channel. At the receiver, the received signal is passed through another square-rootraised-cosine filter, which together with the filter at the transmitter constitutes a Nyquist Filter that is required to achieve zero inter-symbol interference. The filter output is sampled at the optimum sampling point, and subsequently detected by a Minimum-Distance optimum detector. The detected symbols are then compared with the transmitted symbols, and the number of symbols in error is calculated. Throughout the simulation, the signal is maintained at the baseband, and the effect of upconverting it to passband for the purpose of our study is taken into account by characterizing only the nonlinearity effect of power amplifier in the band of interest around the carrier frequency. The result of such simulation is directly affected by the accuracy of the model used to characterize the nonlinearity of the power amplifier. Our model is obtained by examining the transfer characteristic data of an experimental class-A power amplifier fabricated in IBM SiGe 7AP BiCMOS process. Once the best nonlinearity model is found, insightful results can be produced by this simulation, including the performance degradation level, BER, as a function of input back-off from saturation, or the optimum modulation level to use in order to operate at a required power level with a maximum tolerable BER.

On-Chip Cross-Talk Analysis for an Array of Transceivers on a Single Chip

Personnel J. Liang (C.G. Sodini)

Sponsorship

MARCO Focused Research Center on Circuits, Systems, and Software (C2S2) (MARCO/DARPA) and Center for Integrated Circuits and Suystems (CICS)

The Wireless Gigabit Local Area Network (WiGLAN) project is using multiple antennas to increase capacity through Space-Time Coding (STC). Such a system can benefit from integrating multiple RF front ends on a single chip. Each front-end analog circuit consists of a mixer, filtering, and power amplifier at the transmitter side and an LNA, mixer and filtering at the receiver side. In addition, local oscillators at specific frequencies must be synthesized and applied to the mixers. Along the chain of each front end, there are multiple nodes that signal cross-talk can occur.

Such cross-talk between these parallel radios can severely degrade the system performance, imposing major challenges for integration. How much overall-systemperformance degradation does the cross-talk cause? How much can careful circuit-design contribute to cross-talk suppression? How to determine the adequate level for cross-talk suppression? This study will attempt to answer these questions.

The focus of the research is to quantify the effect of cross-talk upon the overall system performance. In particular, a study of the specific nodes which are vulnerable to cross-talk along the radio front ends and the required isolation level for various modulation schemes will be carried out. Two types of signal crosstalk can occur when the WiGLAN system operates. They are the On-Chip Signal Cross-Talk and the Spatial Channel Signal Cross-Talk. The cross-talk level of the first type is related to the particular circuit and system design techniques used. The level of the second type is determined by the characteristics of the space channel or the environment where the system operates. We usually have little control over the characteristics of the space channel. One may ask, since there is already cross-talk in the space channel, why should we be concerned about the on-chip signal cross-talk? The reason lies in that the on-chip signal cross-talk stays almost constant or varies very little with respect to time while the spatial channel signal cross-talk varies randomly in time. The effectiveness of STC depends on the randomness of the overall channel, which includes both the on-chip channel and the space channel. The deterministic nature of the on-chip channel limits the randomness of the overall channel. Therefore, the on-chip signal crosstalk can undermine the overall system performance once it exceeds a certain level. In order to model the WiGLAN correctly and ensure flexibility in quantifying the cross-talk effects, we will make both the modulation coding scheme and the performance characteristics of each circuit block adjustable. Thus, we will be able to quantify the cross-talk level of systems of different digital modulation schemes and different individual circuit performance.

Linear Power Amplifier Design for WiGLAN

Personnel

A. Pham (C.G. Sodini)

Sponsorship:

MARCO Focused Research Center for Circuits, Systems and Software (C2S2) (MARCO/DARPA) and Center for Integrated Circuits and Systems (CICS)

The goal of this research is to build a power amplifier suitable for the WiGLAN (Wireless Gigabit per second LAN) project. In order to get such a high throughput, multiple orthogonal n-QAM modulation channels are used which require power amplifiers with extremely high linearity. However, conventional power amplifiers usually have to trade efficiency for linearity. There are two main sets of solutions to overcome this limitation. First, linearization techniques can be used to boost linearity of highly efficient amplifiers. Examples include envelope-error-restoration, Cartesian-loop feedback, ... Second, adaptive biasing techniques can be used to improve efficiency of linear power amplifiers. The power amplifier proposed in this project uses an adaptive current biasing technique as described below.

Due to their high bias conditions, conventional class A and AB power amplifiers are very inefficient, especially at backed-off power levels. For a conventional fixed-bias amplifier, the dc supply power is constant across the output range at $P_{dc} = V_q I_q$. The efficiency is highest at the maximum output power where both the current and voltage achieve their maximum swings. At lower output levels, the current and voltage amplitudes are reduced. However, the dc supply voltage and current remain unchanged. Therefore, the efficiency decreases rapidly as the output power is reduced as shown in Figure 17.



Fig. 17: (a) Typical Efficiency Curve (b) Fixed-Bias Waveforms (c) Adaptive Current Biasing

To improve efficiency at low output levels, the bias current is adjusted dynamically based on the input signal level. When the output power is reduced, the efficiency can be improved if the bias current (I_q) can be adjusted so that the current swing is to the edge of clipping for every input power level as shown in Figure 17c. Essentially, the bias circuit is a voltage controlled current source, which consists of an input level detector, a rectifier, and an averaging circuit as shown in Figure 18.



Fig 18: Adaptive Current Biasing Block Diagram

The amplifier is designed for the 5.8 GHz UNII-band and fabricated using the IBM SiGe 7HP BiCMOS process. Using the proposed adaptive current biasing technique, the amplifier exhibits significant improvement in efficiency at low output power while maintaining good linearity. The efficiency at low power levels increases as much as two times compared to the fixed-bias version. At maximum linear output power of 19.2 dBm, the power amplifier shows 30% Power-Added-Efficiency (PAE), and 16.45 dB gain with -35.46 dBc Adjacent Channel Leakage Power Ratio (ACPR) operating at a supply voltage of 2.5V.

Smart Active-Matrix Display Drivers For Organic Light Emitting Devices

Personnel M. Powell, J. Yu (V. Bulovic and C. G. Sodini)

Sponsorship

MARCO Focused Research Center for Circuits, Systems and Software (C2S2) (MARCO/DARPA) and Center for Integrated Circuits and Systems (CICS)

In this project we are developing pixilated active matrix "smart drivers" for displays consisting of organic light emitting devices (OLEDs). Organic LEDs are perhaps the most promising novel technology for development of efficient, pixilated, and brightly emissive, flat-panel displays. They naturally emit over large areas, and offer the advantage of growth on lightweight and rugged substrates such as metal foils and plastic, with no requirement for lattice-matching.

Organic LED devices, however, exhibit non-liner light output responses that complicate their implementation in an application requiring a fine control of the output light intensity. Specifically, the I-V characteristics of OLEDs depend on the cathode/anode type, device layer thickness, and operating temperature. The power efficiency of pixels in a display will drift over time due to operational degradation. The individual pixels in a display can then exhibit different aging, in accordance with their use. The brightness non-uniformities due to the differential aging will reduce the useful display lifetime.

Our smart active matrix circuitry compensates for the OLED non-uniformities by monitoring light output and adjusting the driving conditions according to the OLED performance. The adjusted output provides a defect-free picture. In the final design a Si p-n detector integrated behind each pixel will give feedback to the



Fig. 19 (left) An OLED pixel integrated with a "smart" si active matrix driver. The si photodetector monitors the intensity of the OLED pixel during the on state and provides feedback to the driving circuit to keep the light output intensity constant as the device efficiency changes with operation. (Right) pixel design in integrated circuit implementation

driver circuits that will adjust the proper current level to derive a constant brightness output. Figure 20 shows a typical integrated structure in which our patented transparent OLED is used. In this design both OLED electrodes are capable of transmitting the emitted light which is mostly observed on the top, but is also partially absorbed in the detector. The integrated circuit layout of the mini-display is shown in Figure 21. Notice that six transistors control each pixel (Figure 20). Also, each column shares one feedback circuitry. The integrator type of compensation is used for each feedback circuitry to ensure that the light output is matched to the reference input. The values of discrete components were so chosen to stabilize the feedback loop.

The present state of the art of OLED display technology uses a constant current to drive an OLED pixel. In this driving scheme even the most efficient of OLEDs will drop their luminescent output to 90% of the 100 Cd/m² initial brightens in ~ 5000 hours. As a human eye can distinguish brightness change of less than 10%, the 90% operating point indicates the longest useful lifetime of a pixel in a display.

With the circuit developed in this project we compensate for the loss of brightness of an aging OLED pixel by increasing the operating (driving) voltage as



Fig. 20: 28 x 16 pixel integrated circuit layout.

a function of time. The lifetime of such compensated pixel is now primarily limited by the maximum voltage that the driving circuit can deliver. From the data of Figure 22 we project that for the maximum driving circuit voltage of 10V, the constant pixel brightness can be sustained for 30,000 hours by doubling the initial drive current, and for the maximum driving voltage of 12V, the constant pixel brightness can be sustained for 50,000 hours by tripling the initial drive current. Such long projected lifetimes would enable the use of OLEDs in commercially viable displays.



Fig. 21: Driving circuit voltage increase the proposed display driving scheme that compensates for the aging of an OLED.



Fig. 22. α nd β crystal forms of Alq3. (from Brinkmann, et al., J. Am. Chem. Soc., <u>122</u>, 5147 (2000).

A CMOS-Compatible Compact Display

Personnel A. Chen (A.I. Akinwande and H.-S. Lee)

Sponsorship

3M and MARCO Focused Research Center on Integrated Circuits and Systems (C2S2) (MARCO/DARPA)

The proliferation of portable electronic systems has created demand for high-resolution displays which are compact and highly energy-efficient. We have designed and built a proof-of-concept for a display that meets these design constraints. Our display uses a standard digital CMOS integrated circuit to produce a low-brightness image, and an image intensifier to increase brightness to a visible level. Exploiting high level of integration achieved by the CMOS IC, low power techniques such as pixel memory and data compression can be implemented to lower the system power consumption. We are exploring the use of high-accuracy calibration techniques for the display driver circuits. We are designing the driver circuit for compatibility with other emissive display technologies such as organic LEDs so that a single driver circuit can be employed in multiple display technologies. A display using our design should produce a daylight-visible image using approximately half a watt of power.

Silicon devices can convert electrical energy into light, although their efficiency is very low. We use silicon light-emitting diodes to produce a very faint image which is optically coupled into an image intensifier. The image intensifier is a compact vacuum device that uses cathodoluminescence to increase the brightness of an image. It is commonly found in night vision scopes and scientific equipment. The intensifier in principle can be built very compactly by using multiple channel plates with MEMS technology. Cathodoluminescence, using a phosphor to convert electrons to photons, is an established technology used in cathode-ray tubes. Cathodoluminescent devices have high conversion efficiency (40 lumens/watt), high reliability, and can achieve very high output brightness (projection televisions).

We produced a laboratory demonstration of the system. An integrated circuit with light-emitting arrays was fabricated in a commercial 0.18μ m CMOS logic process. Each array measured 16x32 pixels and included a wordline decoder. Each pixel contained a 1-bit digital memory along with light emitter and driver circuits. Sample images were recorded, as shown below.

We are exploring circuit designs to support the integration of light emitters onto CMOS integrated circuits. Memory can be added to the display to eliminate the need for refreshing, thus reducing switching power. In addition, row parallel current level addressing is being



Fig. 23: Circuit board including IC in 0.18µm CMOS process.



Fig. 24: 32-level grayscale image from test system captured with CCD camera.

continued

investigated. Calibration techniques are used to overcome manufacturing process variation and allow precise brightness control of each pixel. In addition, we believe our circuit designs will be capable of supporting multiple emissive display technologies. Our current target is a silicon backplane which can drive (1) silicon light emitters with image intensifier, and (2) an organic LED.



Fig. 25

A Programmable, Wide Dynamic Range CMOS Imager with On-Chip Automatic Exposure Control

Personnel P. M. Acosta Serafini (C. G. Sodini and I. Masaki)

Sponsorship

Intelligent Transportation Research Center (ITRC)

Machine vision applications which use visual information typically need an image sensor able to capture natural scenes which may have a dynamic range as high as four orders of magnitude. Reported wide dynamic range imagers may suffer from some or all of these problems: large silicon area, high cost, low spatial resolution, small dynamic range increase factor, poor pixel sensitivity, small intensity resolution, etc. The primary focus of the proposed research is to develop a singlechip imager for machine vision applications which addresses these problems, but is still able to provide an ultra wide intensity dynamic range by implementing a pixel-by-pixel automatic exposure control. The secondary focus of the research is make the imager programmable, so that its performance (light intensity dynamic range, spatial resolution, light intensity resolution, frame rate, etc.) can be tailored to suit a particular machine vision application.

The imager sensing array has pixels which can be independently read and reset. The proposed brightness adaptive algorithm then predictively scales the voltage in photodiodes that would saturate under normal circumstances based on information gathered in several readout checks. The total integration time is subdivided into several integration times (called integration slots), which are progressively shorter. If in any of the checks it is determined that the pixel will saturate at the end of the current integration slot, then the pixel is reset and is allowed to once more integrate light, but for a shorter period of time. Each pixel has a small associated memory location needed to store an exponent which identifies the actual integration slot used. This information is used to appropriately scale the digitized pixel output.

A proof-of-concept integrated circuit has been fabricated in 0.18um 5M CMOS process shown below in Figure 26. It includes a 1/3'' VGA (640x480) array (7.5 μ m square pixels), 64 cyclic analog-to-digital converters for digital pixel output, an integration controller which implements the described algorithm, 4-bit per-pixel SRAM memory for exponent storage and supporting digital logic.



Fig. 26: Die photo of programmable wide dynamic range CMOS image.

Mixed-Signal Design in Deeply Scaled CMOS Technology

Personnel

J. Fiorenza (H.-S. Lee and C. G. Sodini)

Sponsorship

Center for Integrated Circuits and Systems (CICS) and MARCO Focused Research Center for Circuits, Systems and Software (C2S2) (MARCO/DARPA)

There are tremendous challenges in implementing mixed-signal systems on a single substrate in deeply scaled CMOS technologies primarily due to the negative impact of the technology on analog circuits. Nearly every aspect of scaling except speed goes against analog circuits. Lower power supply voltage severely restricts the signal range, requiring substantially lower circuit noise in order to keep the signal-to-noise ratio. Small geometry transistors exhibit far less voltage gain and greater threshold voltage mismatches than their predecessors. Attempts to overcome device gain limitations with conventional techniques such as cascode and regulated cascode aggravate already slim signal swing. The use of long-channel devices for higher gain inevitably compromises the circuit speed.

In order to overcome the challenges, we are exploring innovative circuit techniques that avoid shortcomings of deeply scaled technologies, and actually exploit them in mixed-signal systems. As the first step we have been investigating circuit techniques that overcome the device gain limitations without penalizing the signal swing or circuit speed. An innovative approach that we have developed employs two signal paths: the main path and the prediction path. The prediction path processes the signal 1/2 clock phase earlier than the main path at a reduced accuracy. The information obtained from the prediction phase is used in the main path in order to compensate for the finite device gain, incomplete settling, and other non-idealities. The two-path approach can be applied to many different classes of analog circuits including data converters, filters, instrumentation amplifiers, and many others. Compared with previous techniques of effective gain enhancement, the proposed technique incurs little penalty in power consumption - an important measure often ignored in the literature. As the initial proof-of-concept, we designed a MOS sample-and-hold amplifier in a standard 0.18 μ digital CMOS process. The simulation predicts the accuracy corresponding to 100dB amplifier gain with no cascading. The chip has been fabricated and shown to be functional, and is currently undergoing performance evaluation.

A CMOS Bandgap Current and Voltage References

Flicker Noise in Scaled CMOS Devices

Personnel M.C. Guyton (H.-S. Lee)

Sponsorship

Center for Integrated Circuits and Systems (CICS)

Most analog circuits require reference voltages and currents that do not vary with power supply voltages and temperature. Bandgap voltage references with an output voltage around 1.2 volts have been popular for this purpose. However, producing current sources referenced to bandgap voltage requires an operational amplifier increasing the complexity and power consumption.

The focus of this research is to develop simple and low power bandgap current references. We have developed a novel bandgap core circuit that produces a bandgap referenced output current directly without an operational amplifier. This simple circuit can even be operated as a 2-terminal bandgap current source. The same core circuit can also be used to generate arbitrary non-integer multiples of bandgap voltage.

A prototype 2-terminal band-gap current source has been designed and fabricated employing only 4 MOS transistors and 2 parasitic PNP transistors in a standard 0.35µ CMOS technology.

This chip shows full functionality. At a nominal 2V supply voltage and 80 μ A output current, the entire circuit dissipates 160 μ W, requiring no excess power consumption other than that of the current source itself. The measured output resistance is 350 k Ω . Since the design was performed without prior temperature characterization of components such as the bipolar transistors and resistors, the trimming range was found to be inadequate for achieving minimum temperature coefficient. Over 5°C – 50 °C temperature range, the temperature coefficient is 142 ppm/V.

Personnel

T. Sepke (H.-S. Lee and C. G. Sodini)

Sponsorship

MARCO Focused Research Center for Circuits, Systems and Software (C2S2) (MARCO/DARPA), Center for Integrated Circuits and Systems (CICS), Maxim Fellowship

Research of flicker noise in MOSFETs encompasses a large body of work spanning several decades. The number fluctuation model explains the source of flicker noise as the trapping of channel electrons in the gate oxide. Traps deeper in the oxide have longer time constants associated with them than traps closer to the channel, resulting in a non-white Power Spectral Density (PSD). A uniform distribution of traps in the oxide produces a PSD that is inversely proportional to frequency. The traps can also be distributed in energy in the oxide resulting in a gate voltage dependence. In addition to channel carrier number fluctuations, the trapped charges create an electric field, which results in mobility fluctuations that are correlated with the trapping events. A complete number fluctuation and correlated mobility fluctuation model was originally presented by Jayaraman and Sodini (IEEE Elec. Dev. 1989), and a simulation adapted form by Hung, Ko, Hu, and Cheng (IEEE Elec. Dev. 1990).

Assuming a simple scaling of devices (L/s, W/s, sC_{ox}), the number fluctuation model predicts no change in the flicker noise voltage at the gate with device scaling. One figure of merit that is of interest to circuit designers is the flicker noise corner frequency which is defined as the frequency where the flicker noise PSD is equal to the thermal noise PSD. Again using the simple scaling defined above, the number fluctuation model, and the square-law drain current model, it can be shown that the corner frequency increases proportional to the scaling factor *s*.

A simulation study has been performed comparing the noise of 0.25µm and 0.18µm transistors. The models predicted flicker noise corner frequencies as high as 100MHz. They also predict a dramatic decrease in the amount of flicker noise when transitioning from strong inversion to subthreshold. In order to investigate the flicker noise, a measurement system for on wafer devices has been designed and is currently being

Device Level Optimization of Phase Noise in Integrated LC VCOs

Personnel A. Jerng (C.G. Sodini)

Sponsorship

Center for Integrated Circuits and Systems (CICS)

implemented. Several devices of different sizes and under different bias conditions will be measured and analyzed. Of special interest are the weak to moderate inversion regimes, and bias points desirable for analog circuit designs.

Upon exploring the behavior of scaled devices, any generalizations or insights will be applied to an analog or RF circuit design application. The circuit design will be made with the goal of mitigating effects of increased flicker noise in future technologies.

Integrated LC Voltage-Controlled Oscillators (VCOs) are essential components in wireless systems. The Wireless Gigabit Local Area Network (WiGLAN) project aims to achieve a data rate of 1 Gb/s using 150 MHz of bandwidth in frequency bands allocated in the 5-6 GHz range. An adaptive M-ary modulation scheme, up to 256-QAM, is chosen to provide the required data rates, imposing stringent accuracy requirements on the Local Oscillator (LO) signal. VCOs with low phase noise and high operating frequencies are required. In addition, there is a desire to maintain compatibility with integration trends such as lower supply voltages. This research focuses on the design of a 5 GHz LC VCO based on 0.18 µm CMOS devices using a 1.8 V supply.

Phase noise analysis is complicated by the non-linear and time-varying nature of an oscillator. Recent research has identified bias circuit noise as an important contributor to phase noise. Upconverted flicker noise from MOS devices degrades close-in phase noise. Understanding the mechanisms through which circuit noise converts into phase noise is essential to the optimization of phase noise performance. Our design approach has been to develop models for the different phase noise conversion processes and to understand the relationship between device parameters such as g_m and f_t and the phase noise mechanisms. The resulting design intuition will allow us to treat and optimize the various phase noise sources independently, and understand the device-level tradeoffs involved in high frequency CMOS VCO design.



We have designed and fabricated an experimental set of seven VCOs using IBM's BiCMOS 7HP process technology (See Figure 27). Key device parameters were varied across the experimental set through choices of device type and device sizing, while keeping circuit parameters such

Fig. 27

Radio Frequency Digital-to-Analog Converter

Personnel

S. Luschas (H.-S. Lee)

Sponsorship

National Semiconductor Fellowship and Center for Integrated Circuits and Systems (CICS)

Dynamic performance of high speed, high resolution, DACs is limited by distortion at the data switching instants. Inter-Symbol Interference (ISI), imperfect timing synchronization, and clock jitter are all culprits. A DAC output current controlled by an oscillating waveform is proposed to mitigate the effects of the switching distortion. The oscillating waveform should be a multiple (k*fs) of the sampling frequency (fs), where k > 1. The waveforms can be aligned so that the data switching occurs at the peak and/or the valley of the oscillating current output. This makes the DAC insensitive to switch dynamics and jitter. The architecture has the additional benefit of mixing the DAC impulse response energy to a higher frequency. Instead of the conventional sinx/x DAC impulse response roll-off, there is a large high frequency lobe near the control oscillating waveform frequency (k*fs). An image of a Low Intermediate Frequency (IF) input signal can therefore be output directly at a high IF or Radio Frequency (RF) for transmit communications applications.

A narrowband sigma-delta DAC with eight unit elements was chosen to implement the RF DAC concept. A sigma-delta architecture allows the current source transistors to be smaller since mismatch shaping is employed. Smaller current source transistors have a lower drain capacitance, allowing large high frequency output impedance to be achieved without an extra cascode transistor. Elimination of the cascode reduces transistor headroom requirements and allows the DAC to be built with a 1.8V supply. The RF DAC is fabricated in a 0.18 μ digital CMOS process. Measured single-tone SFDR is 75dB and SNR is 52dB, while two-tone IMD3 is 70.8dBc over a 17.5MHz bandwidth centered at 942.5 MHz.

as supply current, voltage swing, and LC tank elements fixed in order to allow insightful comparisons of phase noise. The experimental data allows us to quantify the relative contributions of bias noise, MOS thermal noise, and MOS flicker noise to phase noise and to understand the impact of the device parameters on their contributions.

We have found that bias noise and flicker noise contributions to phase noise can be minimized through proper MOS device sizing. In addition, we have observed that 0.18 μ m PMOS devices exhibit better measured phase noise than 0.18 μ m NMOS devices due to lower drain thermal noise (See Figure 28). Our work has resulted in an optimized all-PMOS VCO topology demonstrating low voltage low phase noise performance at 5.3 GHz.



Fig. 28

Low Power RF Front-End for Wireless Microsensor Systems

Substrate Noise Coupling and Reduction Techniques in Mixed-Signal Systems

Personnel A. Y. Wang (C. G. Sodini)

Sponsorship

ABB and NSF Fellowship

The design of wireless microsensor systems has gained increasing importance for a variety of commercial and military applications. With the objective of providing short-range connectivity with significant fault tolerance, these systems find usage in such diverse areas as environmental monitoring, industrial process automation, and field surveillance.

The main design objective is maximizing the battery life of the sensor nodes while ensuring reliable operation. For many applications, the sensors need to "live" for 1-5 years without battery replacement. To achieve this goal, the microsensor system has to be designed in a highly integrated fashion and optimized across all levels of system abstraction. This also means that all the characteristics particular to the microsensor system must be exploited. One such characteristic is that the RF output power is small due to the short transmission distance, which makes the transceiver electronics the dominant source of energy dissipation.

In this research the impact of circuit non-idealities including noise, nonlinearity, and modulation errors upon system performance are analyzed, and these effects are incorporated into the design of key front-end components. In addition, the effect of increasing the RF transmit power, which is small, to compensate for the SNR loss due to circuit non-idealities is investigated. This can potentially lower the performance specification of the RF front-end and reduce the over-all power consumption.

Personnel

M. S. Peng (H.-S. Lee)

Sponsorship

Center for Integrated Circuits & Systems (CICS), MARCO Focused Research Center for Circuits, Systems and Software (C2S2) (MARCO/DARPA)

The demands of lower power, higher speed, and lower cost have driven the integration of all circuits in a system, analog and digital, onto a single chip. In this integration, one of the primary problems is substrate noise coupling. Digital circuits create noise, which couples into the sensitive analog circuits through the shared substrate. Improperly accounted for, this substrate noise can degrade analog performance drastically.

Up to now, most efforts in addressing this problem have been to ensure that analog circuits are robust enough to withstand the digital substrate noise. These techniques include physical separation, differential architectures, and simulation. Little effort has been placed on reducing the substrate noise itself.

With this in mind, the focus of this research is to investigate the characteristics of the substrate noise seen in analog circuits as well as ways to cancel the substrate noise. We have implemented a test chip that includes different digital circuits as substrate noise generators and a deltasigma A/D converter that samples the substrate noise as well as an external signal.

Substrate noise is characterized by observing the output of the delta-sigma converter built on the test chip to measure the substrate noise. By operating one digital circuit, an array of large inverters which mimics digital I/O drivers so that it produces a periodic noise waveform on the substrate, the delta-sigma modulator can be used as an accurate on-chip sampling scope to map the substrate noise as a function of time. The delta-sigma converter is ideal as the sampling scope because of the inherent averaging effect and 1-bit output that simplifies the interface. The sampling edge of the delta-sigma A/D converter is moved relative to the digital clock edge so that the substrate noise waveform can be reconstructed from the delta-sigma converter output. Figure 29 shows an example of the substrate noise generated by low-to-high and high-to-low inverter transitions and measured using this technique.

continued

In order to reduce the effect of the substrate noise, a feedback loop that shapes the substrate noise in bands of interest has been implemented and tested. This type of noise shaping is well suited for band-limited analog applications. While the concept is demonstrated in a low-pass system, the same principle can be used in band-pass systems.

The substrate noise shaping loop is based on a deltasigma modulator loop with the substrate noise treated as quantization noise. The feedback D/A is replaced by an array of transient-injecting inverters. This has the advantages of simplicity, low power, complementarity to existing substrate noise reduction techniques, and no restrictions on analog or digital circuit design. Figure 30 shows the substrate noise reduction seen in a delta-sigma modulator. The substrate noise is seen to be reduced by 20 dB when the substrate noise shaping loop is engaged. Signal to noise plus distortion ratio (SNDR) of a sampled-data analog circuit, in this case another delta-sigma converter, is increased by 10dB. Substrate noise was generated by large inverters which imitate the substrate noise generated by large I/O drivers. We believe further improvement is possible by carefully managing aliasing of the substrate noise in the noise shaping loop as well as in the sampled-data analog circuits.

The test chip with the substrate noise characterization system and the substrate noise shaping system runs at 2.5V and has been fabricated in conventional 0.25µm CMOS technology.



Fig. 29: Measured substrate noise with an on-chip sampling scope.

MTL Annual Report 2003



Fig. 30: Delta-sigma converter outputs with a 7.6kHz input signal. (a) Substrate noise generator off, substrate noise shaping loop off. (b) Substrate noise generator on, substrate noise shaping loop off. (c) Substrate noise generator on, substrate noise shaping loop on.

An Advanced Model Based Vision System for Intelligent Transportation Systems

Personnel

M. Kais (C. Laugier, M. Parent, I. Masaki and B.K.P. Horn)

Sponsorship

Intelligent Transportation Research Center (ITRC), Lounsbery Foundation, and INRIA

In order to offer better safety and increase the capacity of the roads, new concepts of mobility are being developed. These advanced transportation systems are based on autonomous driving and platooning. Platooning applications consist of creating a platoon of electronically coupled vehicles with a very small headway. In a platoon, the first vehicle is manually or automatically driven and the others follow. Another application is driverless fully autonomous electric vehicles called cybercars (See Figure 31) for transportation in urban environment.

These applications require sensors for the guidance and obstacle detection tasks. The goal of the vision system is to build an accurate representation of the environment in front of the vehicle to be used by the planning layer. Stereo vision offers a low cost and easy way to get some range information about the environment. Traditional vision systems are fine for the highway environments, because they are well defined (size, marker). However, urban environments present a challenging task due to the complexity of the scenes. For instance, in a single frame, the road and lane boundaries can take on several primitives, such as a curb, a white lane marker, or a line of cars parked on the side of the road.

Our approach consists in fusing information from the cameras with some a priori knowledge stored in a database. This a priori knowledge consist of a global model i.e. an environment model (how the roads are linked), a geometric model of the road and lane boundaries (shape) and a model of the road and lane boundaries feature (lane marker, curb, planes). This global model is acquired during a learning phase and linked to an existing geographic database using a Geographic Information System. Once the learning phase is completed, the estimate of the position of the vehicle and the a priori knowledge are used to enable some special-ized detectors in specific Regions Of Interest (See Figure 32) in order to extract from the images relevant information that is used for the 3D reconstruction.



Fig. 31: A cybercar



Fig. 32: Results from the lane marker detector

Efficient Traffic Monitoring

Personnel N.S. Love (I. Masaki and B. K. P. Horn)

Sponsorship

Intelligent Transportation Research Center (ITRC)

Traffic control centers use several methods to monitor traffic conditions. Currently, the most popular methods involve the use of cameras (image sensors) at highdensity traffic locations. The cameras are controlled at the traffic control centers. The traffic control center has on the order of 10 monitors and hundreds of cameras; each monitor cycles through a set of cameras while operators watch for any traffic incidents. This system is currently implemented using dedicated analog lines which have a low bandwidth. Bandwidth limitations inhibit the efficient transmission of network data. Consequently, the load due to a continuous transmission of images severely impacts the network's performance.

Our system reduces transmission load by distributing the processing of images to the image sensor and the control of transmission to mobile agents. Each image sensor processes each image and determines the contents of the image, and mobile agents decide if the image, traffic information, or nothing is sent to the traffic control center. The goal of this work is to reduce the transmission load of image sensor networks by distributing processing tasks to image sensors and reducing image transmission using mobile agents.

In the case of traffic monitoring, traffic images are processed at the control center to determine the average speed of vehicles or the number of vehicles that pass through a checkpoint during some time interval (traffic flow). Distributing the processing to the image sensors involves using an image sensor network to perform object recognition and image compression on the images at the image sensor before the image reaches the control center.

Providing select images to the user is achieved thru mobile agents. The control center dispatches mobile agents which search for images according to a user preset priority criteria. At an image sensor, each image is acquired and processed; the contents of the image are determined (number of vehicles, average speed of vehicles, whether there has been an accident or a sharp change in traffic conditions) and a priority is set to the image. The mobile agent from the control center checks when the image is updated and the level of priority of the image. The mobile agent decides if an image is sent back to the control center or information from the image based on the preset priority criteria. Figure 33 shows the components of the network and their interaction. By sending the mobile agent to intelligently decide the transmission of the image or traffic information, the transmission load is reduced.

Each sensor has a processor which acquires images and performs a three-dimensional contour based image compression algorithm on each image. The 3D image compression algorithm is a lossy compression method that retains three components: contour, color, and distance information. Each component can be used together or separately to aid in object recognition without fully decompressing the image.

Edge tracing determines the contours in the image. A modified differential chain coding method is used to further compress the contours. Differential chain coding codes the position of the first edge in the contour and the differential direction of the remaining edges in the contour. Differential chain coding does not follow contours that branch. We modified the differential chain coding to include branching for a more complete representation of the contour and to increase compression of the contour. The modification is a depth-first traversal of the contour with the addition of a marker to signify a split and a marker to signify a return to a split location. The additions of these two markers increase compression by eliminating the need to code start locations at the split. Start locations are costly to encode depending on the size of the image, the larger the image the more bits are needed to encode the start locations. Encoding each marker is 4 additional

bits. The savings of encoding the markers verses the start locations is seen in images that are larger than 64x64.

Finding the color on each side of the contour retains color information. To improve the quality of the decompressed image, the mean color of blocks between contours can also be included in the compressed image.

The distance information is obtained using a stereo vision algorithm. The cameras are aligned on a horizontal bar, and images are captured simultaneously from both cameras. Assuming the relative orientation (rotation and translation) between the cameras is known, the distance from one camera to an object can be determined by finding corresponding points in both images. Figure 34 shows the camera general setup of a stereo vision system.

Object recognition is performed on the compressed images to determine traffic flow and incident detection. The image sensor assigns a priority level to each image based on its contents (i.e. traffic congestion has medium priority, accidents have high priority, etc...).

The vehicles are detected by grouping features of each contour. The features used are depth, motion, and position. The vehicles are modeled with a multivariate Gaussian distribution using the Expectation-Maximization algorithm (EM), which is a maximum likelihood estimate, to approximate the scene as a mixture of Gaussian distributions. EM is a statistical clustering method which gives the mean, standard deviation, and weight of each cluster. Each cluster represents an object in the image. An object is detected by determining which cluster each contour belongs to and placing a bounding box around the contours in the same cluster. A contour is linked to a cluster with the minimum Mahalanobis distance. The Mahalanobis distance is a distance with each dimension scaled by the variance.

Once the vehicles are detected, the vehicles are tracked to determine the average speed, number of vehicles, traffic flow information, and incident detection. Once the image sensor gathers traffic information, mobile agents can determine whether the information is transmitted over the network.

Using the image sensor and mobile agents to complete processing tasks and to retrieve select images reduces the network's transmission load. For example, a police station dispatches a mobile agent to the cameras. The police station requests images with the criteria for high priority level accidents. The mobile agent will only retrieve those images with accidents. Transmission of traffic accident images versus all available images improves the network efficiency. The research develops a demonstration where mobile agents are sent with a given criteria to several camera locations where images are retrieved based on preset criteria.

Reduction of the transmission load will enable more users to obtain information without loss of performance. Distributed processing helps to minimize the transmission load by using the image sensors to complete normal processing tasks as opposed to processing at the control center. Mobile agents are equipped with the appropriate criteria to sift through the traffic information and to provide current traffic images and information to the user. The agents complement the image sensor network by providing select images based on criteria set by the user; using both mobile agents and image sensors, the network performance will be improved.



Fig. 33: Network Components

Fig. 34: Stereo Vision Setup

Sensor Fusion for Automobile Applications

Personnel Y. Fang (I. Masaki and B.K.P. Horn)

Sponsorship

Intelligent Transportation Research Center (ITRC) at MIT's MTL

To increase the safety and efficiency for transportation systems, many automobile applications need to detect detail obstacle information. Highway environment interpretation is important in Intelligent Transportation Systems (ITS). It is expected to provide 3D segmentation information for the current road situation, i.e., the X, Y position of objects in images, and the distance Z information. The needs of dynamic scene processing in real time bring high requirements on sensors in intelligent transportation systems. In complicated driving environment, typically a single sensor is not enough to meet all these high requirements because of limitations in reliability, weather, and ambient lighting. Radar provides high distance resolution, while it is limited in horizontal resolution. Binocular vision system can provide better horizontal resolution, while the miscorrespondence problem makes it hard to detect accurate and robust Z distance information. Furthermore, video cameras could not behave well in bad weather. Instead of developing specialized image radar to meet the high ITS requirements, sensor fusion system is composed of several low cost, low performance sensors, i.e., radar and stereo cameras, which can take advantage of the benefit of both sensors.

Typical 2D segmentation algorithms for vision systems are challenged by noisy static background and the variation of object positions and object size, which leads to false segmentation or segmentation errors. Typical tracking algorithms cannot help to remove the errors of initial static segmentation since there are significant changes between successive video frames. In order to provide accurate 3D segmentation information, we should not simply associate distance information for radar and 2D segmentation information from video camera. It is expected that the performance of each sensor in the fusion system would be better than being used alone.

Our fusion system introduces the distance information into the 2D segmentation process to improve its target segmentation performance. The relationship between the object distance and the stereo disparity of the object can be used to separate original edge map of stereo images into several distance-based edge layers in which we further detect whether there is any object and where the object is by segmenting clustered image pixels with similar ranges. To guarantee robustness, a special morphological closing operation is introduced to delineate vertical edges of candidate objects. We first dilate the edge to elongate the edge length so that the boundaries of target objects will be longer than that of noisy edges. Then an erosion operation deletes short edges. Typically the longest vertical edges are located at the object's boundary. The new distancerange-based segmentation method can detect targets with high accuracy and robustness, especially for the vehicles in highway driving scenarios.

For urban-driving situations, heavy background noise such as trees, etc., usually cause miscorrespondence, leading to edge-separation errors. The false boundary edge lines in the background area can be even longer than the boundary edge lines. Thus it is hard to eliminate false bounding boxes in background areas without eliminating foreground objects. The noisy background adds difficulties in segmenting objects of different sizes. To enhance the segmentation performance, background removal procedure is proposed. Without losing generality, objects beyond some distance range are treated as background. The pixels with small disparity represent the characteristics of the background.

Sometimes in assigning edge pixels to different edge layers, there exists ambiguity. Without further information it is hard to decide among multiple choices. Some algorithms simply pick one randomly, which might not be true in many situations. Typically, to avoid losing potential foreground pixels, edge pixels are assigned to all distance layers and edge-length filters can suppress ambiguity noise. However, when background noise is serious, algorithm picks only edge pixels without multiple choices. Eliminating pixels from the background in this way will lose significant pixels of target objects, making segmented region smaller than its real size. Thus, motion-based segmentation region expansion is needed to compensate for performance degradation. The original segmentation result can be used as initial object segmentation seeds from which larger segmentation boundary boxes will expand. The enlarging process is controlled by the similarity of segmentation seed boxes and surrounding edge pixels. With such region growing operations, the accurate target sizes are captured.

The proposed depth/motion-based segmentation procedure successfully removes the impact of background noise and captures objects of different sizes.

We presented a new sensor-fusion-based 3D segmentation algorithm to detect target distance and 2D location (See Figure 35). The system consists of following components: "distance-based edge layer separation," "background edge pixel removal," "target position detection," and "motionbased object expansion." The system firstly detects the rough depth range of all targets of interest. Then, we propose a new object segmentation method based on both motion and distance information. The segmentation algorithm is composed of two phases. "Distance-based edge-layer separation" and "background detection" are the first phase, which capture significant edge pixels of objects in interested distance layers while rejecting the noise from either background or other distance-based edge layers. Thus, original image edge map will be decomposed into several distance-based edge maps and heavy background noise can be removed. The advantage of this phase is that detecting targets sequentially in different edge maps is easier than segmenting all targets simultaneously in one busy edge map. The second phase is a new depth/ motion-based segmentation/expansion that can accurately capture objects of different sizes. With motion information for decomposed edge layers ("motion-based region expansion"), it further differentiates the target objects from noises in other distance layers, thus helping to detect objects of different sizes or to identify moving objects.

The algorithm successfully increases the accuracy and reliability of object segmentation and motion detection under the impact of heavy background noise. The algorithm can offer precise segmentation in detecting multiple objects of different sizes and non-rigid targets, such as pedestrians. The performance is satisfying and robust while computational load is low. This algorithm not only improves the performance of static image segmentation, but also sets up a good basis for further information tracking in video sequences. It shows that fusing stereo-vision and motion-vision algorithm helps to achieve high accuracy and reliability under the impact of heavy background noise.



Fig. 35: (a) *Highway environment interpretation.* (b) *Segmentation Result for Highway Environment.* (c) *Segmentation Result for Urban Driving Environment.*

Superconducting Bandpass Delta-Sigma A/D Converter

Personnel J. F. Bulzacchelli (H.-S. Lee and M. B. Ketchen — IBM)

Sponsorship

Center for Integrated Circuits and Systems (CICS)

The direct digitization of RF signals in the GHz range is a challenging application for any circuit technology. Traditionally, flash A/D converters have been used to digitize signal frequencies above 1 GHz, but their resolution and linearity are inadequate for most radio systems which must handle signals with a large dynamic range. Semiconductor bandpass delta-sigma modulators are used to digitize IF signals with high resolution, but their performance at microwave frequencies is limited by the speed of semiconductor comparators and the low Q of integrated inductors.

In this program, we present the design and testing of a superconducting bandpass delta-sigma modulator for direct A/D conversion of GHz RF signals. The schematic of the circuit is shown in Figure 36. The input signal is capacitively coupled to one end of a superconducting microstrip transmission line which serves as a high quality resonator (loaded Q > 5000). The current flowing out of the other end of the microstrip line is quantized by a clocked comparator comprising two Josephson junctions. If the current is above threshold, the lower junction switches and produces a quantized voltage pulse known as a Single Flux Quantum (SFQ) pulse. If the current is below threshold, the upper junction switches instead. The pattern of voltage pulses generated across the lower Josephson junction represents the digital output code of the delta-sigma modulator. These voltage pulses also inject current back into the microstrip line, providing the necessary "feedback" signal to the resonator. At the quarter-wave resonance of the microstrip line (about 2 GHz in our design), the resonator shunts the lower junction with a very low impedance; the "feedback" current to the resonator is maximized, and the quantization noise is minimized. Because of the high speed of Josephson junctions and the simplicity of the modulator circuit, the maximum sampling rate exceeds 40 GHz.

While such a high sampling rate improves the performance of the delta-sigma modulator, the challenges of high speed testing in a cryogenic environment are formidable. Even in the best cryogenic sample holders, the long cables used to connect the superconducting chip to room-temperature electronics have significant losses at frequencies above 10 GHz. Experimentally, we found two solutions for clocking the circuit at high frequencies. In the first approach (detailed in previous reports), we employ an optoelectronic clocking technique in which picosecond optical pulses at a 20.6 GHz repetition rate are delivered (via optical fiber) to an on-chip photodetector, the current pulses from which drive a Josephson clock amplifier. In the second approach, the modulator is triggered by an onchip clock source. An increase in bias current turns the Josephson clock amplifier into an oscillator tunable between 20 and 45 GHz. We found that surprisingly good frequency stability could be achieved with the on-chip clock source with careful adjustment of dc bias currents.

Since the modulator output data rate exceeds the capacity of the interface to room-temperature test equipment, on-chip processing of the data is used to reduce the bandwidth requirements for readout. As explained in the 1998 MTL report, two segments of the modulator's bit stream are captured with a pair of 128-bit shift registers. The number of clock cycles skipped between acquiring the two segments is set by an on-chip programmable counter (from 0 to over 8000). Cross-correlation of the two captured segments is used to provide estimates of the autocorrelation function R[n] of the modulator output, from n=0 up to a large value, such as n=8000. Fourier transformation of R[n] then yields a power spectrum with frequency resolution comparable to an 8K FFT of the original bit stream.

Figure 37 shows the block diagram of the modulator test chip. As mentioned above, the bandpass modulator can be clocked either externally by a 20.6 GHz optical source or internally by an on-chip Josephson oscillator. A 1:4 demultiplexer converts the single-bit output of the modulator to 4-bit words at one-fourth the sampling rate. This allows most of the test chip, including the programmable counter and the shift register memory banks, to operate at a reduced clock rate with larger timing margins. Because of the 1:4 demultiplexing, 128-bit memory banks A and B are organized as 4 parallel rows of 32-bit long shift registers. As just discussed, the number of clock cycles skipped between loading the A and B memory banks is set by a programmable counter which is programmed by external control currents. Once the shift registers have been loaded, a readout controller unloads the stored bits and transfers them to "high-voltage" drivers which amplify the output signals up to about 2 mV, which is large enough to be detected by room-temperature electronics. The test chip employs over 4000 Josephson junctions and represents one of the most complex circuits ever designed in this technology.

The test chip was fabricated at HYPRES, Inc. While the chip has been used with the 20.6 GHz optical clock, higher oversampling ratios and SNRs are attained with the on-chip clock source operating near 40 GHz. In the initial experiments, the programmable counter on the test chip was programmed so that the shift registers captured 256 consecutive bits from the modulator, so that 256-point FFTs could be calculated. The output spectra of the modulator at a sampling rate of 42.6 GHz is plotted in Figure 38. The width (about 500 MHz) of the input tone at 1.7 GHz reflects the low frequency resolution of the 256-point FFTs. The Full-Scale (FS) input sensitivity is -17.4 dBm (30 mV rms). Quantization noise is suppressed at 2.23 GHz and at higher frequencies corresponding to higherorder microstrip modes. The SNR (49 dB over a 20.8 MHz bandwidth) is limited by the frequency resolution of the measurements, but still exceeds the SNRs of semiconductor modulators with comparable center frequencies. Other measurements, based on the correlation technique discussed above, show that the in-band noise over a 19.6 MHz bandwidth is -57 dBFS. The center frequency and sampling rate of the experimental modulator are the highest reported to date for a bandpass delta-sigma modulator in any technology.







Circuit and System Level Tools for Thermo-Aware Reliability Assessments of IC Designs

Personnel

S. M. Alam (D. E. Troxel, K.E. Goodson, and C.V. Thompson)

Sponsorship

MARCO Focused Research Center on Interconnect (MARCO/DARPA)

Integrated circuits are often designed using simple and conservative 'design rules' to ensure that the resulting circuits will meet reliability goals. This simplicity and conservatism leads to reduced performance for a given circuit and metallization technology. To address this problem, we had developed a TCAD tool, ERNI, which allows process-sensitive and layout-specific reliability estimates for fully laid out or partially laid out integrated circuits (Y. Chery and S. Hau-Riege) (See Figure 39).

Circuit-level reliability analyses require reliability assessment of a large number of sometimes complexly connected interconnect trees. We have shown through modeling and experiments that the resistance saturation observed in straight via-to-via lines, which can lead to immunity from electromigration-induced failure, also occurs in more complex interconnect trees. We have also shown that trees will be 'immortal' if their effective current-density line-length product, (jL)_{eff}, is below a critical value. The jL product that defines immortality can be determined from experimental characterization or simulation of the reliability of straight via-to-via lines. Simple tests for tree immortality can be used in a hierarchical way to eliminate trees from further more computationally intensive reliability assessments. After filtering of immortal trees, the reliability of mortal trees must be assessed. This can be done through reliability simulations with individual trees, but this computationally intensive method should be reserved for the most problematic trees, those with the least reliability, and which are least convenient to 'fix' through layout modifications. We have suggested computationally simple and conservative 'default' models for assessment of tree reliabilities based on the Korhonen analysis and have tested models and simulations through experiments on simple interconnect trees.

Recent development in semiconductor processing technology has enabled the fabrication of a single integrated circuit with multiple device-interconnect layers or wafers stacked on each other. This approach is commonly referred



Fig. 39: A flowchart for a full hierarchical circuit-level reliability assessment, the basis for the prototype tool ERNI.

to as the Three-Dimensional or 3D integration of ICs. Although there has been some research on the impact of 3D integration on chip size, interconnect delay, and overall system performance, the reliability issues in the 3D interconnect arrays are fairly unknown. We have extended the reliability concepts in ERNI and developed a framework for reliability analysis in 3D circuits with a novel Reliability Computer Aided Design (RCAD) tool, ERNI-3D. Using ERNI-3D, circuit designers can get interactive feedback on the reliability of their circuits associated with electromigration, 3D bonding, and joule heating.

As the 3D integration technology is not yet widespread, and no CAD tool supports IC layouts for such a technology, we first developed a comprehensive 3D circuit layout methodology, the circuit on each wafer or device interconnect layer can be laid out separately with interwafer via information embedded in the layout. The interwafer via information is generalized into three categories sufficient for defining all types of interconnection between wafers in a 3D stack (See Figure 40). A strategy for layoutfile management that incorporates the orientation of each wafer in the bonding process is also proposed. We have implemented the layout methodology in 3D-MAGIC, an extension of MAGIC originally developed at UC Berkeley and widely used in academia. The test circuits designed with 3D-MAGIC are a 3D 8-bit adder and an 8-bit encryption processor mapped into a 3D FPGA.



Fig. 40: Different types of via/contact for 3D ICs.

The reliability CAD tool, ERNI-3D, parses 3D circuit layouts and extracts both conventional and 3D interconnect trees. It employs the Hierarchical Reliability Analysis approach, and filters out a group of immortal trees using their current-density length products. After the filtering process, the stringent reliability models are applied to the remaining interconnect trees to compute their median and mean time to failures. Finally, all the different time to failures are combined using a joint probability distribution to report a single reliability figure for the whole chip. This initial version of ERNI-3D treats 3D circuits with two wafers or device-interconnect layers in the stack. However, the data-structures and algorithms in the tool are generic enough to make it compatible with 3D circuits with more than two device-interconnect layers and to allow the incorporation of more sophisticated reliability models in the future.

As high temperature rise poses as a major challenge in stacked 3D ICs, we are currently working on circuit and system layout for thermal management and its impact on reliability. ERNI-3D provides an infrastructure for such development. A novel feature of this activity is the capability to guide optimal placement of microfluidic *thermal connects* (see Figure 41) at the layout-level. As a demonstration vehicle, we are focusing on structures of the type shown in the figure, in which device layers are bonded face-to-face (high density interconnects) and micromachined wafers are bonded back-to-back (low density through-wafer vias)

to create channels for fluidic thermal connects. One of the key concepts is that while 3D stacked systems produce a heat generation problem, they also provide four more surfaces to use for heat extraction (or two pairs of surfaces for flow-through heat extraction).



Fig. 41: Thermal management in a 3D IC. Here the 3D IC is a 4-wafer bonded stack.

Intelligent Transportation Systems

Personnel

J. F. Coughlin, B. K. P. Horn, J. K. Kucher, T. B. Sheridan, C. G. Sodini, and J. M. Sussman

Sponsorship

Intelligent Transportation Research Center at MIT's MTL

Transportation is important not only economically but also socially. The inter-state highway project built a sound infrastructure for our society. US citizens are spending, on average, about \$1,000 per year on cars, trucks, and roads. What infrastructure do we need for tomorrow? The goal of this project is to develop a technical foundation for tomorrow's transportation systems. Currently we have a number of infrastructures which are independent from each other; examples include infrastructures for transportation, communication, finance, health care, emergency care, and others. In the next generation, these independent infrastructures will be integrated more closely with advanced information technologies. For example, highway tolls can be charged to drivers' bank accounts automatically with electronic toll gates connected to banks' computers. If a car accident occurs, as another example, the accident can be detected by an air-bag sensor and reported automatically through wireless network to ambulance stations. The ambulance and hospital will have a teleconference on the way from the scene to the hospital for rapid medical intervention.

For transportation, safety is very important. We are working on technologies which compensate for the diverse characteristic changes caused by aging, in collaboration with MIT's age laboratory. A typical 50-yearold driver, for example, needs twice as much light to see as a typical 30-year-old driver, and we are developing a pedestrian detection system based on infrared images to make night driving safer.

The Intelligent Transportation Systems project consists of various research topics ranging from small-scale systems to large-scale systems as well as fundamental to application oriented subprojects. An example of small-scale subprojects is an adaptive dynamic range image acquisition chip. Medium-scale systems include a personal-computer-based real-time three-dimensional machine vision system, sensor fusion systems, and an image recognition system for compressed three-dimensional images without decompression. Examples of large-scale systems are an image sensor network and the safety analysis of a fully-automated transportation system.

The research is being carried out at the Intelligent Transportation Research Center in MIT's Microsystems Technology Laboratories. The center is being sponsored by several member companies.

The Low-Power Bionic Ear Project (ICS)

Personnel M. Baker, C. Salthouse, J.J. Sit, and S. Zhak (R. Sarpeshkar)

Sponsorship Advanced Bionics Corporation

The aim of the project is to construct a cochlear-implant processor for the deaf that has the potential to reduce the current power consumption of such processors by more than an order of magnitude via low power analog VLSI processing. In addition, a cochlear implant processor that is based on the architecture of a silicon cochlea, i.e., on an analog electronic model of the inner ear, is being explored for its potential to revolutionize patient's speech recognition in noise (Rahul Sarpeshkar, Lorenzo Turicchia, George Efthivoulidis, and Luc Van Immerseel, "The Silicon Cochlea: From Biology to Bionics", accepted paper, Proceedings of The Biophysics of the Cochlea: Molecules to Models Conference, Titisee, Black Forest, Germany, July 27-August 1, 2002.)

Figure 42 shows the overall system architecture of a current bionic ear system (cochlear implant system). Sound that is transduced from a microphone is eventually converted into electrode stimulation in surgically implanted electrodes. The aim of this project is to reduce the power consumption to levels that will enable fully implanted systems to become a reality.

Several building block circuits for such a processor including a 100uW analog front end, a programmable bandpass filter, and a logarithmic map circuit were designed. Figure 43 shows a chip photograph of a DAC programmable fourth-order programmable bandpass filter that operates on 6uW of power consumption with over 60dB of dynamic range on a 2.8V supply (Christopher Salthouse and Rahul Sarpeshkar, "A Micropower Bandpass Filter for Use in Cochlear Implants", accepted paper, IEEE International Symposium on Circuits and Systems, Arizona, May 2002.)



Fig. 42



Fig. 43

The Visual Motion and Inertial Motion Sensing Project

Personnel M.Tavakoli-Dastjerdi (R. Sarpeshkar)

Sponsorship

Caltech Subcontract of DARPA Funding

This project maps the distributed feedback loops of biological photoreceptors to silicon to create low-power high-performance silicon photoreceptors. Such photoreceptors are useful as front ends in VLSI motion sensors; important in robotic and active-vision applications. An ultra-low-noise MEMS vibration sensor, which provides inertial information to a vibrating visual sensor being built by collaborators at Caltech, has also been built.

Figure 44 shows the VLSI layout of a visual motion sensor that yields the speed and direction of a globally moving visual image along the "Y" direction. The array contains both photodiodes and analog VLSI processing circuitry that is inspired by similar circuitry in the housefly. Figure 45 shows the experimental setup for testing a capacitive MEMS vibration sensor with associated ultra low noise offset-compensating electronics on an electronics die which is wirebonded to the MEMS die. The sensor achieved an electronic noise floor equivalent to 30ug/rt(Hz) over a 1Hz-100Hz bandwidth, a specification that appears to be 3 times better than any equivalent commercial or research system in spite of its separate-die solution for mechanical and electrical systems. The system was able to detect a change of 1 part per 5 million in capacitance. The offset-compensating electronics has been briefly described in "A Low-Noise Nonlinear Feedback Technique for Compensating Offset in Analog Multipliers", Maziar Tavakoli-Dastjerdi and Rahul Sarpeshkar, accepted paper, IEEE International Symposium on Circuits and Systems, Arizona, May 2002.



Fig. 44





Spike-Based Hybrid Computers Project

Personnel M. O'Halloran, A. Mevay, H. Yang, R. Sarpeshkar

Sponsorship

Office of Naval Research

This project attempts to combine the best of analog and digital computation to compute more efficiently than would be possible in either paradigm of computation. (Rahul Sarpeshkar and Micah O'Halloran, "Scalable Hybrid Computation with Spikes", in press, Neural Computation, 2002). This project is inspired by the duality of analog spike-time and digital spike-count codes of the brain's neurons. It is being applied to create lowpower time-based analog-to-digital converters, analog memories, and novel event-based control architectures. Several design issues that are important in mixed-signal systems including good power supply rejection are being explored.

Figure 46 shows the layout of a low power analogto-digital converter that uses time as a signal variable rather than the traditional variables of voltage or current to perform quantization. A technique for achieving good power supply rejection without sacrificing the gain bandwidth product of an amplifier has been reported (Micah O'Halloran and Rahul Sarpeshkar, "A Low Open-Loop Gain High-PSRR Micropower CMOS Amplifier for Mixed-Signal Applications", paper, IEEE International Symposium on Circuits and Systems, Arizona, May 2002).



Fig. 46