Radiocommunications

1. Draw and explain the main blocks of a radio receiver.

Course nb.1 slide 12.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

The receiver block incorporate many amplifier and processing stages, and one of the most important is the oscillator stage.

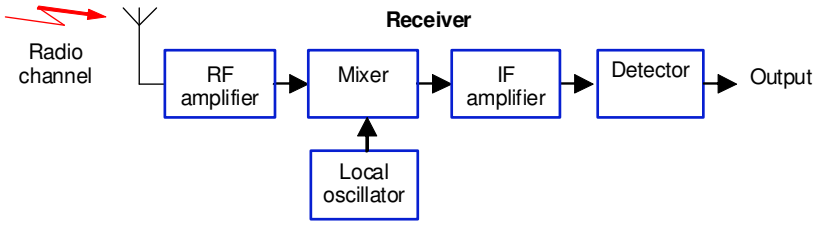

The receiver oscillator is called the local oscillator as it produces a local carrier within the receiver which allows the incoming carrier from the transmitter to be down converted for easier processing within the receiver. The mixer and the IF amplifier are used to extract the intermediate frequency. The user information is obtained after detection.

2. What wavelength corresponds to a frequency of 600 MHz?

Course nb.2 slide 45.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

$\lambda = c/f = (3x10^8 m/s)/(6x10^8 Hz) = 0.5$ m

3. What are the radiation regions of an antenna?

Course nb.3 slide 36-39.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

The antenna radiation field is divided into:

- reactive near-field (objects within this region will result in coupling with the antenna and distortion of the ultimate far-field antenna pattern),

- radiating near-field (transition region),

- far-field (the gain of the antenna is a function only of angle).

4. Describe the directivity of a half-wave dipole antenna.

Course nb.4 slide 35-36.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

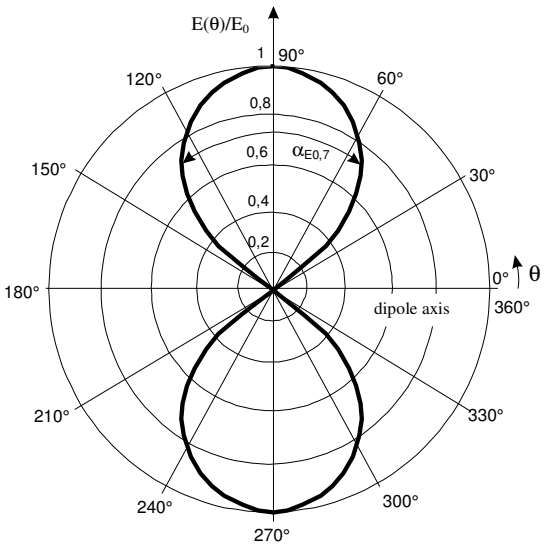

Answer:

A half-wave dipole has an antenna gain of 1.64 or $G = 2.15$ dBi.

It has an Omni directional pattern in the H-plane.

In E-plane the directivity is bidirectional.

5. What represents the array factor?

Course nb.5 slide 15-16.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

An array of antennas working simultaneously can focus the reception or transmission of energy in a particular direction, which increases the useful range of a system.

The array influence is contained inside of an array factor AF:

Array pattern = Element pattern x Array factor

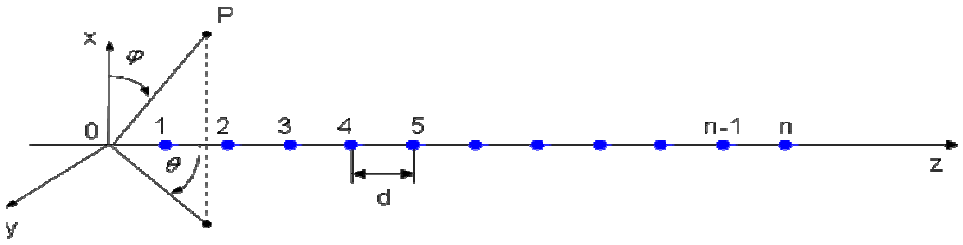

6. Draw and explain the main blocks of a superheterodyne receiver.

Course nb.6 slide 9-11.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

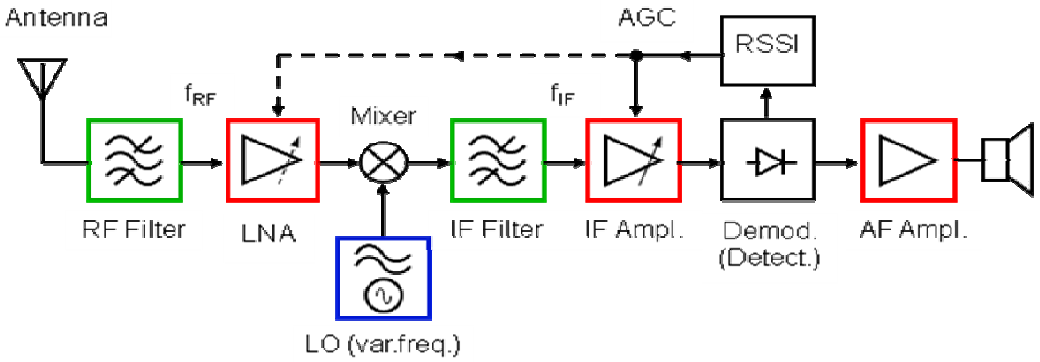

The RF front end consists in a band-pass filter and low-noise amplifier for radio bandwidth selection. The local oscillator and the mixer followed by an IF filter are used for heterodyne process performing channel selection. IF blocks amplify the signal to ensure the right level at the demodulator input. An automatic gain control loop is used to maintain a constant level in case of fading.

7. Draw and explain the QPSK modulator.

Course nb.7 slide 33-34.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

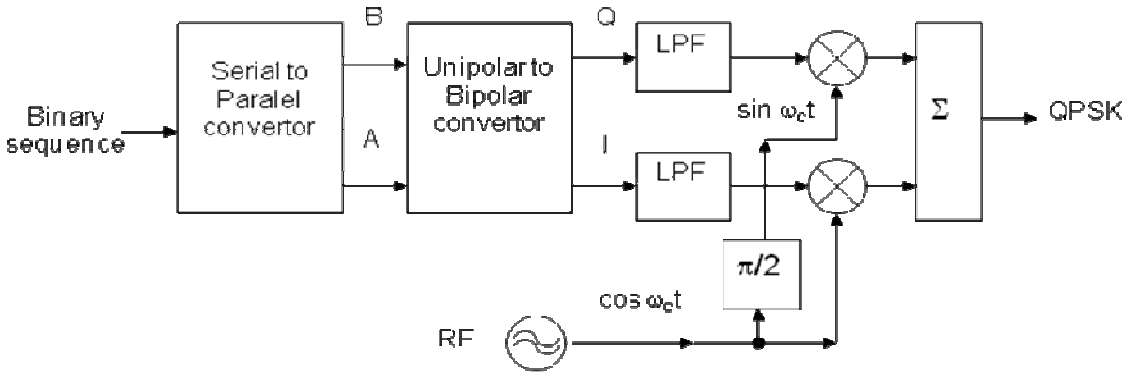

The baseband coding of information consists first in parallel conversion. The two streams resulted are converted in bipolar signals with a symbol duration being twice the bit duration. Each of them is low-pass filtered and used to modulate in phase a RF carrier. The two carriers have a 90^{o} phase shift to be orthogonal. The result is a constant amplitude signal with 4 possible states of its phasor (4 state constellation).

8. Explain the image reject mixer with Hartley architecture.

Course nb.8 slide 17-23.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

The principle of image-rejection is to process the desired channel and its image in such a way that the image can be eliminated eventually by *signal cancellation*.

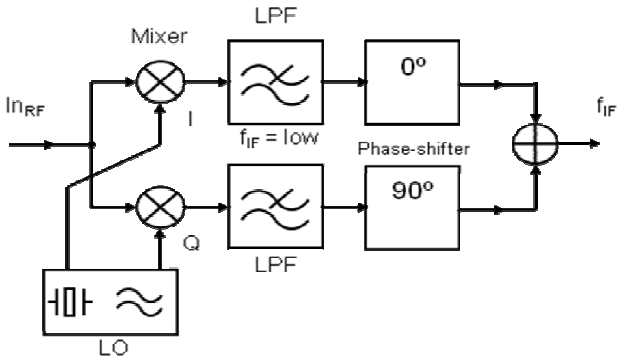

The RF signal in the downconversion is split into two components by using two matched mixers and quadrature LO signals.

The resulting IF signals, namely in phase (I) and quadrature phase (Q), are then lowpass-filtered and after one is phase-shifted by 90°, the IF signals are combined.

In this process, depending on the IF path that is subjected to the 90° phase-shifter, either the image band or the receive band is cancelled after the summation of I and Q outputs.

9. Define the receiver's selectivity.

Course nb.9 slide 19.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

A receiver's selectivity performance is a measure of the ability to separate the desired band about the carrier, from unwanted interfering signals received at other frequencies.

This situation is most often characterized by a weak received desired signal in the presence of a strong adjacent or alternate band user.

Receiver selectivity may be defined also as the ability to reject unwanted signals on adjacent channel frequencies.

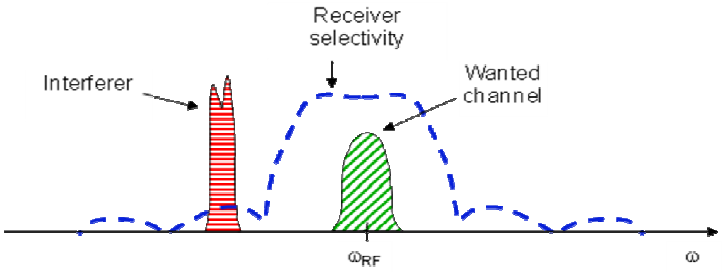

10. Which are the major components of the PLL (Phase-locked loop) frequency synthesizer?

Course nb.10 slide 27-29.

Site address: https://intranet.etc.upt.ro/~RADIOCOM/

Answer:

A PLL is a circuit which causes a particular system to track with another one. More precisely, a PLL is a circuit synchronizing an output signal (generated by an oscillator) with a reference or input signal in frequency as well as in phase.

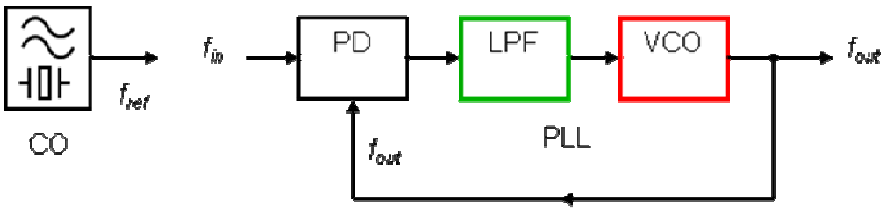

The major components are the Phase Detector (PD), the LPF = Loop Filter (Low-pass filter LPF), and the Voltage-controlled oscillator (VCO).

Power Electronics

1. Define the total harmonic distortion coefficient (*THD*) and the power factor (*PF*). Power factor formula for sinusoidal input voltage and nonlinear load.
   Chapter 2, slides 4, 13, 15.

Total harmonic distortion (*THD*) coefficient

Given a periodic signal $x(t)$, the total harmonic distorion coefficient is defined as:

$$(\text{THD}) = \frac{(\text{rms without fundamental})}{(\text{rms fundamental value})} = \frac{\sqrt{X_0^2 + \sum_{n=2}^{\infty} \frac{X_n^2}{2}}}{\frac{X_1}{\sqrt{2}}}$$

where X_0 is the dc component and X_i are the amplitudes of harmonics.

Power factor

For efficient transmission of energy from a source to a load, it is desired *maximize average power*, while *minimizing rms current and voltage* to (and hence minimizing losses).

Power factor (*PF*) is a *figure of merit* that measures how energy is efficiently transmitted. For a port defined by voltage *v(t)* and current *i(t)*, both of period *T*, the power factor is defined as:

$$PF = \frac{P}{V_{\text{rms}} \cdot I_{\text{rms}}} = \frac{\int_{t_0}^{t_0+T} v(t) \cdot i(t) dt}{\sqrt{\frac{1}{T}\int_{t_0}^{t_0+T} v(t)^2 dt} \cdot \sqrt{\frac{1}{T}\int_{t_0}^{t_0+T} i(t)^2 dt}}$$

Power factor always lies between zero and one if the real energy flow sense is chosen

$$|PF| \leq 1$$

Nonlinear dynamical load, sinusoidal voltage

With a sinusoidal voltage, current harmonics do not lead to average power. However, current harmonics do increase the rms current, and hence they decrease the power factor.

$$P = \frac{V_1 I_1}{2} \cos(\varphi_1 - \theta_1)$$

$$I_{\text{rms}} = \sqrt{I_0^2 + \sum_{n=1}^{\infty} \frac{I_n^2}{2}}$$

$$\text{PF} = \left(\frac{\frac{I_1}{\sqrt{2}}}{\sqrt{I_0^2 + \sum_{n=1}^{\infty} \frac{I_n^2}{2}}} \right) \left(\cos(\varphi_1 - \theta_1) \right) = K_d \cdot K_\theta$$

$$\text{PF} = \text{(distortion factor)} \cdot \text{(displacement factor)}$$

2. The three *basic dc-dc nonisolated converters*: buck, boost, buck-boost: schematics with SPDTs and their static conversion ratios values. Boost and Ćuk converters practical schematics and Ćuk static conversion ratio.
   Chapter 3continued, slides 5, 18, 24, 28, 35.

Buck

(a)

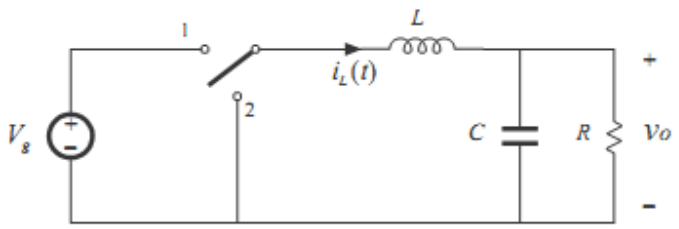

$$M(D) = D$$

Boost

(b)

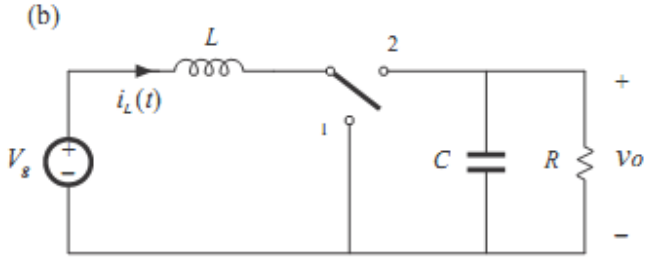

$$M(D) = \frac{1}{1-D}$$

Buck-boost

(c)

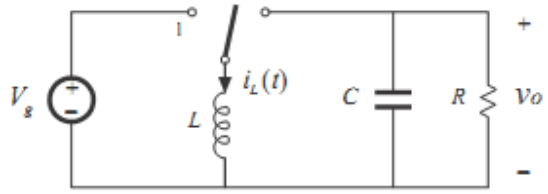

$$M(D) = \frac{-D}{1-D}$$

Boost converter:

Realization using power MOSFET and diode

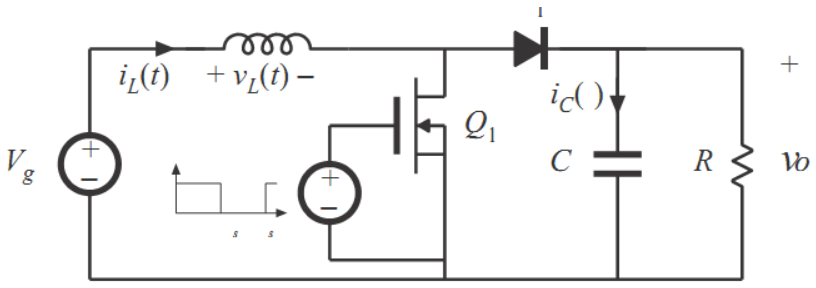

Cuk converter:

Cuk converter: practical realization using MOSFET and diode

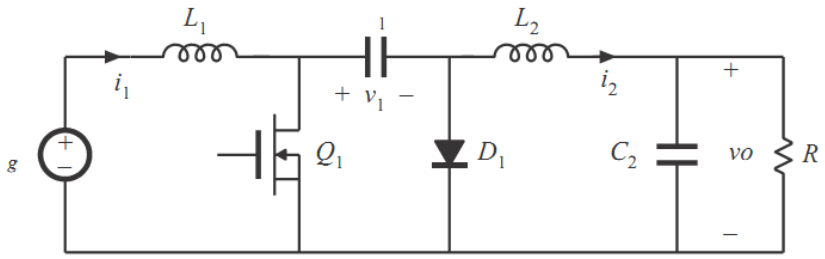

$$M(D) = \frac{V_o}{V_g} = -\frac{D}{1-D}$$

3. The *sychronous rectifier* – schematics, advantages.
   Chapter 5, slide 26.

Buck converter with synchronous rectifier

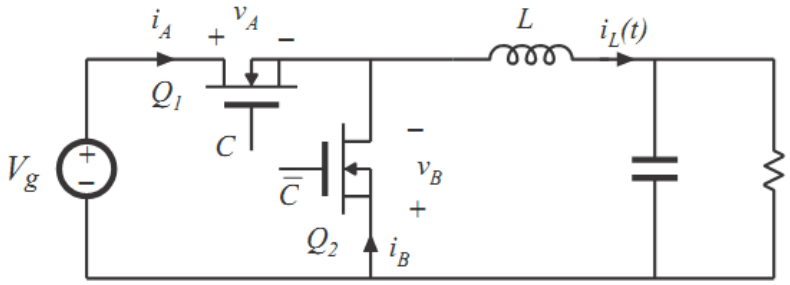

- MOSFET Q_2 is controlled to turn on when diode would normally conduct
- Semiconductor conduction loss can be made arbitrarily small, by reduction of MOSFET on-resistances
- Useful in low-voltage high-current applications

4. Nonideal converters analysis. Boost equivalent model including *semiconductor conduction losses*. Output voltage calculation.
   Chapter 3, slides 25, 29, 30.

Inclusion of semiconductor conduction losses in the boost converter model

Boost converter example

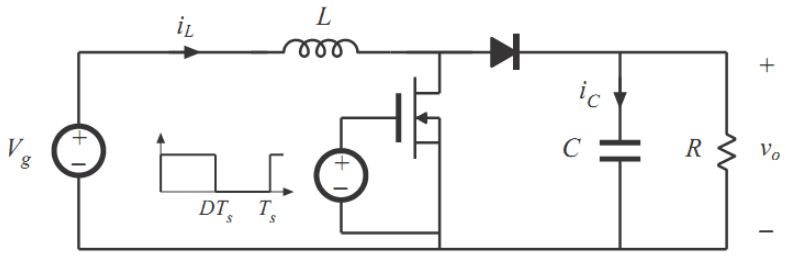

Models of on-state semiconductor devices:

MOSFET: on-resistance R_{on}

Diode: constant forward voltage V_D plus on-resistance R_D

Insert these models into subinterval circuits

Complete equivalent circuit

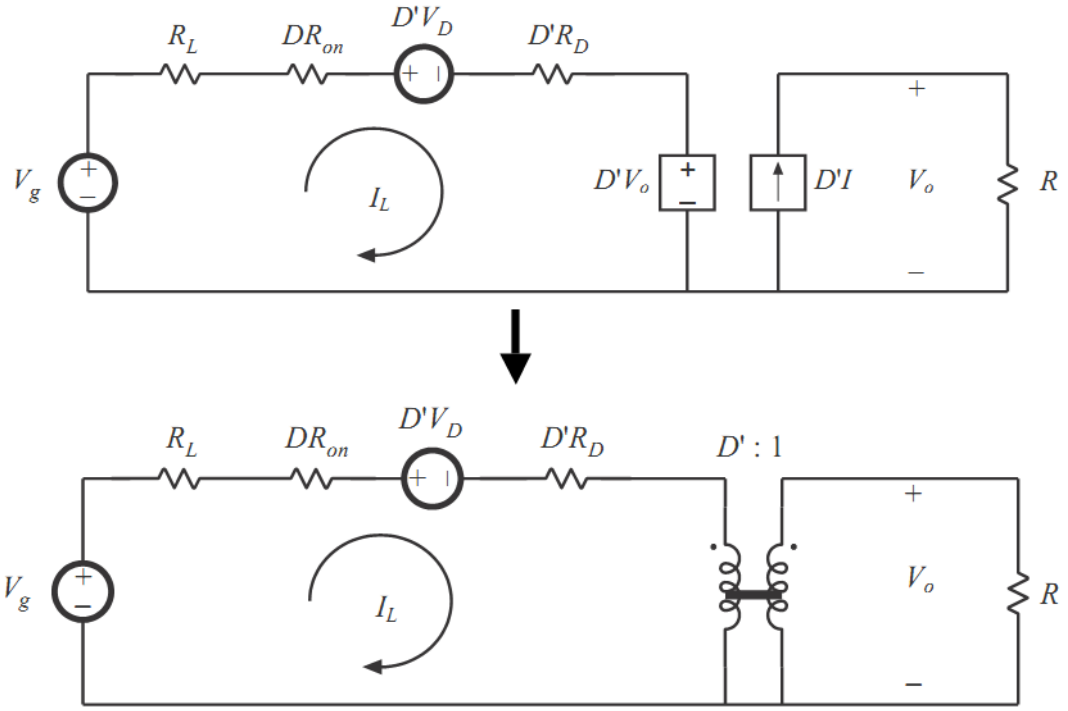

Solution for output voltage

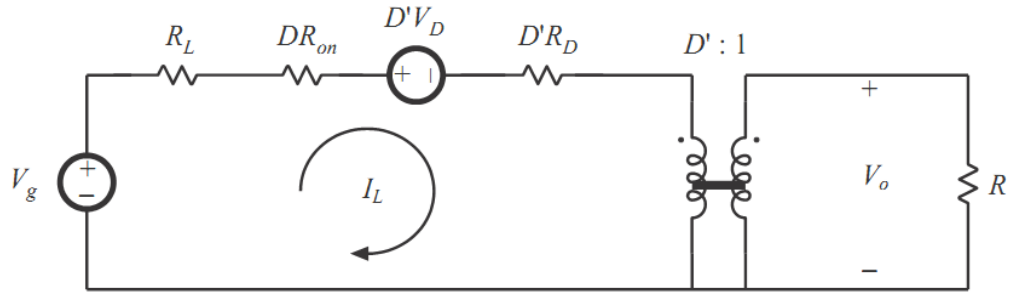

$$V_o = \left(\frac{1}{D'}\right)\left(V_g - D'V_D\right)\left(\frac{D'^2R}{D'^2R + R_L + DR_{on} + D'R_D}\right)$$

$$\frac{V_o}{V_g} = \left(\frac{1}{D'}\right)\left(1 - \frac{D'V_D}{V_g}\right)\left(\frac{1}{1 + \frac{R_L + DR_{on} + D'R_D}{D'^2R}}\right)$$

5. The *flyback* converter: schematics, static conversion ratio, applications and limitations. Chapter 6, slides 59, 53, 65.

Flyback converter schematics:

Flyback converter having a 1:n turns ratio and positive output:

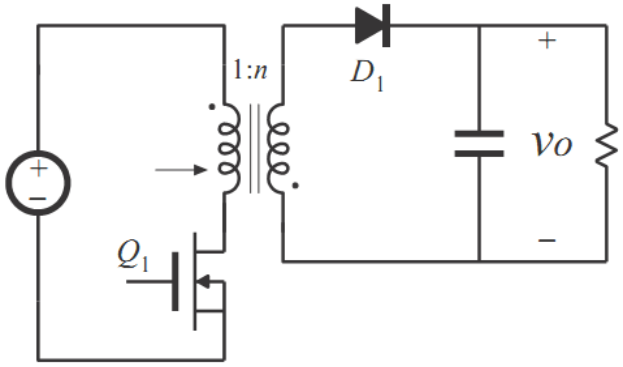

Conversion ratio $M(D)$

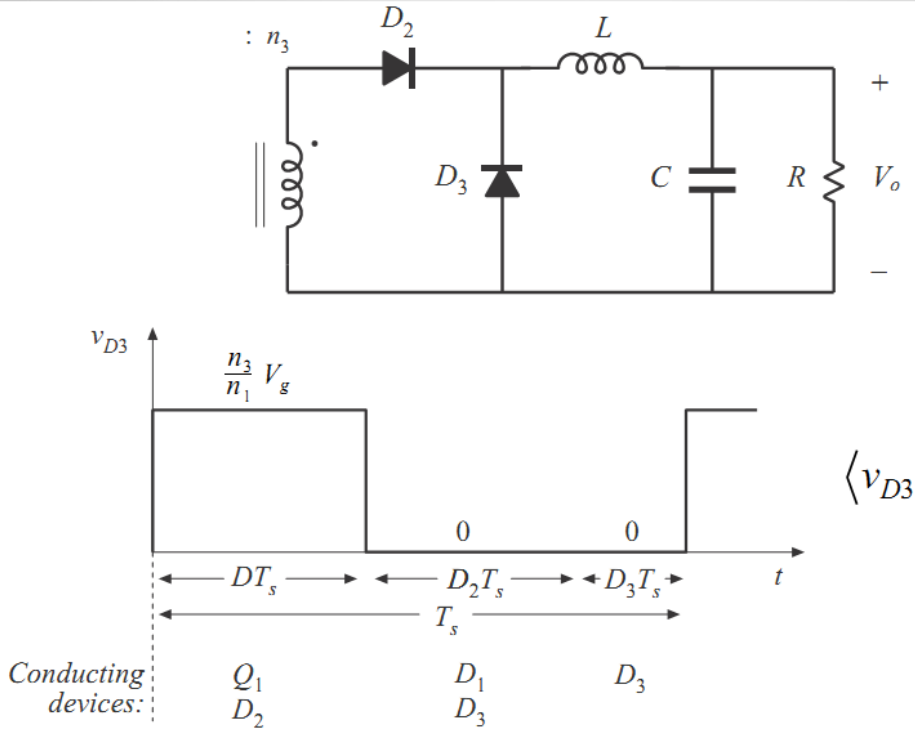

$$\langle v_{D3} \rangle = V_o = \frac{n_3}{n_1} D V_g$$

Discussion: Flyback converter

- Widely used in low power and/or high voltage applications
- Low parts count
- Multiple outputs are easily obtained, with minimum additional parts
- Cross regulation is inferior to buck-derived isolated converters
- Often operated in discontinuous conduction mode
- DCM analysis: DCM buck-boost with turns ratio

6. The classical single-transistor *forward* converter: schematics, main waveforms, maximum duty cycle.
Chapter 6, slides 44, 46, 51, 53.

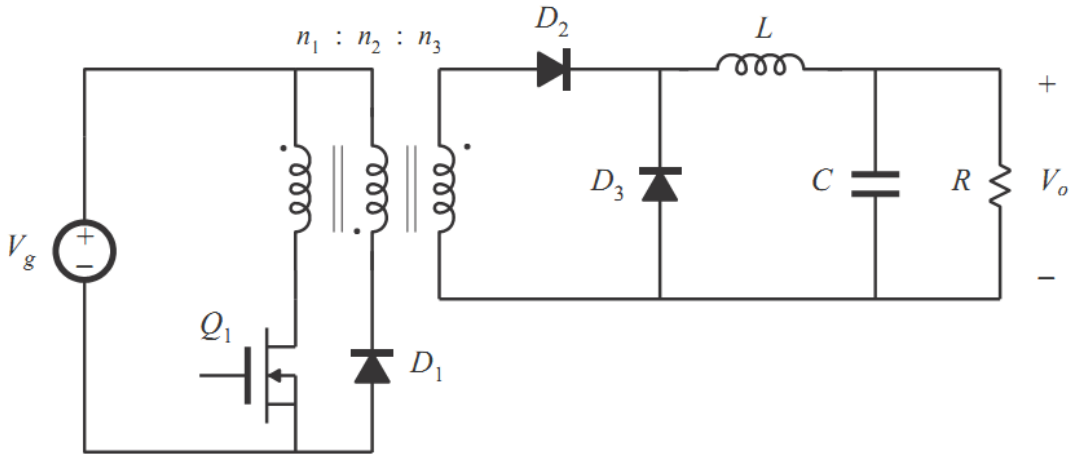

- Buck-derived transformer-isolated converter
- Single-transistor and two-transistor versions
- Maximum duty cycle is limited
- Transformer is reset while transistor is off

Forward converter: waveforms

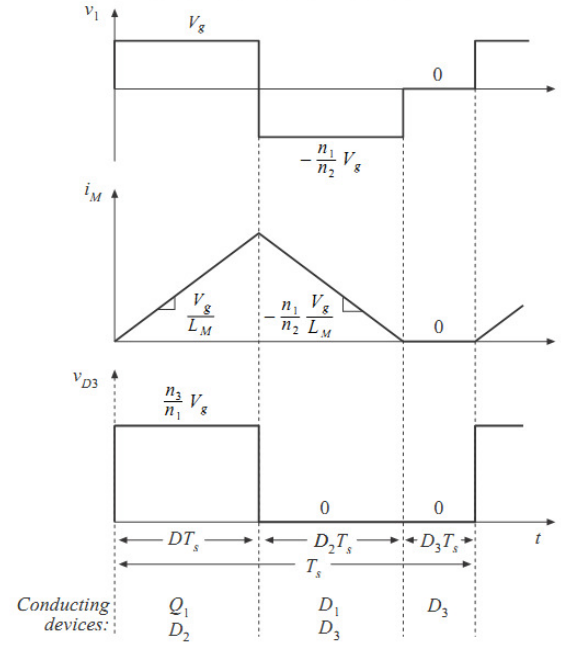

- Magnetizing current, in conjunction with diode D_1, operates in *discontinuous* conduction mode
- Output filter inductor, in conjunction with diode D_3, may operate in either CCM or DCM

Transformer reset

From magnetizing current volt-second balance:

$$\langle v_1 \rangle = D\left(V_g\right) + D_2\left(-V_g n_1/n_2\right) + D_3\left(0\right) = 0$$

Solve for D_2:

$$D_2 = \frac{n_2}{n_1} D$$

D_3 cannot be negative. But $D_3 = 1 - D - D_2$. Hence

$$D_3 = 1 - D - D_2 \geq 0$$

$$D_3 = 1 - D\left(1 + \frac{n_2}{n_1}\right) \geq 0$$

Solve for D

$$D \leq \frac{1}{1 + \frac{n_2}{n_1}} \qquad \text{for } n_1 = n_2: \quad D \leq \frac{1}{2}$$

7. The *full-bridge* isolated buck converter: schematics, main waveforms, solutions for preventing core saturation.
   Chapter 6, slides 37, 39, 40, 42

Full-bridge isolated buck converter

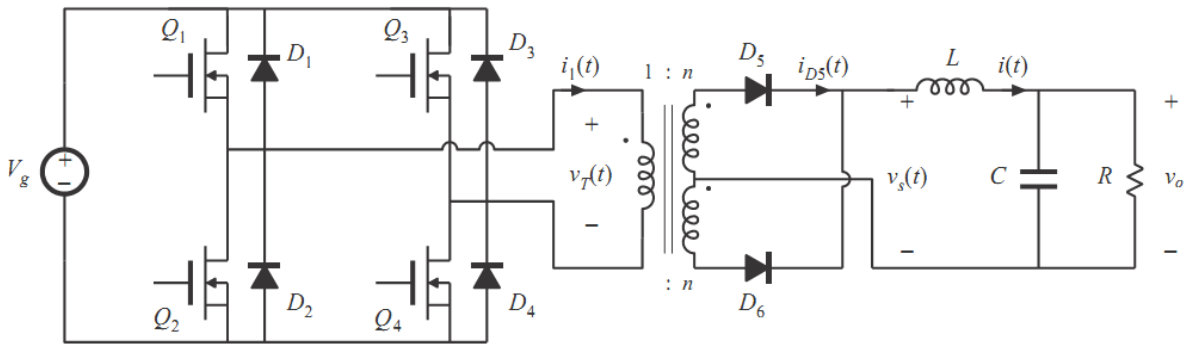

Full-bridge: waveforms

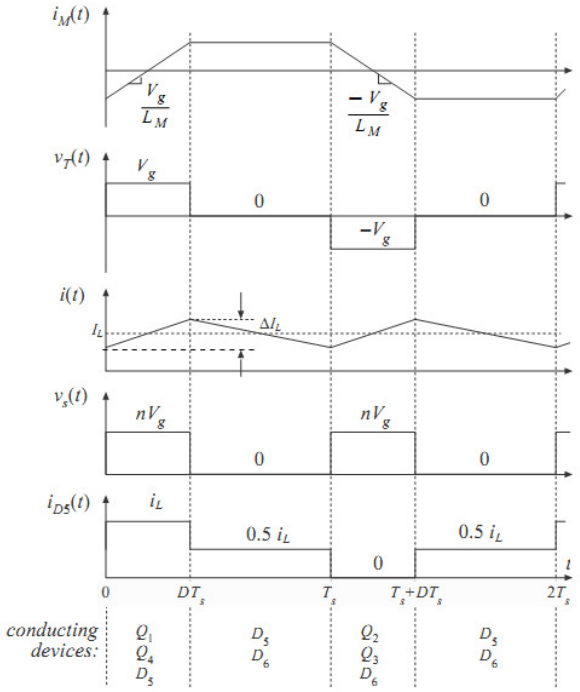

- During first switching period: transistors Q_1 and Q_4 conduct for time DT_s , applying volt-seconds $V_g\,DT_s$ to primary winding
- During next switching period: transistors Q_2 and Q_3 conduct for time DT_s , applying volt-seconds $-V_g\,DT_s$ to primary winding
- Transformer volt-second balance is obtained over two switching periods
- Effect of nonidealities?

Effect of nonidealities on transformer volt-second balance

Volt-seconds applied to primary winding during first switching period:

$$(V_g - (Q_1 \text{ and } Q_4 \text{ forward voltage drops})) \text{ x } (Q_1 \text{ and } Q_4 \text{ conduction time})$$

Volt-seconds applied to primary winding during next switching period:

$$-(V_g - (Q_2 \text{ and } Q_3 \text{ forward voltage drops})) \text{ x } (Q_2 \text{ and } Q_3 \text{ conduction time})$$

These volt-seconds never add to *exactly* zero.

Net (nonzero) volt-seconds are applied to primary winding

Magnetizing current slowly increases in magnitude

Saturation can be prevented by placing a *capacitor* in series with primary, or by use of the so called *current programmed mode*

8. *Push-pull isolated buck* and *isolated Ćuk* converters: schematics, type of control, advantages and disadvantages regarding transformer magnetizing current.
Chapter 6, slides 56, 57, 76.

Push-pull isolated buck converter

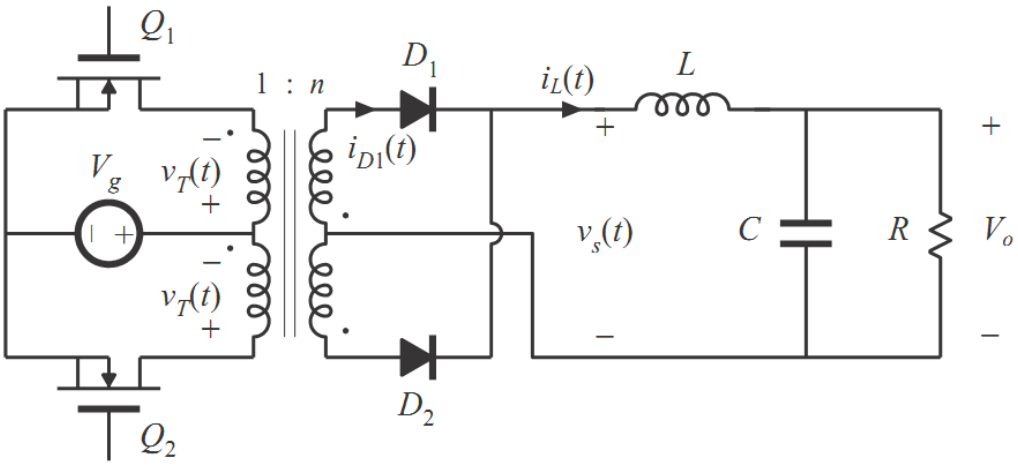

$$V_o = nDV_g \qquad 0 \le D \le 1$$

Waveforms: push-pull

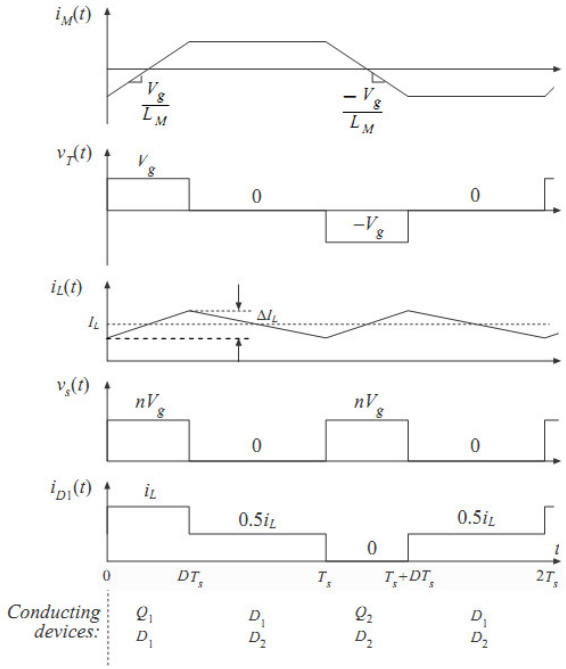

- Used with low-voltage inputs
- Secondary-side circuit identical to full bridge
- As in full bridge, transformer volt-second balance is obtained over two switching periods
- Effect of nonidealities on transformer volt-second balance?
- Current programmed control can be used to mitigate transformer saturation problems. Duty cycle control not recommended.

Isolated Cuk converter

Insert transformer between capacitors C_{1a} and C_{1b}

$$M(D) = \frac{V_o}{V_g} = \frac{nD}{D'}$$

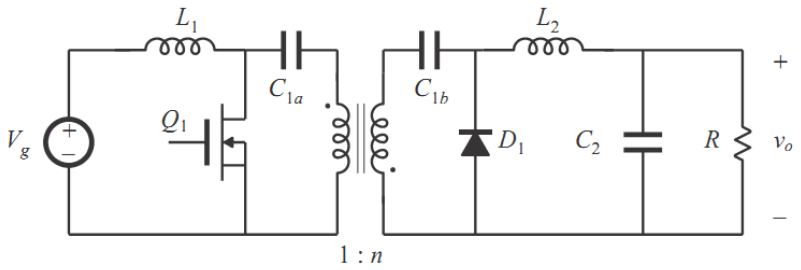

Discussion

- Capacitors C_{1a} *and* C_{1b} ensure that no dc voltage is applied to transformer primary or secondary windings
- Transformer functions in conventional manner, with small magnetizing current and negligible energy storage within the magnetizing inductance

9. Principle of *current programmed control*: block diagram, advantages, stability.
   Chapter 8, slides 1, 2, 3.

Current Programmed Control (CPM)

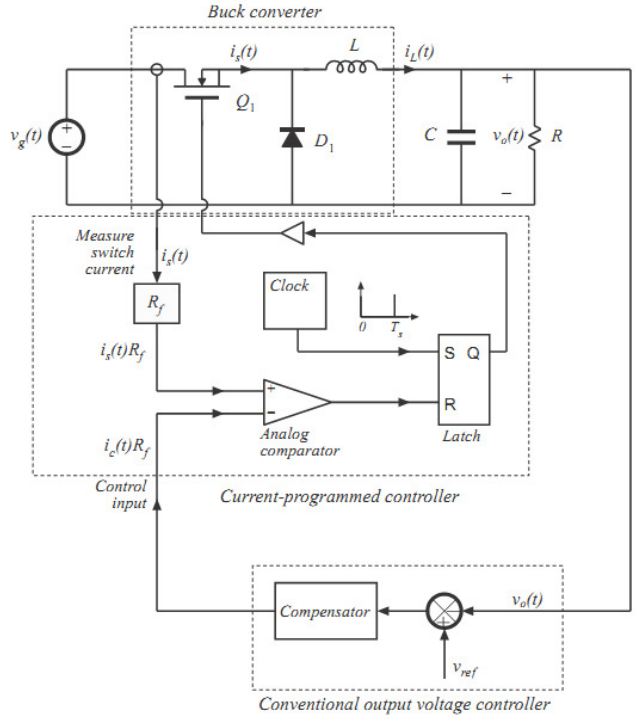

The peak transistor current replaces the duty cycle as the converter control input.

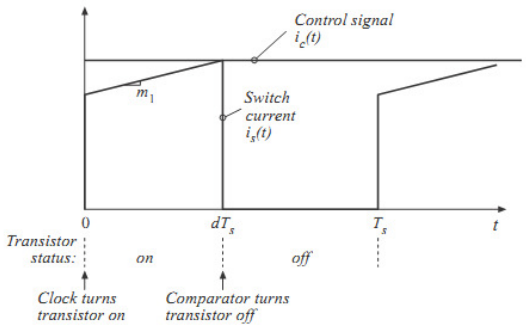

Current programmed control vs. duty cycle control

Advantages of current programmed control:

- Simpler dynamics —inductor pole is moved to high frequency
- Simple robust output voltage control, with large phase margin, can be obtained without use of compensator lead networks
- It is always necessary to sense the *transistor* current, to protect against overcurrent failures. We may as well use the information during normal operation, to obtain better control
- Transistor failures due to excessive current can be prevented simply by limiting $i_c(t)$
- Transformer saturation problems in bridge or push-pull converters can be mitigated

A disadvantage: susceptibility to noise

Oscillation for $D > 0.5$

- The current programmed controller is inherently unstable for $D > 0.5$, regardless of the converter topology
- Controller can be stabilized by addition of an artificial ramp

10. The *three-phase bridge rectifier (six pulse rectifier)*: schematics, line current for high load inductor, harmonic content of the line current and output voltage. Chapter 7, slides 10, 11.

The Three-Phase Bridge Rectifier (six pulse rectifier)

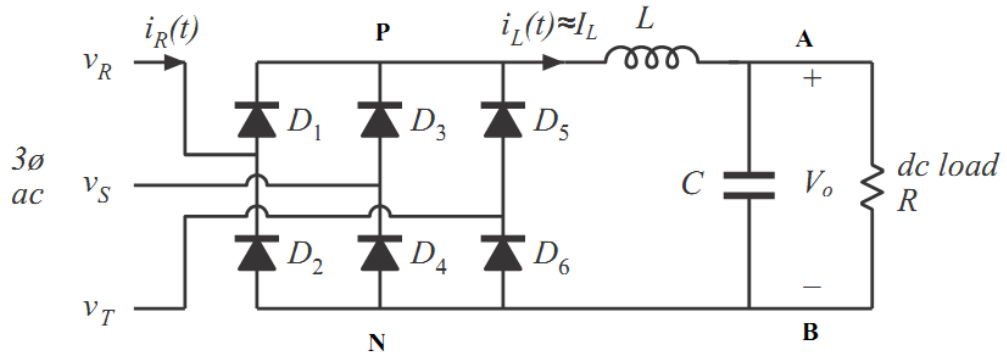

Line current waveform for infinite L

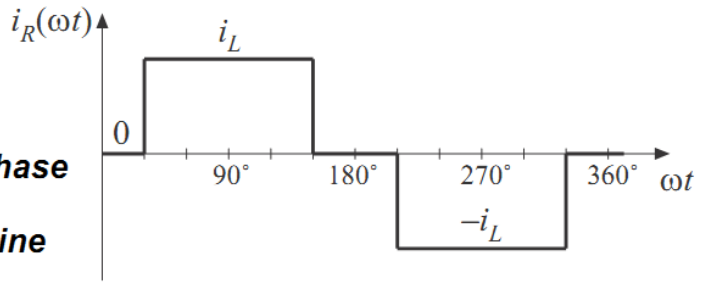

Operation: maximum *phase* voltage criteria !!!
Output voltage equals *line* voltages.

The Three-Phase Bridge Rectifier

- Dc and a fundamental of *six* times the line frequency in the output voltage: 0, 6, 12, 18, etc.
- Odd non-triplen harmonics in the ac line current.
- No more than two of the six diodes can conduct simultaneously. Therefore *line current* exhibits intervals during which the current is zero.
- Unlike the single-phase case, the ac line current *contains distortion* even with the filter elements removed (pure *resistive case*).

Embedded Systems

1. The general architecture of an embedded system.

CH. 1, slide 21.

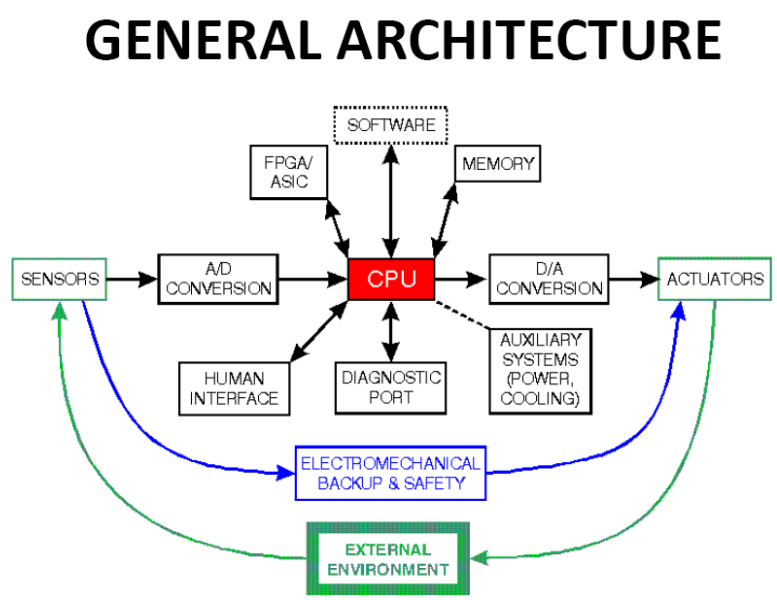

Fig. 2. The General Architecture of an Embedded System

2. What are the relative advantages/disadvantages of RISC versus CISC architectures?

CH. 3, slides 13-16.

INSTRUCTION SET ARCHITECTURE

- The most common types of general-purpose ISA architectures implemented in embedded processors are:

- **Complex Instruction Set Computing (CISC) Model**
- **Reduced Instruction Set Computing (RISC) Model**

COMPLEX INSTRUCTION SET COMPUTING (CISC)

- A **large number of instructions** each carrying out different permutation of the same operation
- Instructions provide for **complex operations**
- Different instructions of **different format**
- Different instructions of **different length**
- Different **addressing modes**
- Requires **multiple cycles** for execution

REDUCED INSTRUCTION SET COMPUTING (RISC)

- **Fewer instructions** aiming simple operations that can be executed in a **single cycle**
- Each instruction of **fixed length** – facilitates instruction pipelining
- **Large** general purpose **register set** – can contain data or address
- **Load-store Architecture** – no memory access for data processing instructions

RISC-CISC

- Over time, improvements in chip fabrication techniques have improved performance exponentially, according to Moore's law, whereas architectural improvements have been comparatively small.
- Modern CISC implementations have adopted many of the performance improvements introduced by RISC, such as single-clock instructions.
- The RISC-CISC distinction has blurred significantly in practice.

3. The main features of an ARM7TDMI are:

T: ……..

D: ……..

M: ……

I: ……

Explain them.

CH. 4, slide 9.

ARM Architecture Versions

Family	Architecture Version	Core	Feature	In application
ARM1	ARMv1	ARM1	26 bit addressing, no multiply or coprocessor	ARM Evaluation System
ARM2, ARM3	ARMv2/2a	ARM2/250/2a	Includes 32-bit result multiply co-processor	Acorn Archimedes
ARM6, ARM7	ARMv3	ARM600…7500 FE	32 bit addressing	Acorn Archimedes Apple eMate 300 Psion Series 5

4. Enounce and explain the role of the following ARM registers: r13, r14 and r15, status registers.

CH. 4, slide 16, 20

Registers

- General Purpose registers hold either data or address
- All registers are 32 bits
- In user mode 16 data registers and 2 status registers are visible
- Data registers: r0 to r15
- Three registers r13, r14, r15 perform special functions: r13 – **stack pointer**, r14 – **link register** (where return address is put whenever a subroutine is called, r15 – **program counter**)

Status Registers

- CPSR: monitors and control internal operations

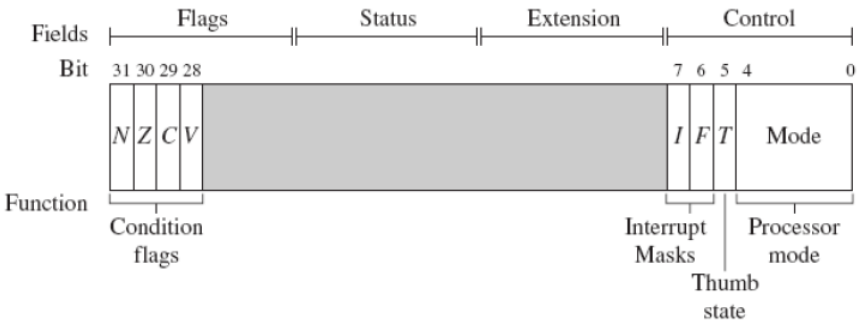

Fig. 4.3. A generic program status register (*psr).*

5. Which is the role of the barrel shifter? Enumerate the basic operations which could be performed with it. Illustrate the concept with an assembly language example.

CH. 4, slide 42-44.

Using Barrel Shifter

- Enables shifting 32-bit operand in one of the source registers left or right by a specific number of positions within the cycle time of instruction
- Basic Barrel shifter operations
  - Shift left, right, rotate
- Facilitates fast multiply, division and increases code density
- Example: mov r7, r5, LSL #2
  - Multiplies content of r5 by 4 and puts result in r7

Using Barrel Shifter

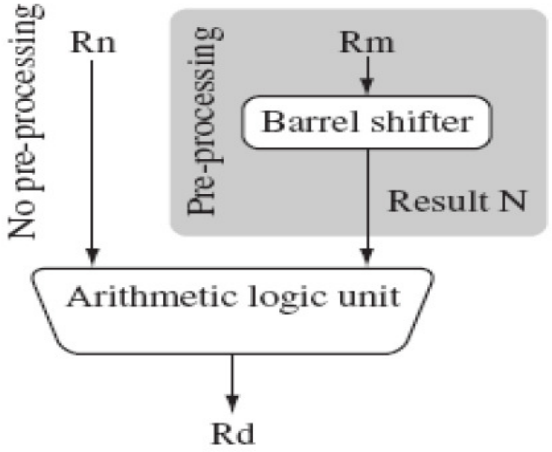

Fig. 4.6. Barrel shifter and ALU.

Using Barrel Shifter

Mnemonic	Description	Shift	Result	Shift amount y
LSL	logical shift left	*x*LSL *y*	$x \ll y$	#0–31 or *Rs*
LSR	logical shift right	*x*LSR *y*	$(\text{unsigned})x \gg y$	#1–32 or *Rs*
ASR	arithmetic right shift	*x*ASR *y*	$(\text{signed})x \gg y$	#1–32 or *Rs*
ROR	rotate right	*x*ROR *y*	$((\text{unsigned})x \gg y) \mid (x \ll (32 - y))$	#1–31 or *Rs*
RRX	rotate right extended	*x*RRX	$(c \text{ flag} \ll 31) \mid ((\text{unsigned})x \gg 1)$	none

Note: *x* represents the register being shifted and *y* represents the shift amount.

Fig. 4.6b. Barrel shifter operations.

6. Present possible implementations for the non-volatile memory. What could be store in it?

CH. 5, slide 10-12.

Non-volatile Memory

- Mask ROM
  - Used for dedicated functionality
  - Contents fixed at IC fab time (truly write once!)
- ERPOM (erase programmable)
  - Requires special IC process (floating gate technology)
  - Writing is slower than RAM, EPROM uses special programming system to provide special voltages and timing
  - Reading can be made fairly fast
  - Rewriting is slow
    - Erasure is first required, EPROM – UV light exposure, EEPROM – electrically erasable
- Flash

Flash Memory

- Uses single transistor per bit (EEPROM employs two transistors)
- A flash memory provides high density storage with speed marginally less than that of SRAM's
- Write time is significantly higher compared to DRAM

Embedded Non-volatile Storage

- On-chip non-volatile storage is used for storage of:
  - Configuration information
  - Executable code that runs on core processors
  - Recorded data: repeated write

7. The I2C protocol (features, connections, protocol, advantages, disadvantages).

CH. 6, slide 19-24.

I2C

- Shorthand for an "Inter-integrated circuit" bus
- Developed by Philips Semiconductor for TV sets in the 1980's
- I2C devices include EEPROMs, thermal sensors, and real-time clocks
- Used as a control interface to signal processing devices that have separate data interfaces, e.g. RF tuners, video decoders and encoders, and audio processors

I2C Features

- **Bi-directional**
  - Data can flow in both directions
- **Synchronous**
  - **Data** is clocked along with a clock signal
  - **Clock** signal controls when data is changed and when it should be read
  - **Clock rate can vary** unlike asynchronous (RS-232 style) communication
- I2C bus has three speeds:
  - Slow (under 100 Kbps)
  - Fast (400 Kbps)
  - High-speed (3.4 Mbps) – I2C v.2.0

I2C Connections

- **Two wired bus**
  - **Serial data (SDA) line**
  - **Serial Clock (SCL) line**
- Voltage Levels
  - High -1
  - Low - 0
- **Bit transfer**
  - SCL = 1 implies SDA = valid data
  - **Stable data during high clock**
  - **Data change during low clocks**

I2C Signals

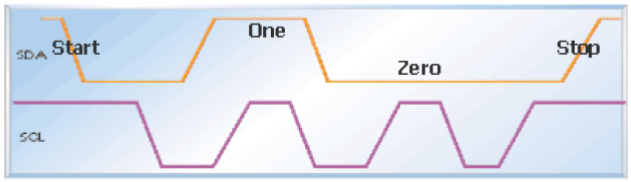

Sample Signals

- Start – high-to-low transition of the SDA line while SCL line is high
- Stop – low-to-high transition of the SDA line while SCL line is high
- Ack – receiver pulls SDA low while transmitter allows it to float high
- Data – transition takes place while SCL is slow, valid while SCL is high

Basic Protocol

- I2C is a master slave protocol
  - Master devices controls the clock (SCL)
  - Slave device may hold the clock low to prevent data transfer
  - No data is transferred unless a clock signal is present
  - All slaves are controlled by the master clock
  - Master and slave are interchangeable

I2C Tradeoffs

- Advantages:
  - Good for communication with on-board devices that are accessed occasionally
  - Easy to link multiple devices because of addressing scheme
  - Cost and complexity do not scale up with the number of devices
- Disadvantages:
  - The complexity of supporting software components can be higher than that of competing schemes (for example, SPI)

8. The SPI protocol (bus configuration, comparation with I2C, protocol, applications).

CH. 6, slide 25-28.

SPI

- Shorthand for “Serial Peripheral Interface”
- Defined by Motorola on the MC68HCxx line of microcontrollers
  - Generally faster than I2C, capable of several Mbps
- Applications:
  - Like I2C, used in EEPROM, Flash, and real time clocks
  - Better suited for “data streams”, i.e. ADC converters
  - Full duplex capability, i.e. communication between a codec and digital signal processor

SPI Bus Configuration

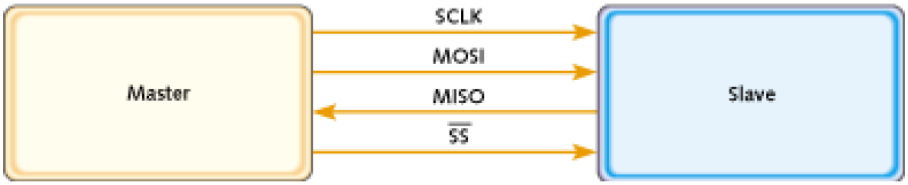

- Synchronous serial data link operating at full duplex
- Master/slave relationship
- 2 data signals:
  - MOSI – master data output, slave data input
  - MISO – master data input, slave data output
- 2 control signals:
  - SCLK – clock
  - /SS – slave select (no addressing)

SPI vs. I²C

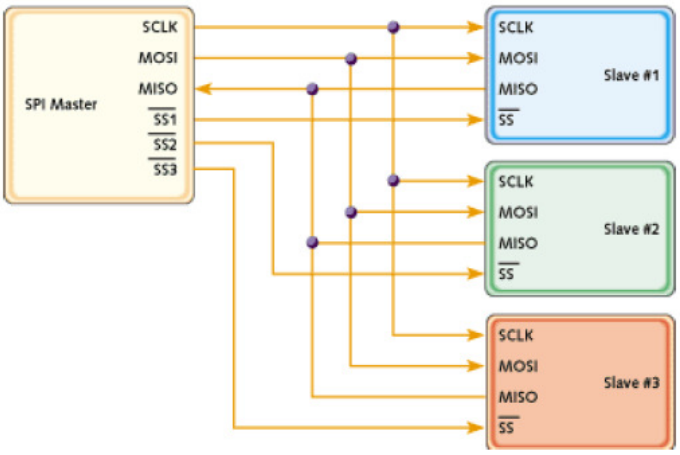

- For point-to-point, SPI is simple and efficient
  - Less overhead than I2C due to lack of addressing, plus SPI is full duplex.
- For multiple slaves, each slave needs separate slave select signal
  - More effort and more hardware than I2C

SPI Protocol

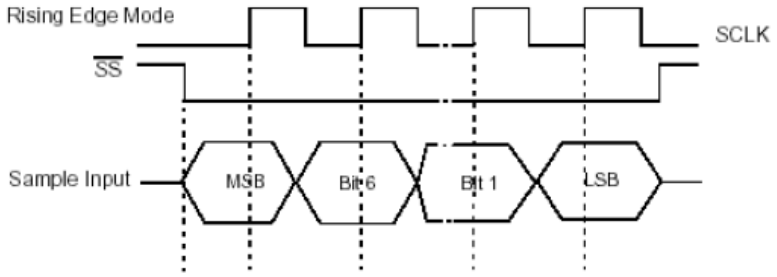

- 2 Parameters, Clock Polarity (CPOL) and Clock Phase (CPHA), determine the active edge of the clock

CPOL	CPHA	Active edge
0	0	Rising
0	1	Falling
1	0	Falling
1	1	Rising

- Master and slave must agree on parameter pair values in order to communicate

9. Steps in software building process (compiling, object files, linking).

CH. 7, slide 15, 17, 18.

Compiling

- Parsing, semantic analysis and object code generation
- Processor specific characteristics to be exploited for efficient code generation
- Cross-compiler or cross-assembler
  - Run on host to produce code for target

Object files

- Contents of an object file can be thought of as a very large, flexible data structure which contains instructions and data resulting from translation process
  - In standard formats
    - COFF: Common Object File Format
    - ELF: Extended Linker Format
  - Object file contains
    - Code blocks – text block
    - Initialize global variable – data block
    - Usually symbol table

Linking for Embedded System

- Merge sections from multiple object files
- Unresolved reference to symbols replaced by reference to actual variables or function calls
- Binding to correct relative address
- A special object file that contains compiled start-up code included
- Start-up code
  - Small block assembly language code that prepares the way for execution of the code

10. Define the kernel and its responsibilities.

CH. 8, slide 12, 13.

Kernel

- Most frequently used portion of OS
- Resides permanently in main memory
- Runs in privileged mode
- Responds to calls from processes and interrupts from devices

Kernel's responsibility

- Managing Processes
- Context switching: alternating between the different processes or tasks
- Various scheduling algorithms
  - Scheduling: deciding which task/process to run next
- Various solutions to dealing with critical sections
  - Critical sections = providing adequate memory-protection when multiple tasks/processes run concurrently