Subject

Cell-based back-end design

Floorplanning

Guideline

Floorplanning is a continuous process that accompanies all steps of VLSI design.

Floorplanning keeps track of

- Partitioning into major building blocks
- Anticipated sizes, shapes and placement of blocks
- Package selection and pin/pad utilization
- Wide busses and electrically critical signals
- Clock domains (frequency, conditional clocking)
- Voltage domains (power dissipation, density, local current needs)
- On-chip power and clock distribution schemes

Example: Add-drop switch for a packet-based data network

- Connects to up to four identical network branches.
- Each port is bidirectional with two pairs of differential signals.
- Actions upon receiving a data packet:
  1. analyze address header,
  2. determine destination,
  3. process header accordingly,
  4. send updated packet off over appropriate network branch.
- Further connects to a local host via a microcomputer interface:
  - $\uparrow\uparrow$ translate data from that interface into the standard packet format before transmitting them through the network (add),
  - $\downarrow\downarrow$ strip the header off each packet addressed to the local host and deliver payload data to that interface (drop).

Example step 0: Overall architecture

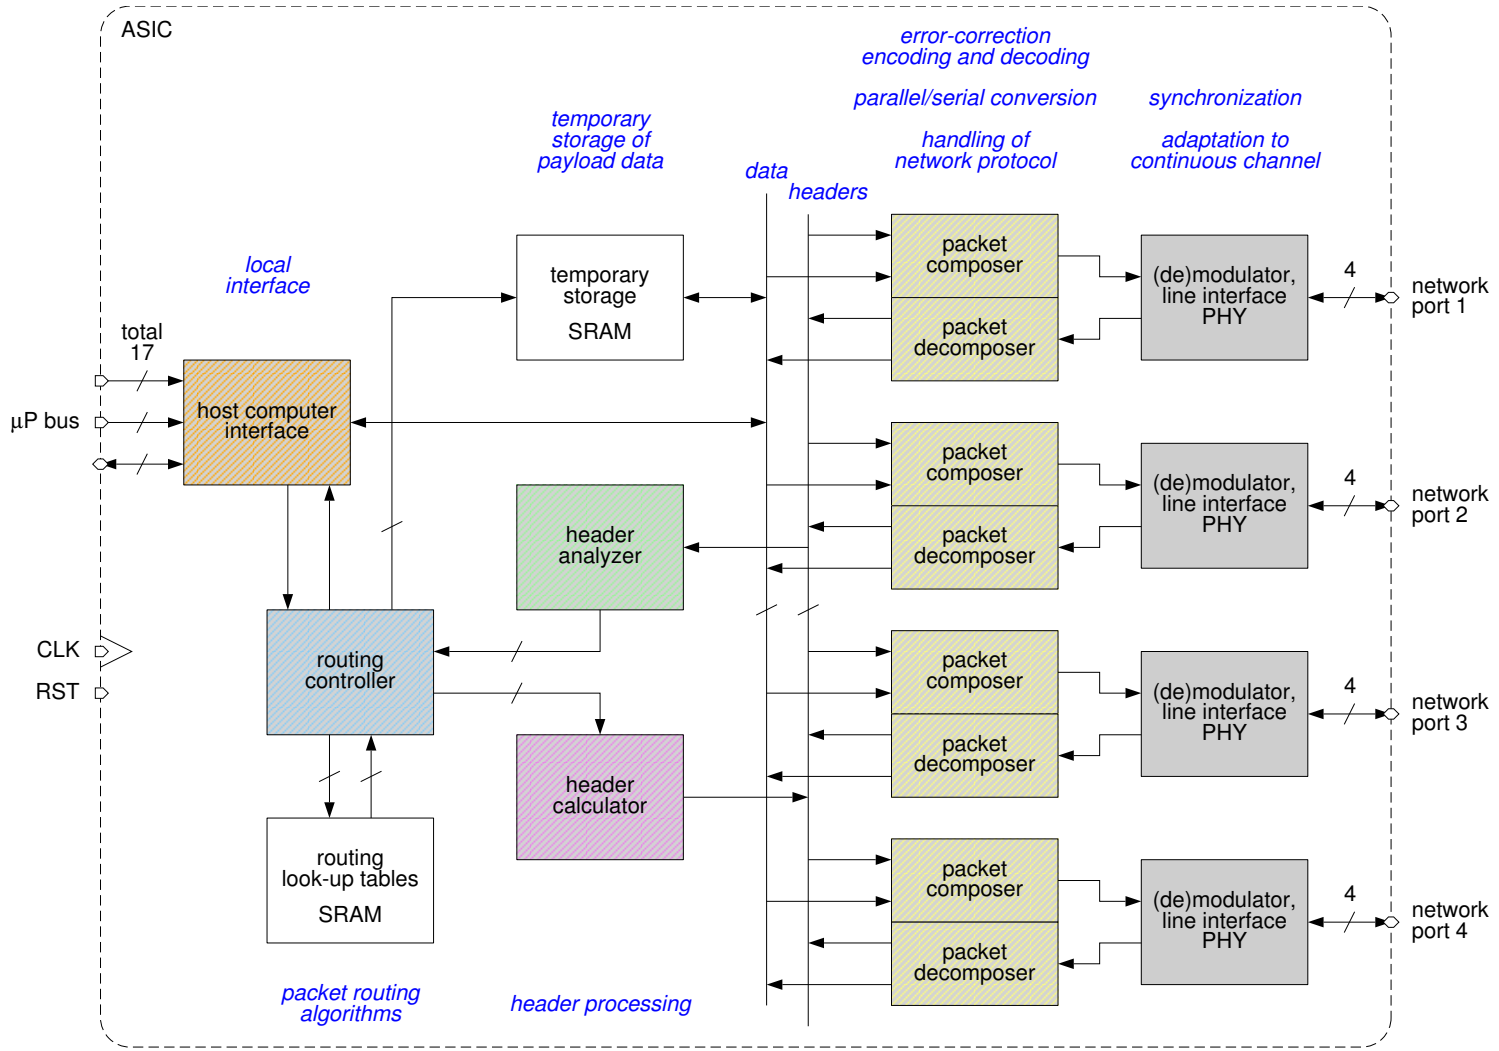

Example step 1: Circuit partitioning

Physical link interface (PHY): combines a modulator/demodulator with line drivers/receivers, 4 instances required, layout to be reused from a previous project.

Temporary data buffer: SRAM macrocell.

Lookup tables for holding routing information: another identical SRAM macrocell.

Packet composer/decomposer: to be implemented with standard cells, to be placed next to PHYs.

Header analyzer/calculator: to be implemented with standard cells.

Routing controller: to be implemented with standard cells.

Host computer interface: to be implemented with standard cells.

Example step 1: Circuit partitioning

Physical link interface (PHY): combines a modulator/demodulator with line drivers/receivers, 4 instances required, layout to be reused from a previous project.

Temporary data buffer: SRAM macrocell.

Lookup tables for holding routing information: another identical SRAM macrocell.

Packet composer/decomposer: to be implemented with standard cells, to be placed next to PHYs.

Header analyzer/calculator: to be implemented with standard cells.

Routing controller: to be implemented with standard cells.

Host computer interface: to be implemented with standard cells.

Clock domain: 1.

Voltage domains: 2 (core plus padframe).

Example step 2: Pin budget

Function	Pins	Subtotal
clock and reset	2	
boundary scan test	5	
Basic equipment		7
network ports	$4 \cdot 4$	
μP port data bits	8	
μP port address bits	4	
μP port control lines	5	
Functionality		33
core ground	4	
core power	3	
padframe ground	7	
padframe power	7	
Power supply		21
Spare		3
Total		64

Example step 2: Pin budget

Function	Pins	Subtotal
clock and reset	2	
boundary scan test	5	
Basic equipment		7
network ports	$4 \cdot 4$	
μP port data bits	8	
μP port address bits	4	
μP port control lines	5	
Functionality		33
core ground	4	
core power	3	
padframe ground	7	
padframe power	7	
Power supply		21
Spare		3
Total		64

Observation

Pin count relative to circuit complexity impacts architectural choices.

Core-limited and pad-limited floorplans.

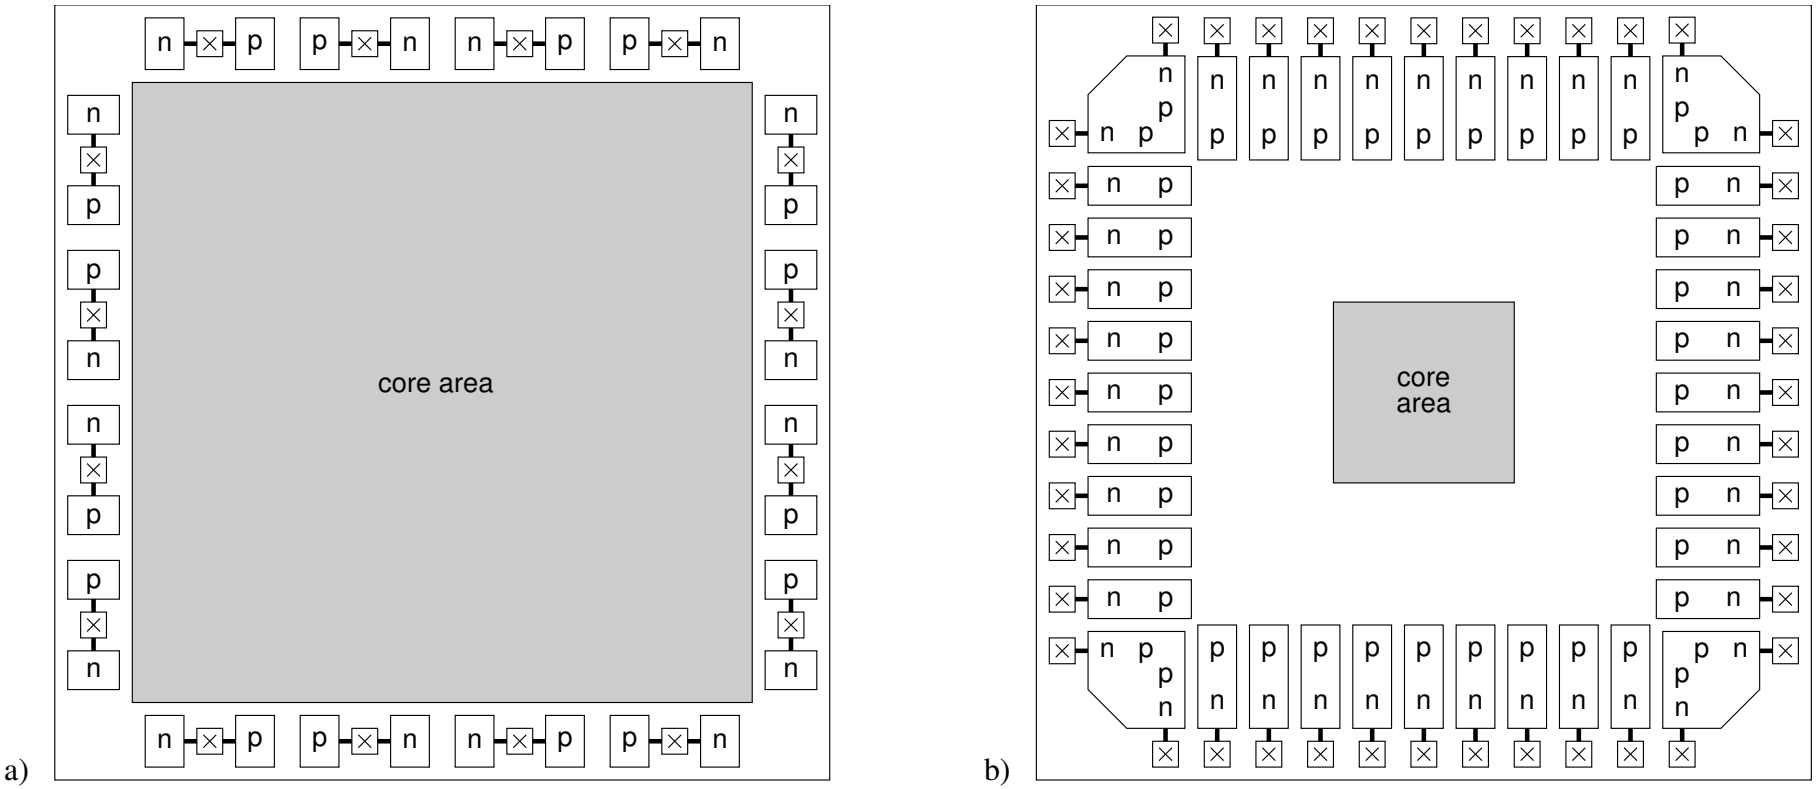

Figure: Die size imposed by core area (a) and by pad count (b).

Observation

Highly pad-limited situations are a waste of die area.

Example step 3: Arrangement of building blocks

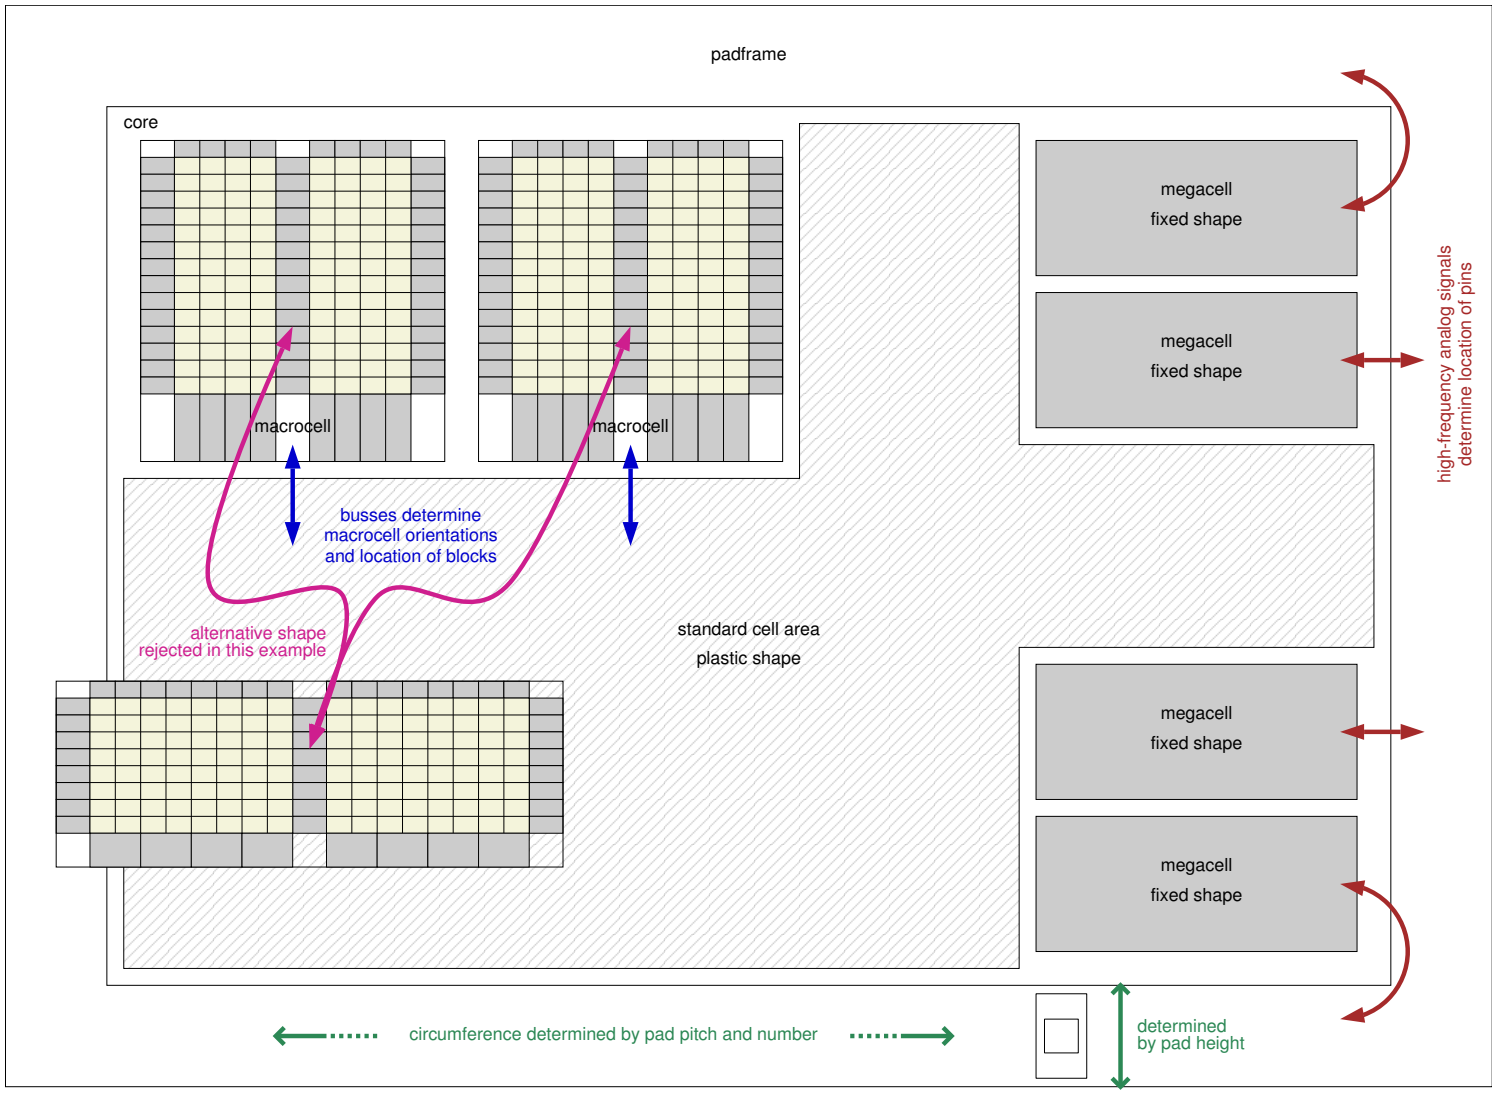

Plan power, clock and signal distribution

Interconnect planning implies holding back five adverse effects

relative importance of adverse effect	power and ground	clock distri-bution	signal wires
resistive voltage losses (*IR*-drops)	strong	less	less
ground bounce and supply droop	strong	n.a.	n.a.
electromigration	strong	less	less
crosstalk	less	strong	strong
interconnect delays	n.a.	strong	strong

- Interconnect resources are best allocated following a priority list, that starts with long and critical nets, such as supply and global clocks, and ends with uncritical local nets.
- Interconnect planning determines the floorplan and sometimes even affects architectural choices (actually more and more so).

Example step 4: Target floorplan prior to place and route

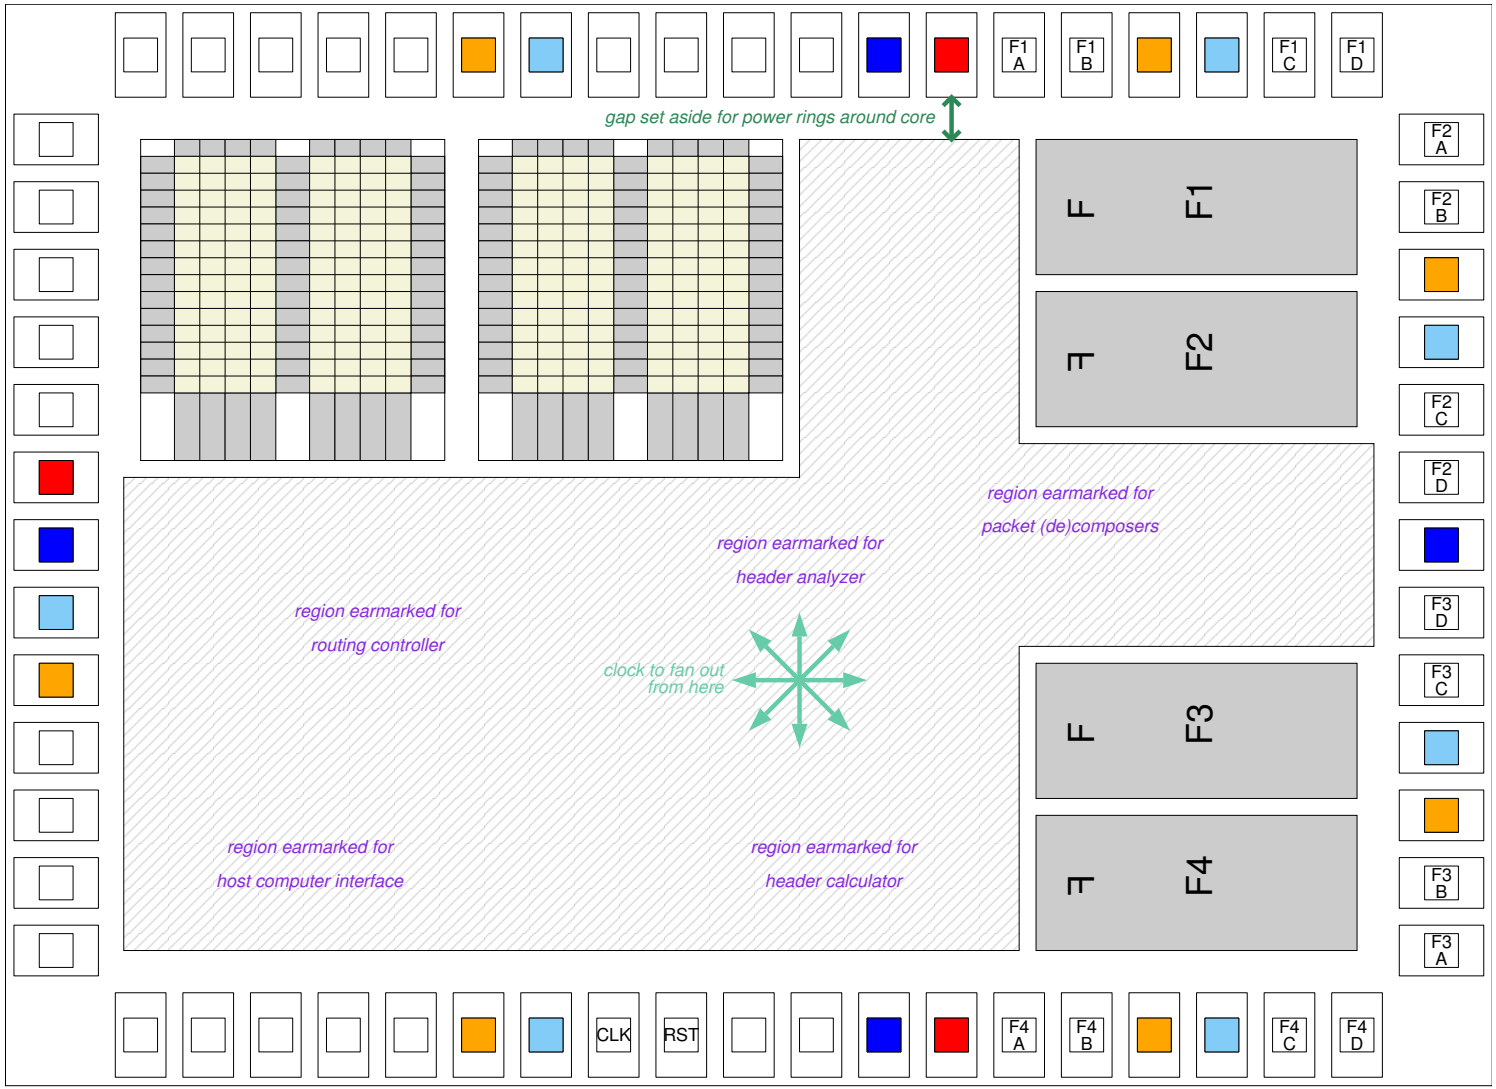

Example step 5: Place and route and chip assembly

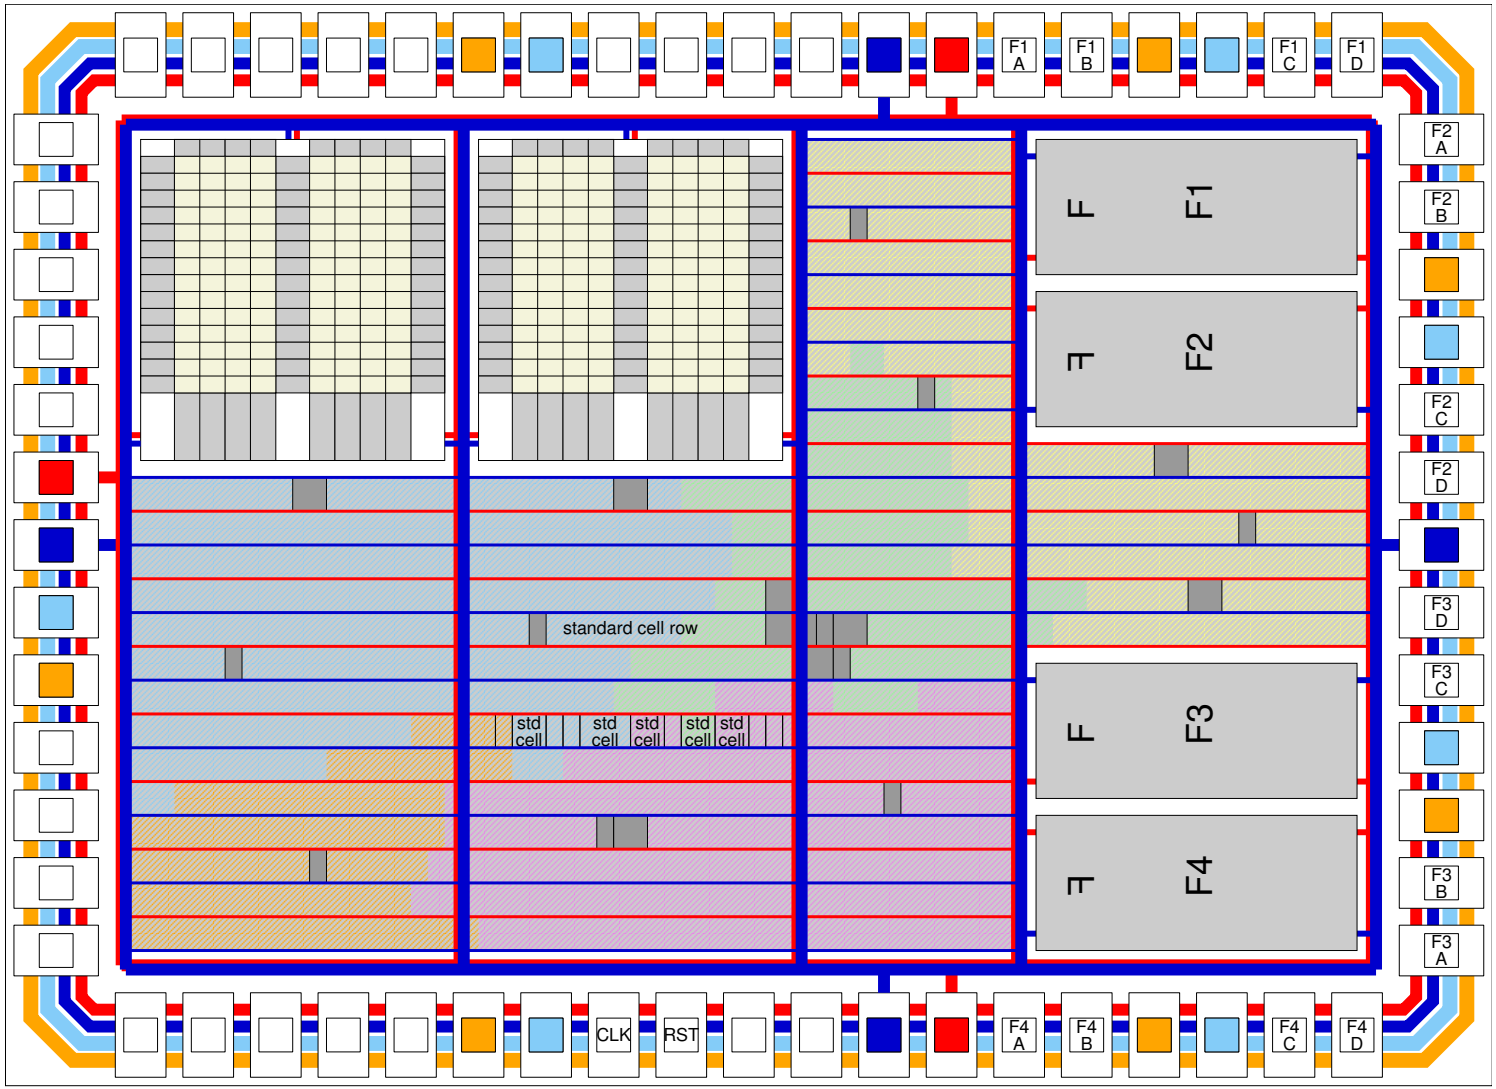

Subject

Packaging

Packaging

What purposes does an IC package serve?

Packaging

What purposes does an IC package serve?

- Protect the die against mechanical stress and environmental attacks.
- Carry away the thermal power
  while keeping the die at an acceptable temperature.
- Provide electrical connection with the surrounding circuitry.
- Expand connector geometry to match next higher assembly level.
- Facilitate the handling of parts during shipping and board assembly.

Package examples

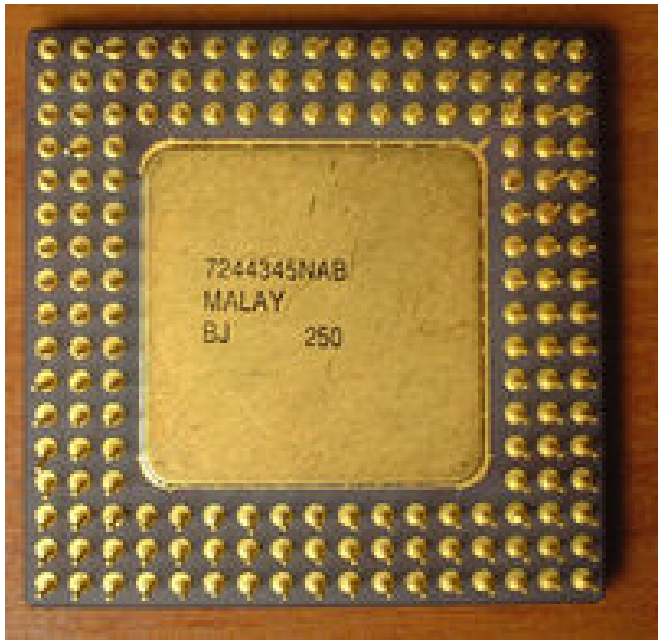

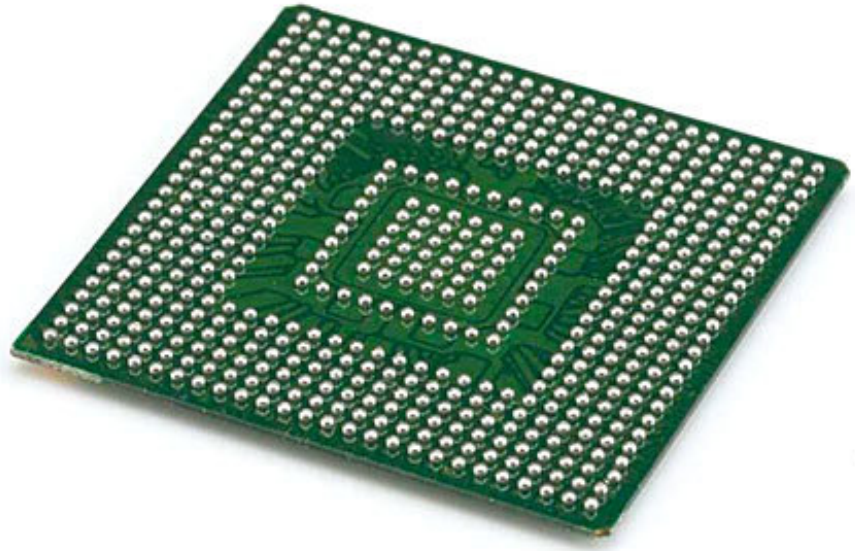

An alphabet soup of popular package types

Package type	Package terminals			
	mount	location	typ. pitch [mm]	typ. count
single in line package (SIP)	TH	1 edge	2.54	2 ... 12
dual in line package (DIP, DIL)	TH	2 edges	2.54	6 ... 40
pin grid array (PGA)	TH	surface	2.54, 1.27	72 ... 478
small outline package (SOP, SOIC)	SM	2 edges	1.27	8 ... 44
thin SOP (TSOP)	SM	2 edges	1.27, 0.8, 0.5	24 ... 86
shrink SOP (SSOP)	SM	2 edges	0.8, 0.65, 0.5	8 ... 70
leadless chip carrier (LLCC, LCC)	SM	4 edges	1.27	16 ... 84
quad flat package (QFP, QFN)	SM	4 edges	1.27, 0.8, 0.65	28 ...
fine-pitch QFP (FQFP)	SM	4 edges	0.5, 0.4, 0.3	... 376
small outline J-leaded pack. (SOJ)	SM	2 edges	1.27	24 ... 44
leaded chip carrier (LDCC, JLCC)	SM	4 edges	1.27	28 ... 84
land grid array (LGA)	SM	surface	1.27, 1.0	48 ...
fine-pitch land grid array (FLGA)	SM	surface	0.8, 0.5	... 1933
ball grid array (BGA)	SM	surface	1.5, 1.27, 1.0	36 ...
fine-pitch BGA (FBGA)	SM	surface	0.8, 0.65, 0.5, 0.4	... 2577

TH stands for through-hole and SM for surface mount.

Testing and encapsulation steps I

Wafer sorting: Generic test structures (such as layout items, isolated devices, and tiny "circuits") typically hidden in the saw lanes[1] between adjacent dies get contacted and measured.

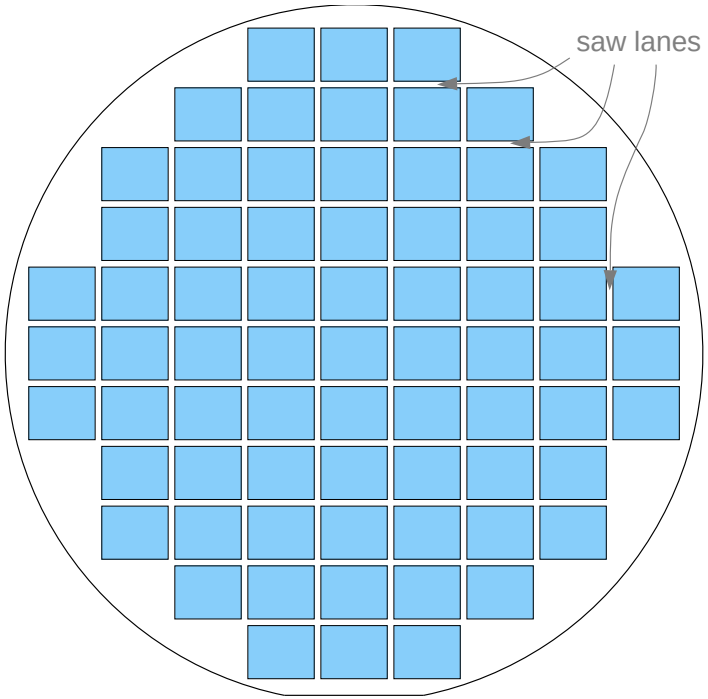

Figure: Processed wafer.

[1] "Kerf" and "Scribe line" are synonyms. Typical lane width 75...120 μm.

Process Control Monitors (PCM)

- Test structures are collected into PCMs and serve to monitor
  - sheet resistances
  - line widths
  - dielectric thicknesses
  - threshold voltages
  - gain factors
  - leakage currents
  - contact/via resistances
  - continuity of long contact/via chains

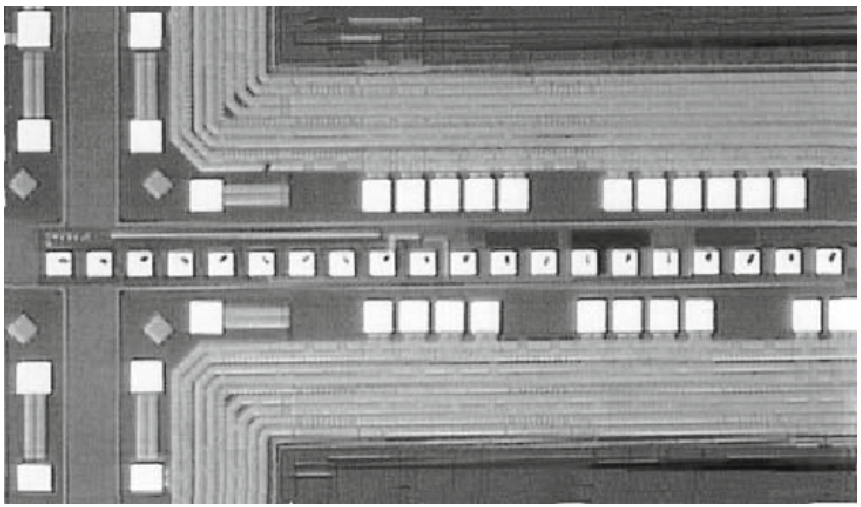

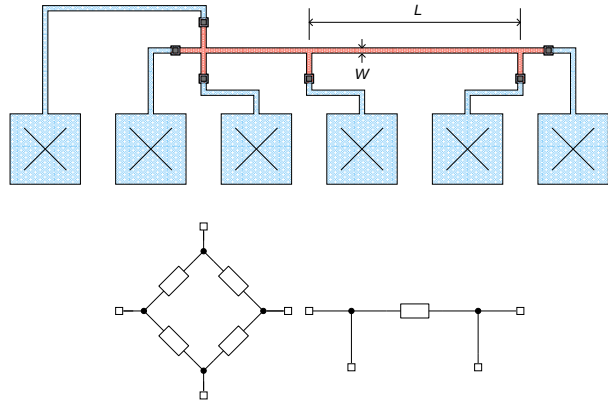

Figure: PCMs located in the saw lane between two dies (photo: William Mann).

Process Control Monitors (PCM)

- Test structures are collected into PCMs and serve to monitor
  - sheet resistances
  - threshold voltages
  - contact/via resistances
  - line widths
  - gain factors
  - continuity of long contact/via chains
  - dielectric thicknesses
  - leakage currents

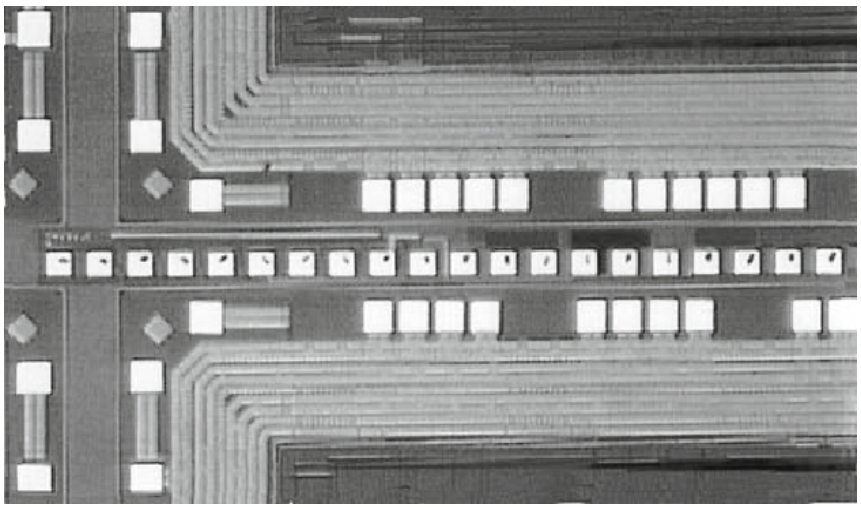

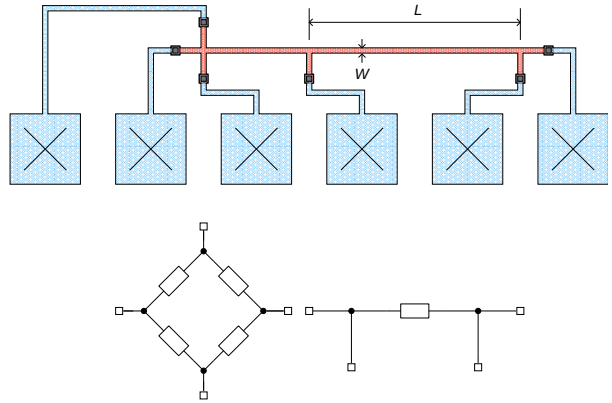

Figure: PCMs located in the saw lane between two dies (photo: William Mann).

$\rightsquigarrow$ Wafers found to suffer from defects or excessive parameter variations are scrapped.

Testing and encapsulation steps II

Wafer testing: Dies contacted one by one for functional tests by way of stimulation and response checking $\mapsto$ "go/no go".

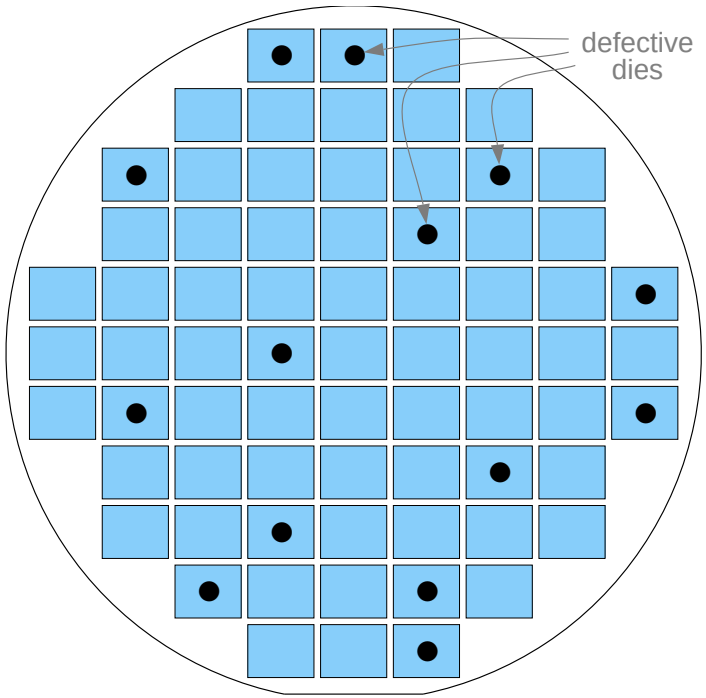

b)

Figure: Defective circuits inked.

Testing and encapsulation steps II

Wafer testing: Dies contacted one by one for functional tests by way of stimulation and response checking $\mapsto$ "go/no go".

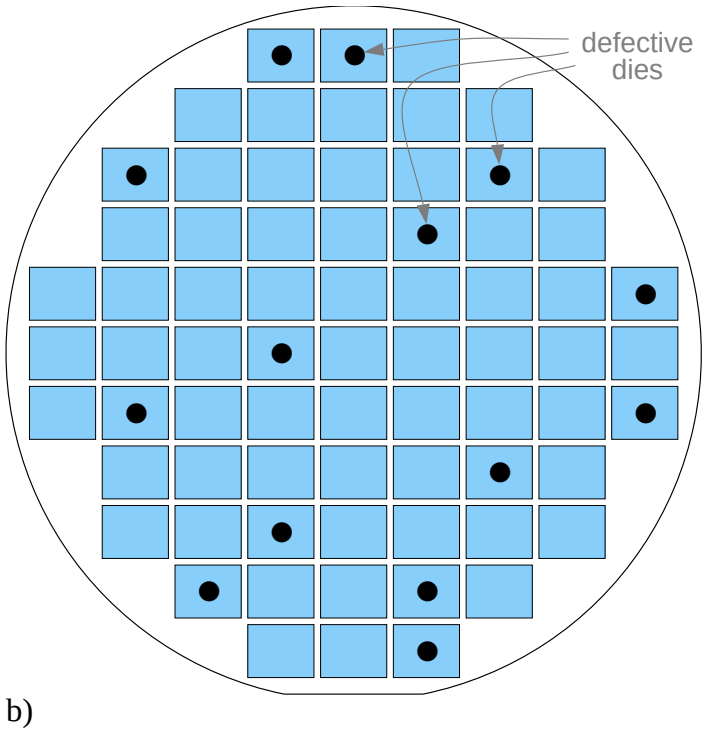

Figure: Defective circuits inked.

$\rightsquigarrow$ Dies found to malfunction are barred from packaging.

Probe cards

- A set of ultrafine needles lowered on the wafer to electrically contact one circuit at a time.

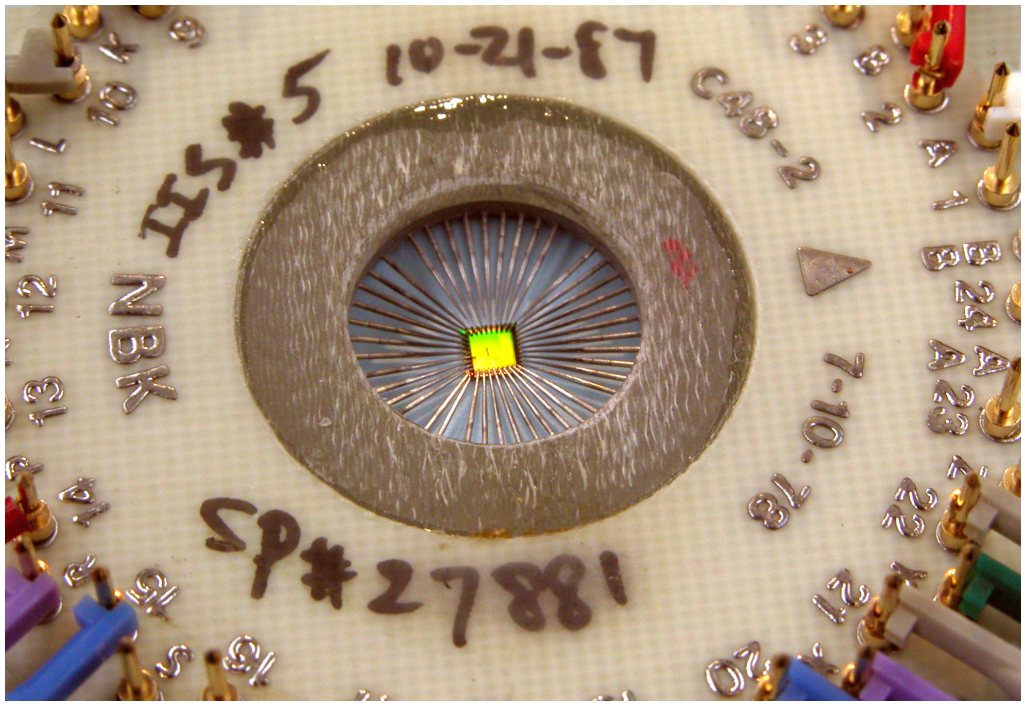

Figure: A circuit under test (CUT) contacted by a probe card.

$\rightsquigarrow$ Results kept in electronic records (no more inking).

Testing and encapsulation steps III

Backgrinding: from ≈750 μm to 150 μm or even down to 50 μm.

Singulation: Two orthogonal series of parallel saw cuts followed by stretching of elastic carrier tape (aka "blue film").

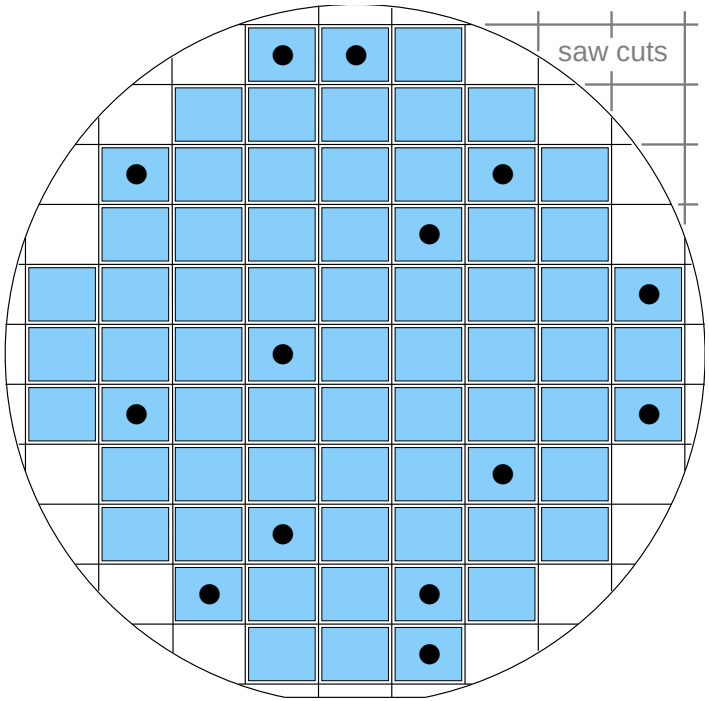

Figure: Wafer after sawing.

Testing and encapsulation steps IV

Die bonding: Die gets fastened in cavity by means of solder or epoxy compound.

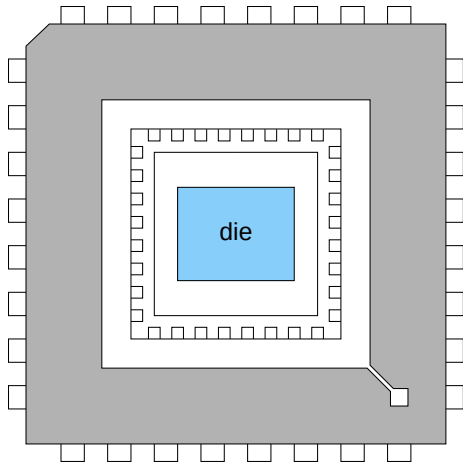

Figure: Good die attached to package cavity.

Testing and encapsulation steps V

Wire bonding: Bond pads on the silicon die get connected to package leads with Au, Ag or Cu wires (of diameter 15, 17, 20, 25 or 33 μm).

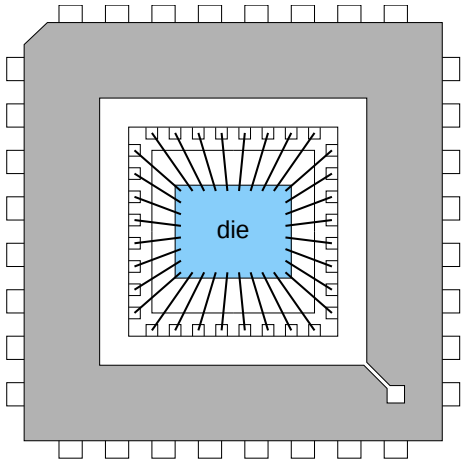

e)

Figure: Wire bonding completed.

Automatic equipment operates at a rate $\geq$ 6 wires/s.

Testing and encapsulation steps VI

Sealing: Metal or ceramic lid (ceramic package) or molding (plastic package).

Final testing:

- Another series of functional "go/no go" tests.
- Measurement of electrical and timing parameters.
- Binning according to speed, power and/or leakage (optional).

Stamping/laser marking: Part designation, speed/power grade, production lot.

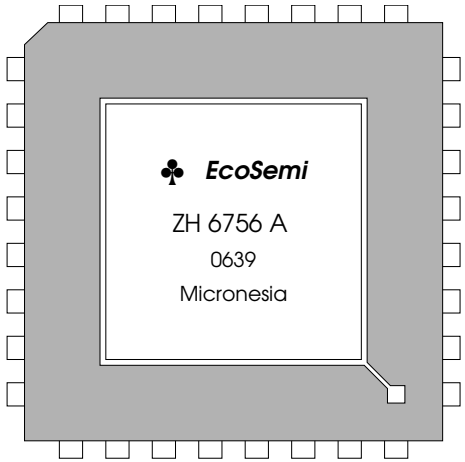

f)

Figure: Final IC.

Bonding diagram with related terms and rules

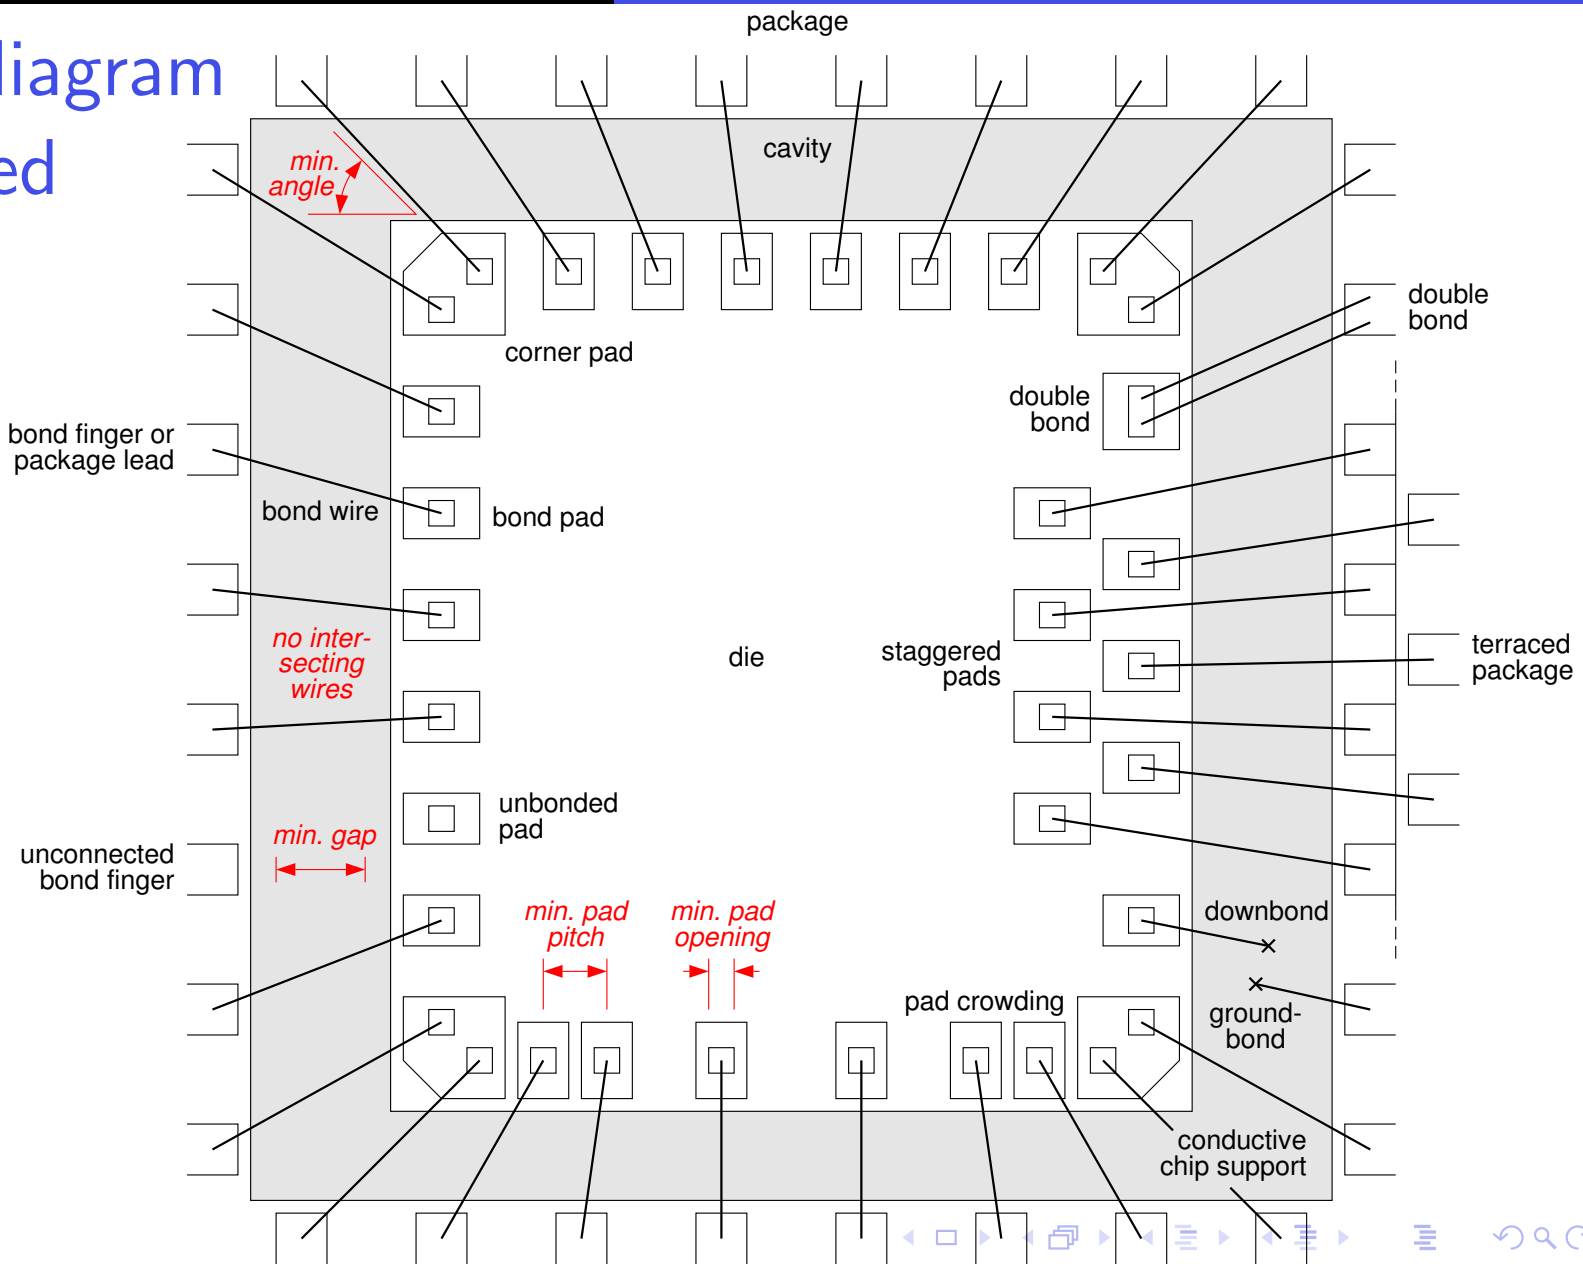

A set of conservative bonding rules

Bonding normally poses no problems if

- No intersecting wires
- Die not overly elongated ($\frac{1}{2} \leq \textit{aspect ratio} \leq 2$)
- Minimum gap of 0.6 mm between cavity and die on all four sides
- 25 μm bond wires no shorter than 1.0 mm and no longer than 4 mm
- No downbonds, groundbonds and double bonds
- Minimum angle of 45° between bond wires and chip edge
- Square bond areas with overglass opening of 75 μm by 75 μm or more
- Minimum pad pitch 90 μm

Pinout and bonding I

Hints from practical experience

- Some package types provide special low-impedance leads for power and ground. $\rightsquigarrow$ Take advantage of them!
- Some packages have the chip support internally connected to one of their pins. $\rightsquigarrow$ Watch out for shorts, connect to VSS or VDD depending on whether an n- or a p-well process is being used!
- Components mistakenly rotated by 90, 180 or 270° are a frequent mishap, both with dies in cavities and with IC packages on PCBs.
  $\rightsquigarrow$ Clearly identify pin 1 on the die and on the package!

Pinout and bonding II

Beware of traps

- The inner routing of package leads need not necessarily be the same for packages of the same type and pin count, PGAs are notorious.
- Two packages that share the same geometric size, pin count, and pitch do not necessarily have their pins located at the same places.[2]

Lesson learned

- Make sure to know all electrical, mechanical and thermal details of your packages, sockets and test sockets before ordering any parts. Check availability too.

[2] As an example, 8 mm by 8 mm 56-pin quad flat no lead (QFN) packages that conform with the JEDEC MO-220K standard have 14 pins along each edge. Yet, industry has also produced QFN-56 packages with 13 pins on each side plus one pin at each (slant) corner.

High-performance packages I

Wire bonding has drawbacks for high-frequency high-power VLSI chips:

- Connections confined to chip periphery, hence limited in number.
- On-chip power distribution difficult.
- Bond wires and package lead fingers add to parasitic impedance.

Better approach:

- Cover the entire die area with connections (bumps).
- Flip the chip to make its bottom surface available for heat evacuation.
- Place local bypass capacitors next to the die.

High-performance packages II: LGA, BGA and PGA

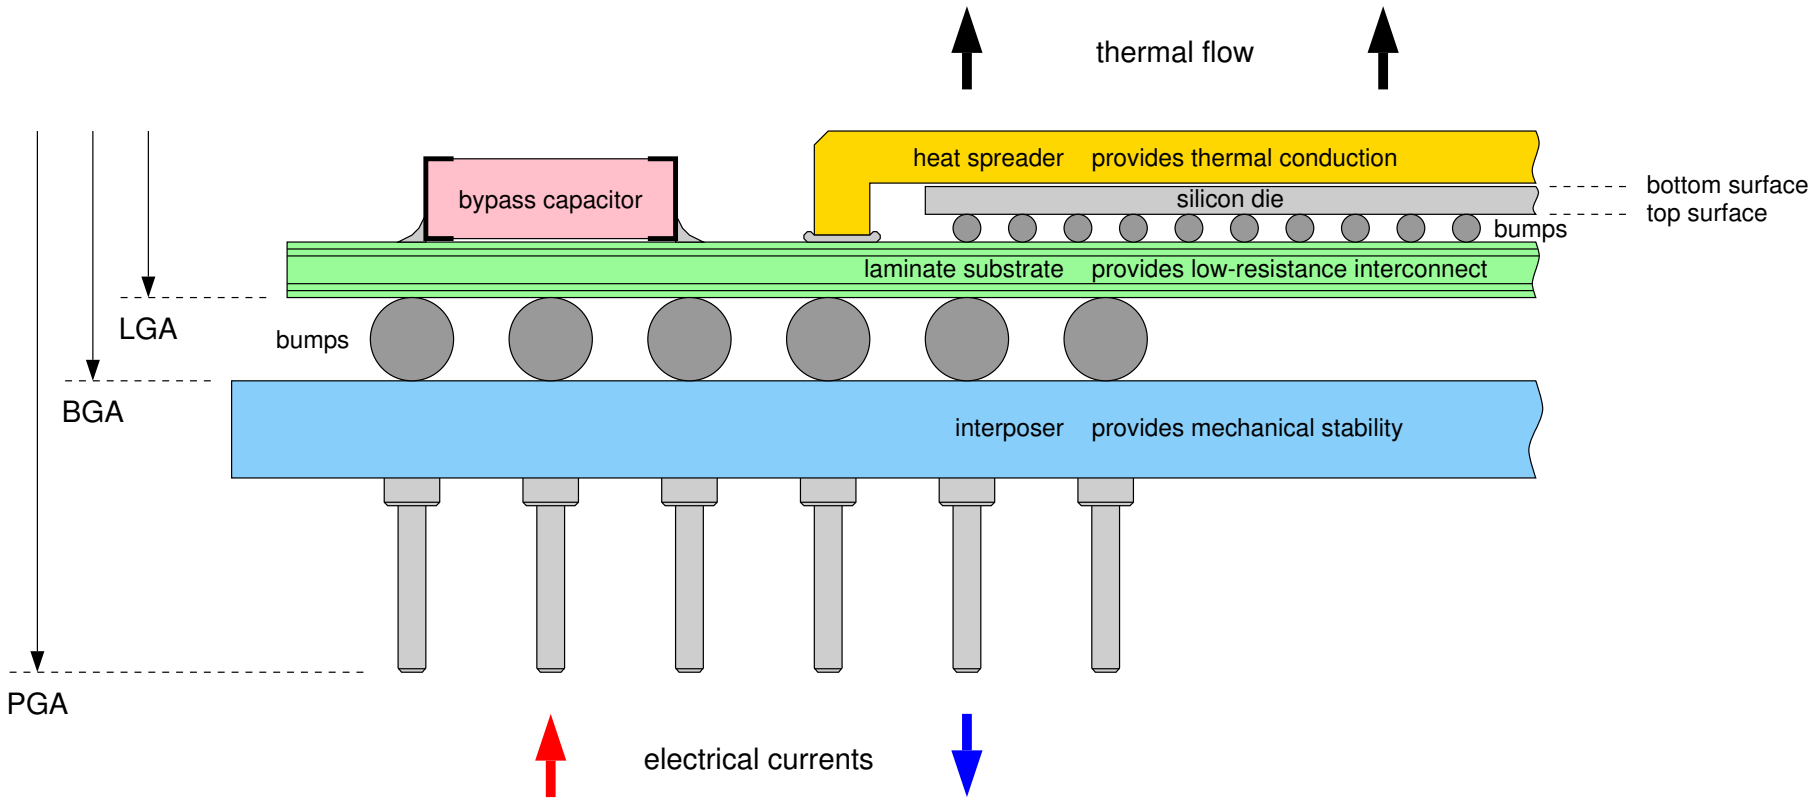

Figure: Cross section through a high-performance package.

LGA = land grid array, BGA = ball grid array, PGA = pin grid array.

High-performance packages III

Observation

Laminate substrates can be understood as an extension of the silicon chip that makes available more metal layers of even lower sheet resistance.

Compare on- and off-chip interconnect resources:

Location	conduct. material	min. pitch [μm]	thick-ness [μm]	sheet resist. [m$\Omega/\square$]	insulat. material	rel. diel. constant [ ]
on-chip upper metal	Al, Cu	0.8	0.9...1.5	30..50	SiO_2	3.9
ceramic substrate	Ag	250	8...10	0.2	Al_2O_3	7...8
epoxy laminate	Cu	70	17	0.01	FR4	≈ 4

High-density packaging I: multi-chip module (MCM)

- Bare dies and tiny SMD components mounted and interconnected on a small substrate before being encapsulated in a common package.

+ Reduction in overall size.
+ Lower parasitics.
+ Each subsystem manufactured with its optimum technology.
+ Each subsystem tested separately.
+ Lower initial costs.
+ Shorter turnaround time.
+ No fragmentation of fabrication volume into multiple varieties.
– Substrate and mounting costs.

Example: Satellite TV Switch and Equipment Control

- An ASIC for switching RF between 5 inputs and 4 outputs existed.
  A new product was to handle 9 inputs and to include four DiSEqCs.

Example: Satellite TV Switch and Equipment Control

- An ASIC for switching RF between 5 inputs and 4 outputs existed.
  A new product was to handle 9 inputs and to include four DiSEqCs.
- System-in-package instead of developing new mixed-signal ASIC.
  - 2 RF switch ASICs
  - 1 glue logic chip
  - 4 DiSEqCs (Digital Satellite and Equipment Control)
  - passives (coupling capacitors, termination resistors)

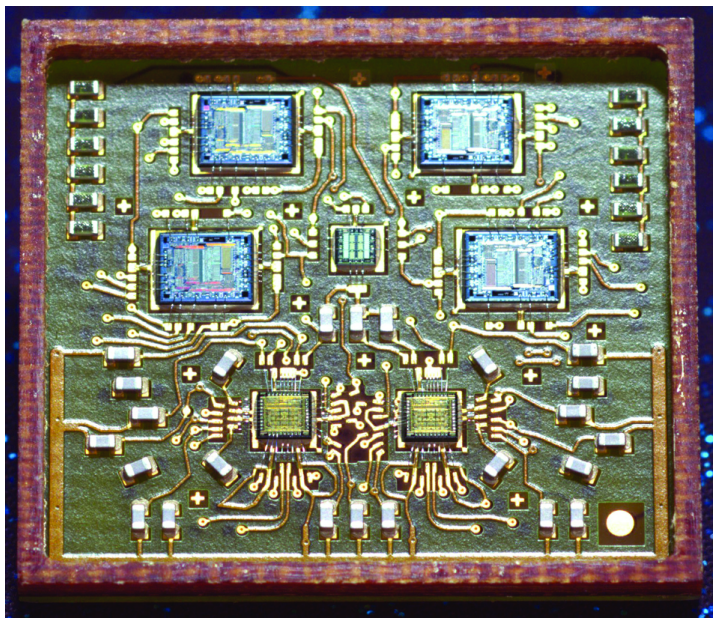

Figure: Seven dies plus discretes mounted on a laminate substrate.

High-density packaging II: Folded flexiprints

- Tested chips and discretes are surface-mounted on a flexible film substrate.
- The flexiprint is then cut and folded.

\+ Can be designed to fit into an irregular volume.

Examples: Hearing aids, digital cameras, automotive electronics.

High-density packaging III: Chip stacking

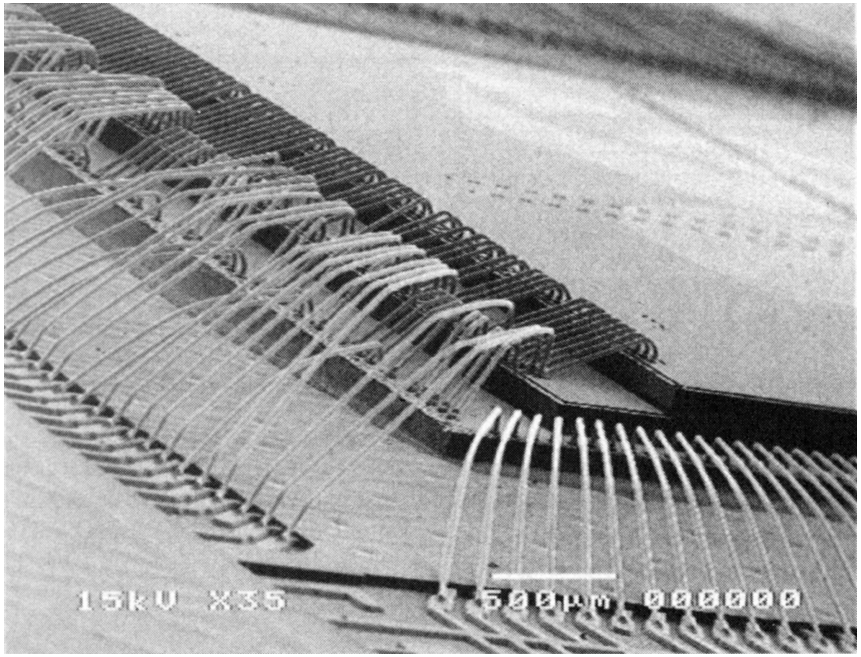

Figure: Three stacked chips interconnected by wire bonding

+ Supports mix-and-match composition of technologies.
  E.g. digital logic, RF front end, commodity RAM, and flash memory, each fabricated with a specialized process and in larger quantities.

High-density packaging IV: Through-silicon vias (TSV)

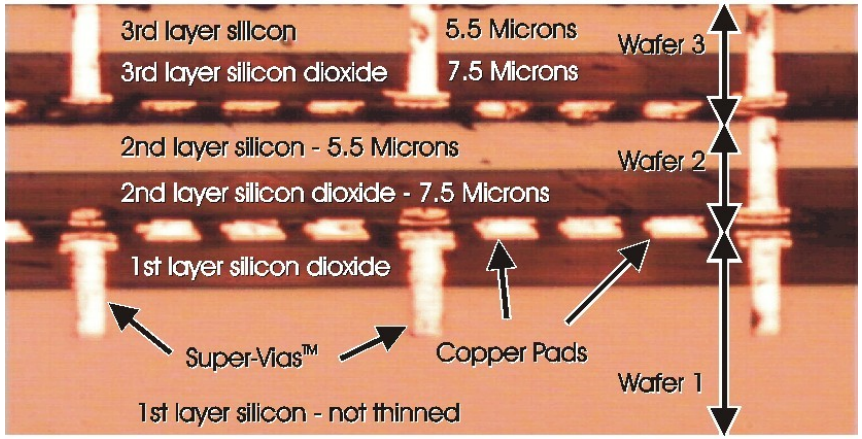

Figure: Cross section through chip stack with TSVs (photo by Tezzaron).

+ Minimal wiring parasitics and, hence, maximum energy efficiency.
+ May become more economic than further (horizontal) technology scaling.
∼ A few problems still need to be solved
  - Heat evacuation
  - Mismatch between thermal expansion of Cu and Si
  - Cu must not be allowed to diffuse into Si (poisonous to MOSFETs)
  - 3D "floorplanning", routing, and parasitics extraction
  - Industry standards

Example: Stacked Silicon Interconnect (aka 2.5D packaging)

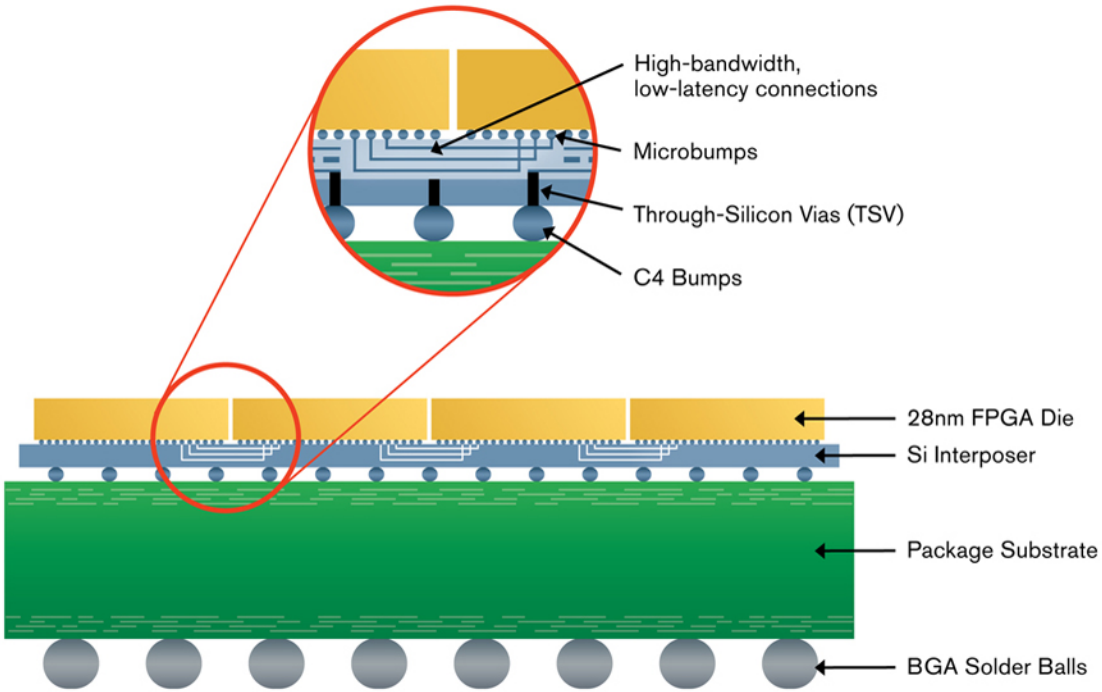

Figure: Cross section through a Virtex-7 FPGA (drawing by Xilinx).

- Extra silicon interposer: manufactured in 65 nm technology, purely passive, includes TSVs

\+ Allows for finer geometries than a laminate or ceramic substrate
$\rightsquigarrow$ more and faster die-to-die connections.

Opting for an adequate packaging technique

Criteria:

- Space available for encapsulated chip(s) on printed circuit board
- Board mounting technique (SM, TH, COB, etc.)
- Number of pins required and closest number actually available
- Cavity size along with range of acceptable die sizes
- Electrical characteristics, i.e. package parasitics
- Maximum power, ambient temperature, thermal resistance
- Resistance against mechanical, thermal, and environmental stress
- Expected lifetime, aging, reliability
- Graded product range from a limited inventory of parts
- Required equipment for automatic packaging, testing and mounting
- Yield losses due to packaging and mounting operations
- Packaging and mounting costs

Insight gained

- Several options are available for high-density packaging
  - System-on-Chip (SoC)
  - Multi-chip module (MCM)
  - System-in Package (SiP)
  - Folded flexiprints
  - Cubing and Package-on-Package (PoP)
- High density packaging is an alternative to monolithic integration for products that sell in moderate quantities and/or multiple varieties, or that combine incompatible technologies.

Guideline

Chip, package, board-level wiring, signal integrity, and heat evacuation must be understood together as critical design elements.

- The future may bring combinations of VLSI, chip stacking, and composite substrates that embed bypass capacitors and other passives.

The costs of Integrated Circuits

"What makes up the overall costs of a microelectronic subsystem?"

The costs of Integrated Circuits

"What makes up the overall costs of a microelectronic subsystem?"

As always, expenses fall into either of two categories.

Non-recurring costs cover everything that must be paid for before volume production can begin.

Recurring costs depend on the quantity produced.

Non-recurring costs of Integrated Circuits

- Project management including negotiations with business partners
- Circuit specification
- Purchase and assimilation of virtual components
- Circuit design and verification
- Preparation of testbenches and vectors for simulation
- CAE/CAD related expenses
- Sign-off procedure (NRE)
- Preparation of fabrication masks (NRE)
- Setting up of fabrication and testing facilities (NRE)
- Prototype fabrication and testing
- Redesigns, if any
- Product qualification (life-cycle tests, JEDEC standards, etc.)

Recurring costs of Integrated Circuits

- Supply chain management
- Raw semiconductor wafers
- Wafer processing
- Volume testing
- Process monitoring and yield enhancement
- Packaging
- Royalties for virtual components
- Board space
- Component mounting

Cost structure of Integrated Circuits

Overall costs per unit are

$$c = \frac{c_0}{n} + c_1$$

where

- c_0 sum over all non-recurring expenses
- n number of working circuits produced
- c_1 cost increment per circuit produced
  (for raw wafers, fabrication, testing, packaging, etc.)

Cost structure of Integrated Circuits

Overall costs per unit are

$$c = \frac{c_0}{n} + c_1$$

where

c_0	sum over all non-recurring expenses
n	number of working circuits produced
c_1	cost increment per circuit produced (for raw wafers, fabrication, testing, packaging, etc.)

The actual numbers and the optimum choice are intimately related to

- Circuit size
- Fabrication process
- Production volume
- whether implementation is as an ASIC or with field-programmable devices

IC manufacturing costs I

Putting aside testing and packaging, manufacturing one functional die costs

$$c_f = \frac{c_{wr} + c_{wp}}{n_f}$$

where

c_f	cost of one functioning die
c_{wr}	cost of one raw wafer
c_{wp}	cost of processing one wafer
n_f	defect-free dies obtained per wafer

Packing square cookies into a round box

Rectangular dies are being manufactured from round wafers

$$n_m \approx \frac{\pi}{A_d}\left(\frac{d_w}{2} - \sqrt{A_d}\right)^2$$

where

n_m	number of dies manufactured per wafer
d_w	wafer diameter
A_d	die area (including scribe lines)

Yield losses due to functional defects I

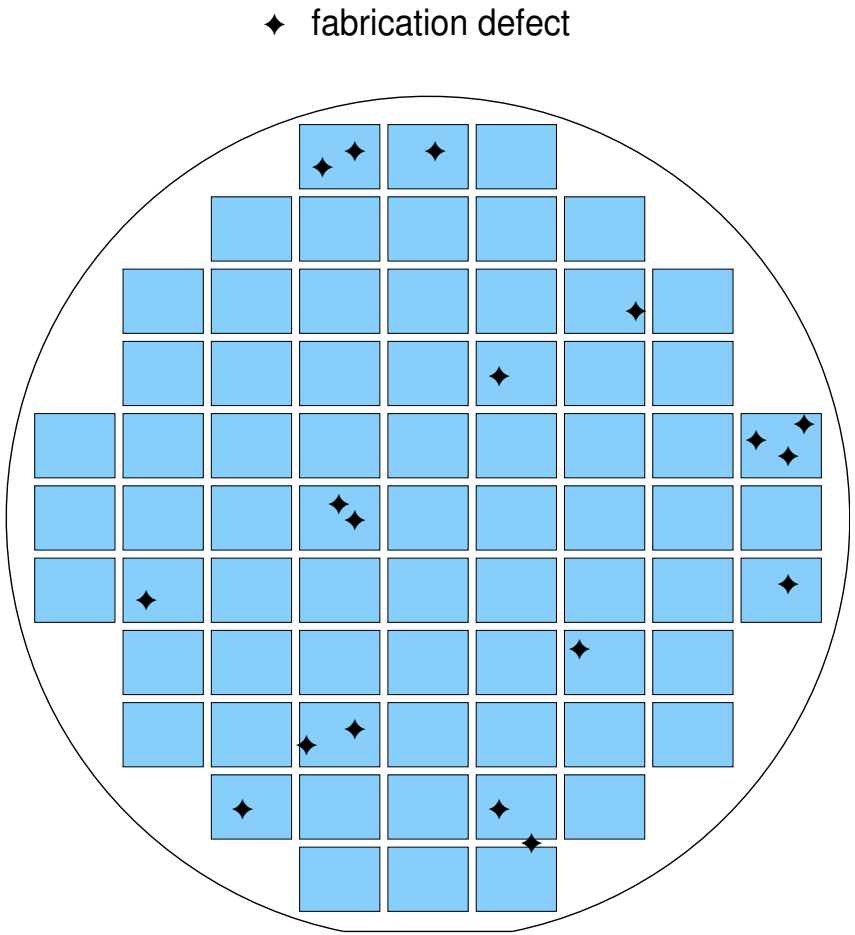

Figure: Fabrication defects occur distributed across wafer surface.

Yield losses due to functional defects II

- As die area gets larger, the probability of finding a defect increases.

$\rightsquigarrow$ Empirical yield models such as the negative binomial model:

$$y_f = \frac{n_f}{n_m} \approx (1 + \frac{DA_c}{\alpha})^{-\alpha}$$

where

y_f	fabrication yield ($y_f \leq 1$)
n_f	number of functional dies per wafer ($n_f \leq n_m$)
n_m	number of manufactured dies per wafer
D	defect density (improves with maturity of manufacturing process)
	$D \approx 0.004$ mm^{-2} was typical for the 90 nm generation in 2006
α	clustering factor (loosely related to number of lithographic steps)
	$\alpha \approx 4.0$ for the 90 nm generation
A_c	die area occupied by actual layout structures ($A_c < A_d$)

IC manufacturing costs II

With edge effect and fabrication yield put together, the manufacturing of one functional die costs

$$c_f = \frac{c_{wr} + c_{wp}}{n_m\, y_f} \approx (c_{wr} + c_{wp})\,(1 + \frac{DA_c}{\alpha})^{\alpha}\, \frac{A_d}{\pi(\frac{d_w}{2} - \sqrt{A_d})^2}$$

IC manufacturing costs II

With edge effect and fabrication yield put together, the manufacturing of one functional die costs

$$c_f = \frac{c_{wr} + c_{wp}}{n_m\, y_f} \approx (c_{wr} + c_{wp})\,(1 + \frac{DA_c}{\alpha})^{\alpha}\,\frac{A_d}{\pi(\frac{d_w}{2} - \sqrt{A_d})^2}$$

Observation

Only for small circuits do fabrication costs grow roughly proportionally with die size. For truly complex circuits, in contrast, expenses increase in a highly progressive way.

- Defect density determines the largest die that can be manufactured at some given cost.
- High-capacity RAMs often include redundant rows or columns that get selectively switched-in after fabrication in lieu of malfunctioning locations.

IC manufacturing costs III

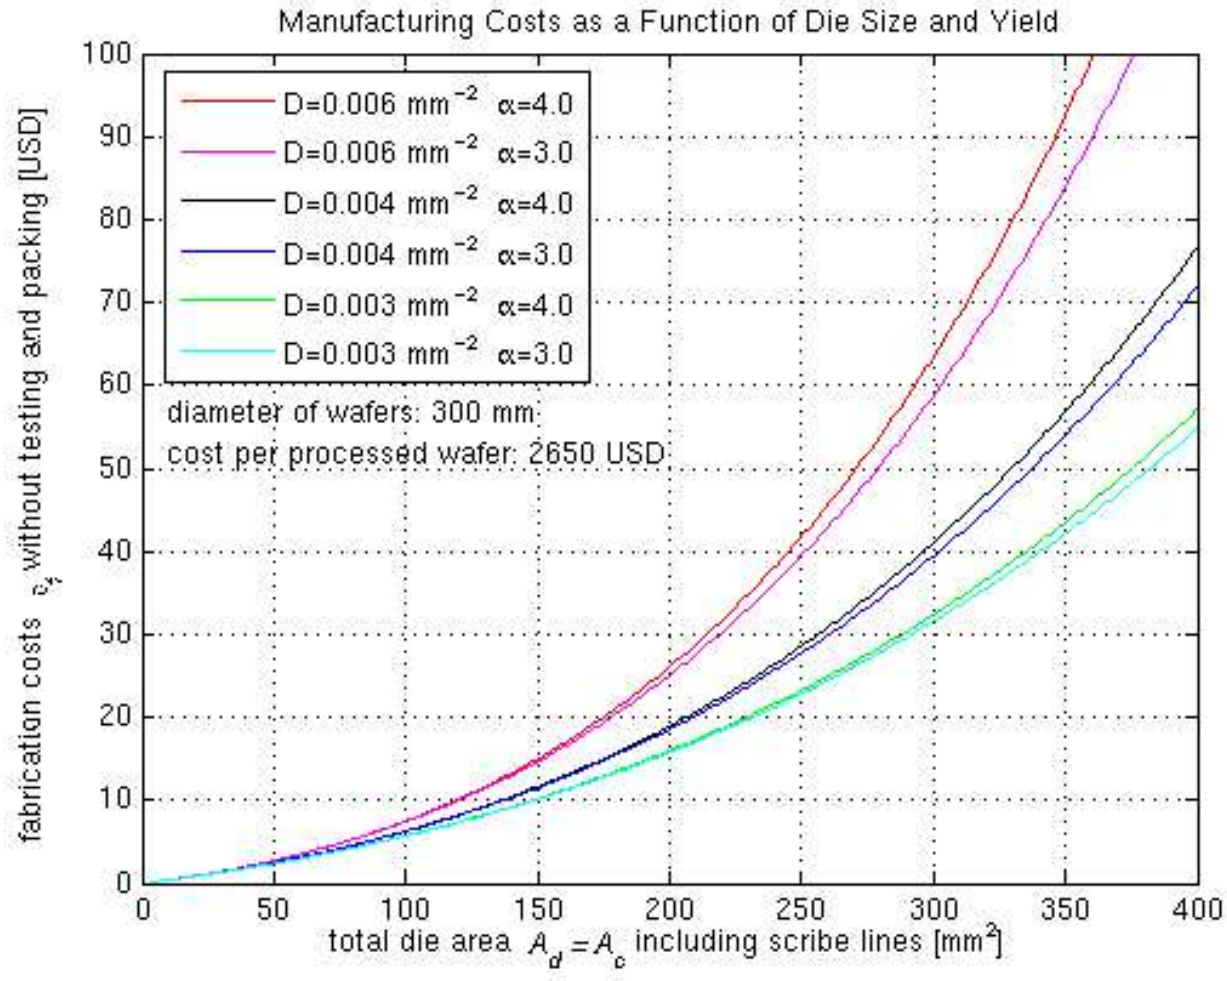

Figure: Manufacturing costs a function of die size and yield.

The impact of the fabrication process I

Costs for a mask set explode from one generation to the next

- Sub-wavelength resolution mandates sophisticated resolution enhancement techniques (PSM, OPC)
- Number of photolithographic masks keeps growing

year of introduction	process generation [nm]	metal levels	range of costs for a mask set [USD]
1995	350	4	50k ... 60k
1997	250	5	90k ... 110k
1999	180	6	250k ... 300k
2001	130	7	500k ... 700k
2004	90	8	800k ... 1.1M
2007	65	9	1.2M ... 1.5M
2010	45	10	1.6M ... 2M

The impact of the fabrication process II

Rough cost model

$$c_f = \frac{c_{wr} + c_{wp}}{n_f} \approx \frac{c_{wr} + \sum_{p=1}^{P} c_{ls}(p)}{n_m \, y_f}$$

where

c_f	cost of one functioning die
c_{wr}	price of one raw wafer
c_{wp}	costs of processing per wafer
n_f	number of functional dies per wafer
P	number of lithographic steps
$c_{ls}(p)$	costs associated with the p'th lithographic step
n_m	number of manufactured dies per wafer
y_f	fabrication yield

Numerical example

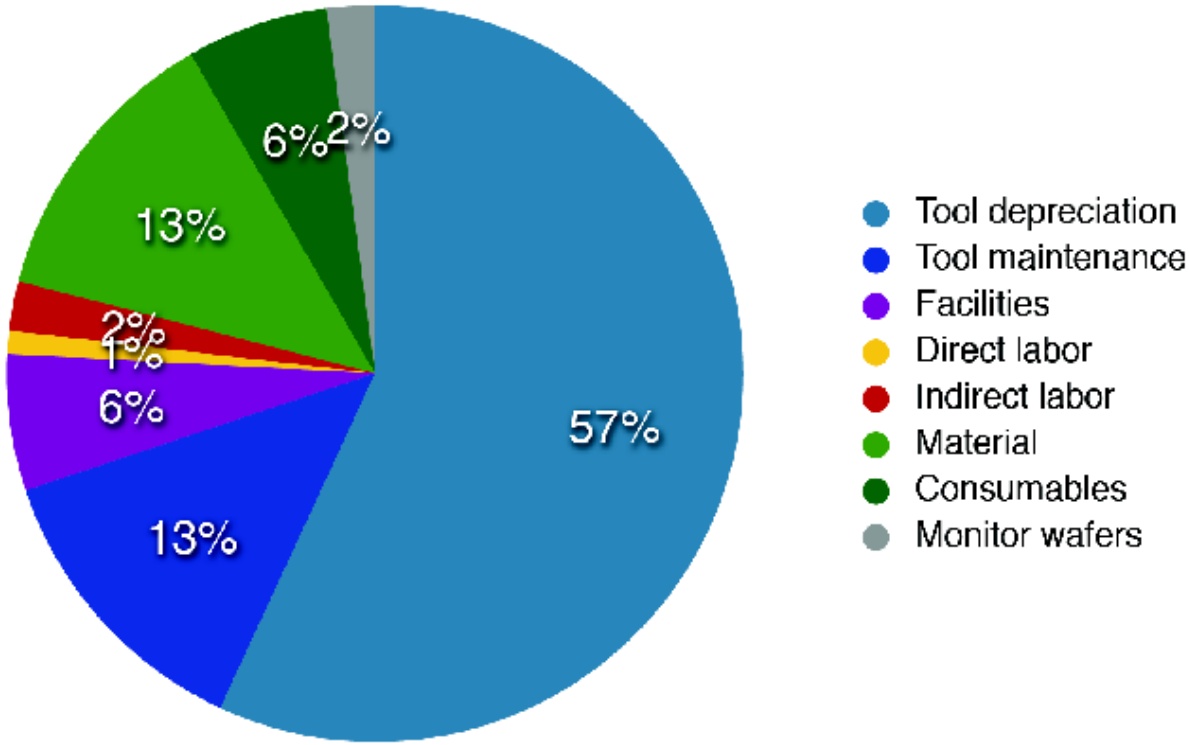

Figure: Cost breakdown for 300 mm wafers in 65 nm 1P9M Cu CMOS technology processed by a Taiwan-based fab (numbers courtesy of IC Knowledge).

Insight gained

Observation

Industry moves over to the next process generation when the savings from fabricating a given circuit in a denser process compensate for the more expensive masks and wafer processing.

- The fact that yield is lower initially is taken into account.

Insight gained

Observation

Industry moves over to the next process generation when the savings from fabricating a given circuit in a denser process compensate for the more expensive masks and wafer processing.

- The fact that yield is lower initially is taken into account.

Observation

The high proportion of capital expenses renders processing costs highly dependent on a wafer fab's utilization, and the same applies to profitability.

Yield losses due to parametric variability I

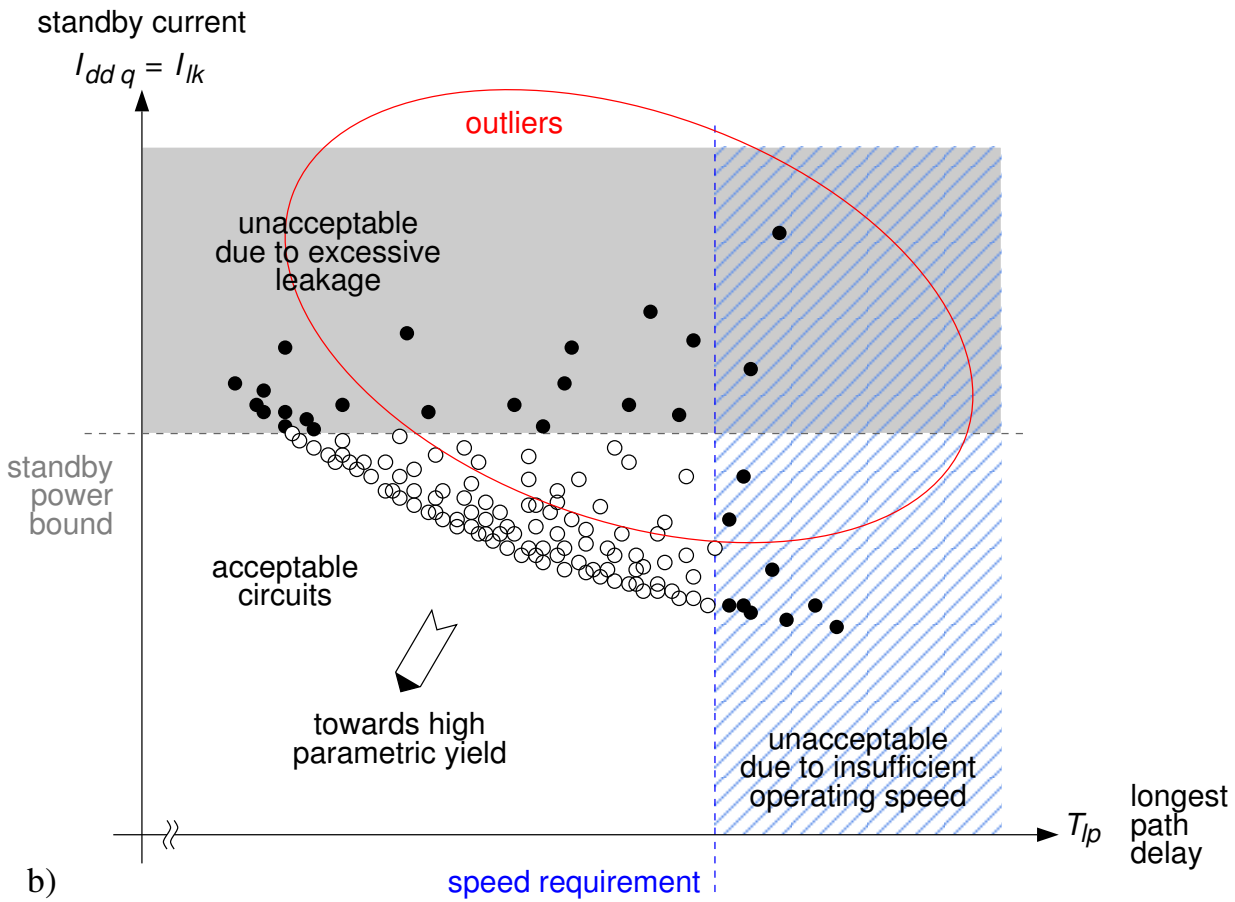

Figure: Yield gets also affected by important parameter variations.

Yield losses due to parametric variability II

- Yield is not only limited by fabrication defects, but also by unpredictable parameter variations.
  - MOSFET threshold voltages, current drives, and leakage currents vary because of local variations of doping concentration and oxide thickness.
  - Metal thickness varies because of non-idealities of chemical mechanical polishing.
- Fabricated parts may have to be rejected because they do not meet specifications in terms of operating speed, I/O timing, standby power, or the like.
- Outliers exhibit higher failure rates than the normal population.
- Manufacturability is a vital issue since the 90 nm generation.

Caveat

Variability is not accounted for in the yield and cost models presented here.

The impact of volume

Adequate strategy for

Large chips produced in huge quantities:

cut die size and package costs by investing more into manual labor and elaborate software tools.

Small to moderate fabrication volumes:

keep non-recurring expenses low,
refrain from over-optimization and
accept non-minimum die size.

The impact of volume

Adequate strategy for

Large chips produced in huge quantities:

cut die size and package costs by investing more into manual labor and elaborate software tools.

Small to moderate fabrication volumes:

keep non-recurring expenses low,
refrain from over-optimization and
accept non-minimum die size.

Typical mask usage

- commodity RAMs and CPUs: 5000 wafers per mask set
- USICs: 500 wafers per mask set

Numerical example

Consider the subsequent full-custom IC design

- die size 4 mm by 4 mm
- scribe lines of 60 μm on all four sides
- wafer diameter 300 mm
- yield 82%
- process monitors hidden in scribe lines

Numerical example

Consider the subsequent full-custom IC design

- die size 4 mm by 4 mm
- scribe lines of 60 μm on all four sides
- wafer diameter 300 mm
- yield 82%
- process monitors hidden in scribe lines

$\rightsquigarrow$ Approximately 3229 functioning dies per wafer.

27% of total wafer surface lost due to edge effect, wafer cutting, and imperfect fabrication combined.

Numerical example (continued)

The design is being fabricated in a

- 90 nm CMOS technology
- 28 mask layers
- 350 USD for raw wafer
- 2300 USD for wafer processing
- extra 1.48 USD for packaging and volume testing

$\rightsquigarrow$ Approximately 2.30 USD for manufacturing one functioning unit.

Numerical example (continued)

Total up-front costs amount to $7 \cdot 10^6$ USD.

- $3.5 \cdot 10^6$ USD for 17 man years of development effort
- $1 \cdot 10^6$ USD for one production mask set
- $1 \cdot 10^6$ USD for the EDA infrastructure
- Rest for virtual components, cell libraries, probe cards, trainings, sign-off, NRE charges, etc.

Numerical example (continued)

Total up-front costs amount to $7 \cdot 10^6$ USD.

- $3.5 \cdot 10^6$ USD for 17 man years of development effort
- $1 \cdot 10^6$ USD for one production mask set
- $1 \cdot 10^6$ USD for the EDA infrastructure
- Rest for virtual components, cell libraries, probe cards, trainings, sign-off, NRE charges, etc.

Overall picture:

	production volume n					
costs per unit [USD]	1k	10k	100k	1M	10M	100M
non-recurring c_0	7 000	700	70	7.00	0.70	0.07
recurring c_1	2.30	2.30	2.30	2.30	2.30	2.30
overall c	7 002	702	72.30	9.30	3.00	2.37
mission-critical	design effort			manufacturing costs		

The impact of configurability

- Applications that mandated a custom ASIC a few years ago fit into a single FPL device today, and this trend is to carry on
- The sales volume necessary to justify custom parts is moving up.

The impact of configurability

- Applications that mandated a custom ASIC a few years ago fit into a single FPL device today, and this trend is to carry on
- The sales volume necessary to justify custom parts is moving up.
- Yet, field-programmable logic is unlikely to rival hardwired logic on the grounds of integration density and recurring costs.
- From a comparison of SRAM-based FPGAs against cell-based ASICs by Kuon & Rose in 90 nm CMOS:
  - Area overhead factor 35
  - Slowdown by a factor between 3.4 and 4.6
  - Dynamic power ratio of 14

  Hardwired multipliers and antifuse technology improve the situation, but a significant penalty remains.

Observation

Configurable logic is best confined to circuit parts subject to frequent changes.

Cost effectiveness

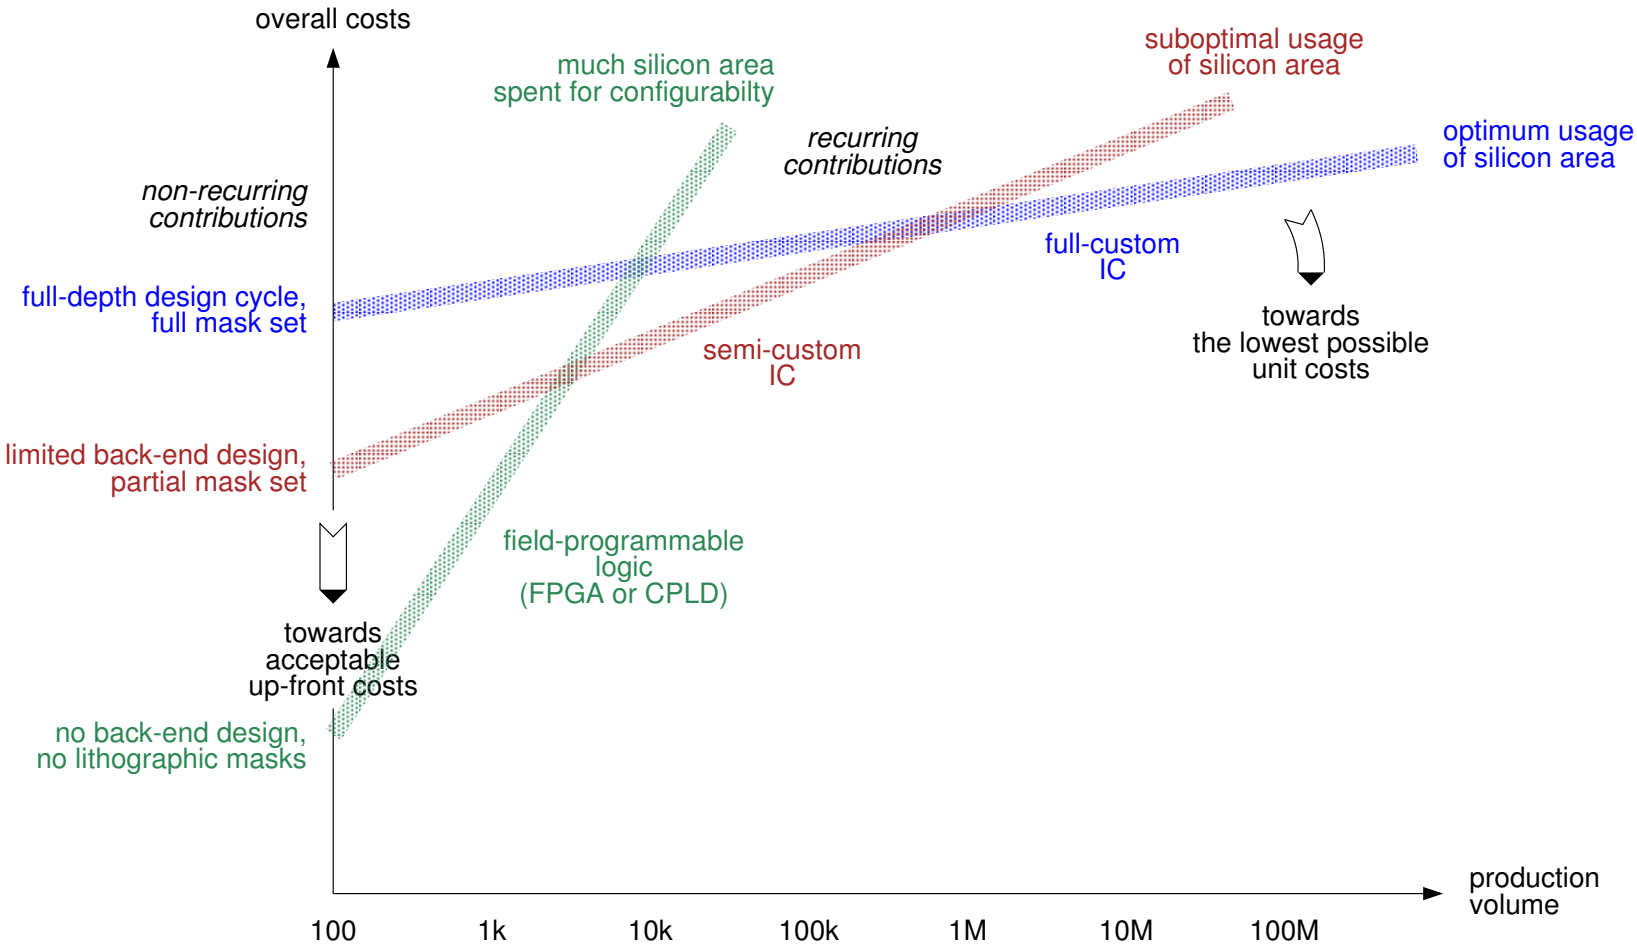

Figure: Overall costs as a function of fabrication depth and volume.

Cost structure of Integrated Circuits

Conclusion

The massive non-recurring and the low recurring costs of full-custom ICs

- favor large volume fabrication but, at the same time, also
- put at a disadvantage products that sell in minor quantities.

- Semi-custom ICs and FPL provide more favorable cost structures for ASICs that are expected to sell in moderate and small quantities.

Fabrication avenues for small quantities I

- Initial costs of IC manufacturing discourage USICs for products with low to moderate sales volumes.
- Many silicon foundries do not accept orders below 1000 wafers per year.
- Field-programmable logic fills the gap, but
  - provides no support for analog circuit blocks,
  - devices come with a very limited selection of packages,
  - is often inferior in terms of energy efficiency.

"What are the other options?"

Sharing mask costs among several designs

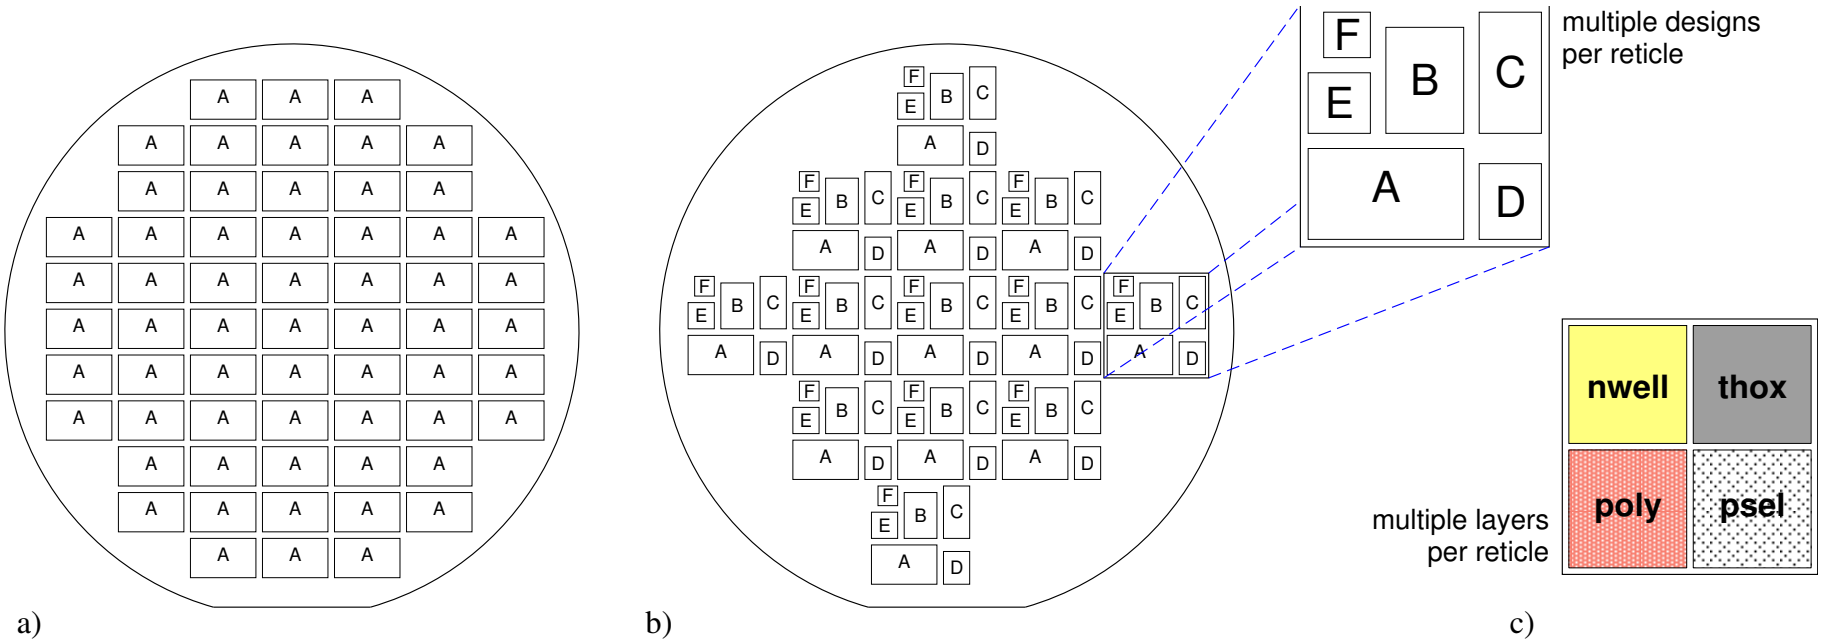

Figure: Single-project (a) vs. multi-project wafer (b). Multi-level mask (c).

Fabrication avenues for small quantities II

Options summary:

- Multi-project wafers
- Multi-level masks
- Electron beam lithography
- Laser programming
- Hardwired FPGAs (pin-compatible semi-custom replacements for FPGAs)
- Structured ASICs
- Combined approaches (such as eASIC)

Fabrication avenues for small quantities II

Options summary:

- Multi-project wafers
- Multi-level masks
- Electron beam lithography
- Laser programming
- Hardwired FPGAs (pin-compatible semi-custom replacements for FPGAs)
- Structured ASICs
- Combined approaches (such as eASIC)

Digest

- The exploding costs of a full mask set works in favor of techniques that achieve customization from a small subset or with no masks at all.
- A handful of companies specializing in the business offer full-custom fabrication with more favorable NRE charges and short turnaround times.