Reliability of complex earth-bound machines

Moore's law drives continuous improvement

System Architecting Study Group meeting at Dutch Space 1-oct-2013

Tom Hoogenboom, ASML

Background & purpose

- Wafer steppers are complex earth-bound machines
  - Must expose > 1000000 silicon wafers/year
  - Each wafer contains 100-1000 individual chips
  - Each chip contains millions of transistors
  - Each exposure must be accurate within a few nm
    else a transistor can fail
    - For reference, grass grows at a rate of 33 nm/s
  - Exposures have to be done at speed (10 fps)
  - Many different subsystems make the stepper complex,
    all have to work together
- Wafer steppers must be reliable:
  - Economy – down time is expensive
    → may miss a complete market cycle(e.g. Christmas sale)
  - Trust – is related to yield (good chips).
    Example: the transistor that triggers your airbag must work ok
- Today we examine some reliablity aspects of
  this class of machines

Summary

- Reliablility = 1 / (failure rate * impact)
- Moore’s law challenges reliablility
- Achieving reliablility is a design challenge
- MBKM write the test specification first

Reliablility = 1 / (failure rate * impact)

- Un-reliable means $\Sigma(\frac{1}{\text{failure rate} * \text{impact}})$ low

- Requirements for ground-based and space-born applications are different:

Design goal	Ground-based	Space-born
Failure rate	Manageable	low
Impact	Manageable (but safety OK)	low
Time/effort	1 year	10 years

Examples of ‘manageable’

- Manageable failure rate:
  - For a new machine, not yet in full production, customer may accept a restart every 48 hours
    - If this buys time to perform printing performance evaluation for the next generation chips
  - Within (say) 1 year the machine is upgraded in the field to prepare for full production

- Manageable impact:
  - When a machine happens to go down, it is known which chips were printed correctly and which were not.

This is the machine-centric view

‘This printer is 99.9% reliable’

But what about the chips?

The obvious rule for chips

like the ones in your airbag controller

is ‘0 tolerance for error’

So when the printer works,
it must work 100% repeatable

Summary

- Reliablility = 1 / (failure rate * impact)
  so we must work to reduce both
- Moore’s law challenges reliablility
- Achieving reliablility is a design challenge
- MBKM write the test specification first

Moore is more failures? No*!*

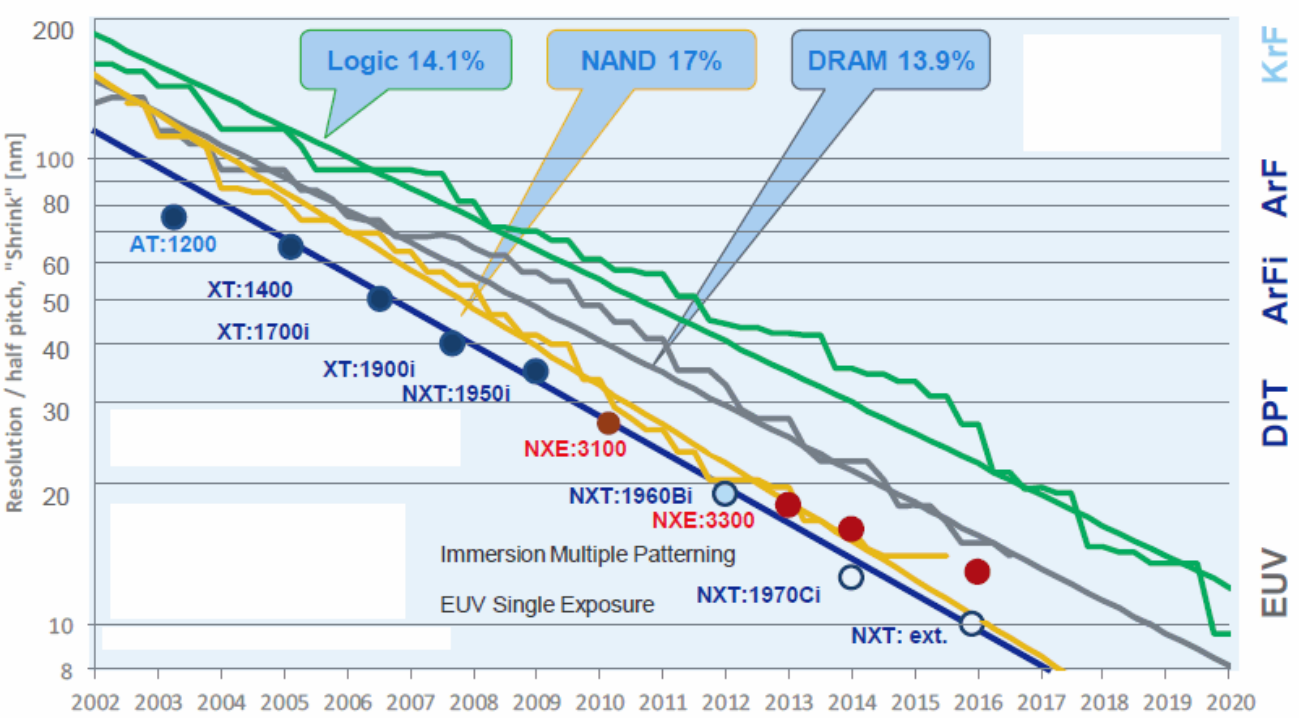

- 2x shrink → 4x transistor count → reliab/transistor improves
  - → Uniformity drive – all transistors made equal
  - → Redundancy at all system levels (printer *and* chip)

Supplier and end-user win-win: Uniformity helps printer *and* chip

Typical integrated circuit

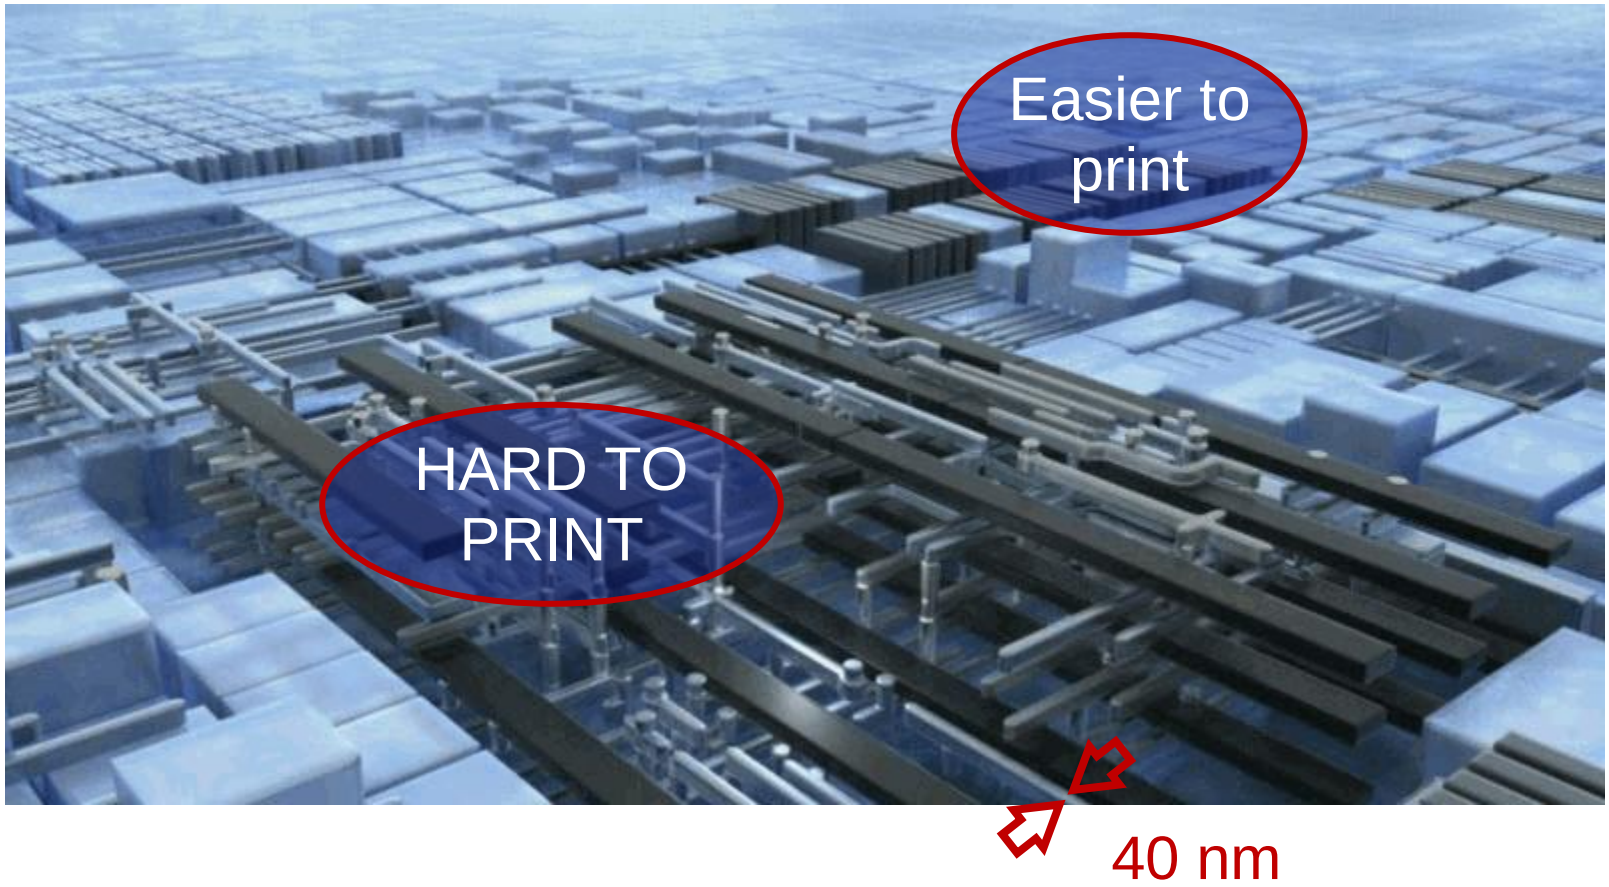

Redundancy helps – chip example

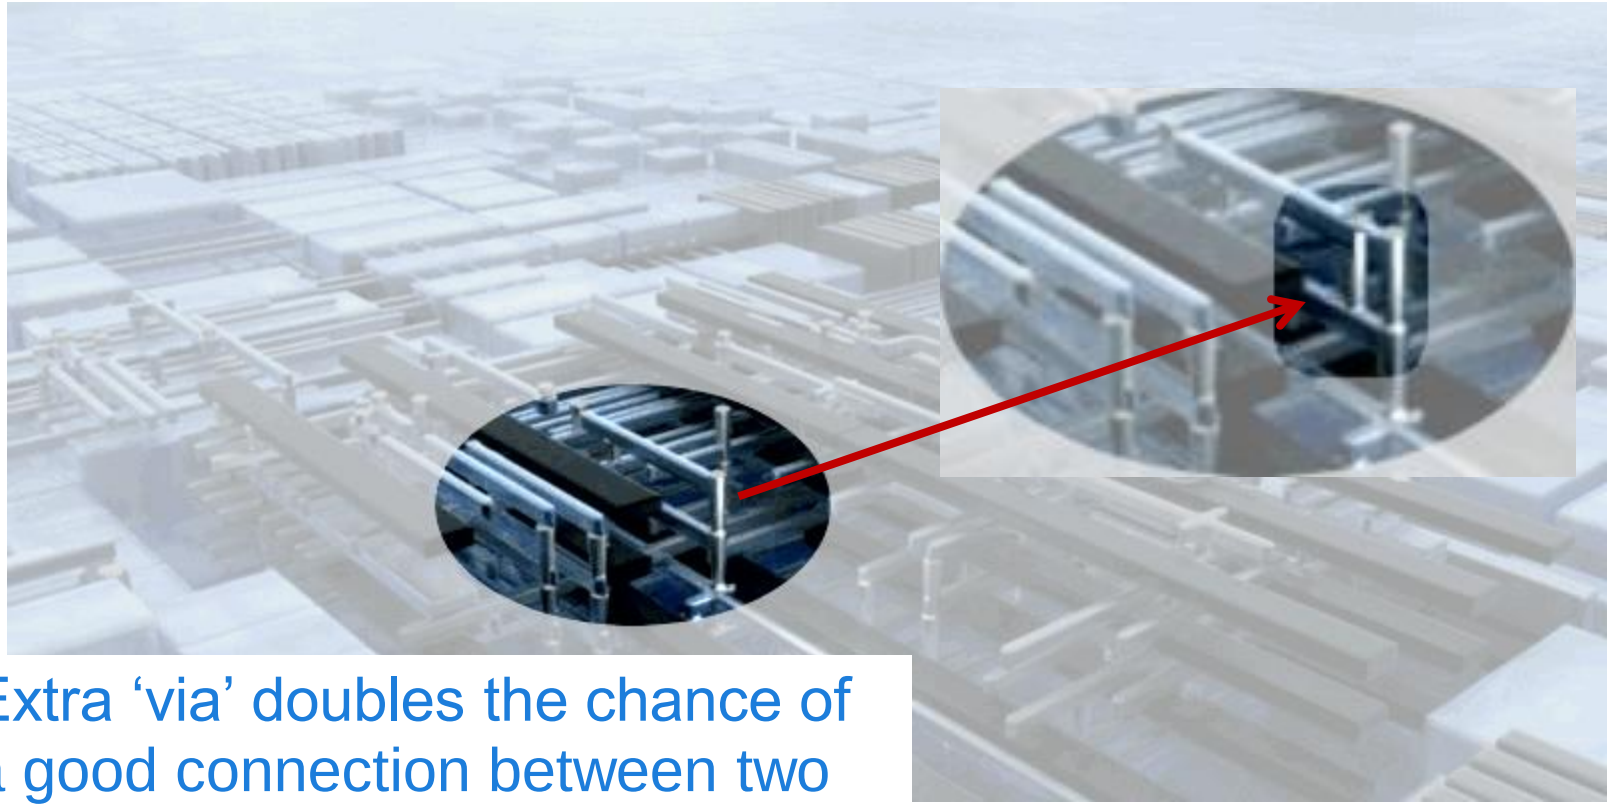

Extra ‘via’ doubles the chance of a good connection between two metal layers

But this cannot be all – we also need a reliable printer

Summary

- Reliablility = 1 / (failure rate * impact)
  so we must work to reduce both
- Moore's law challenges reliablility
  Suppliers and end-users must work together
- Achieving reliablility is a design challenge

- MBKM write the test specification first

No ‘proven’ design for chip-making machines

- Lithography is driven by shrink
  - continuous nm improvements needed
- Impact of failure is bigger, e.g.:
  - Same defect on mask affects more transistors
  - Machines are more expensive – down is more costly
  - Finished wafers more expensive – rework/scrap is more costly
- Therefore we technicians (supplier and end-user) fight the ‘rate’ part
  - ‘simple’ issues become a nuisance and need to be understood, then removed
    - Even as machine size and complexity increases in view of new machine demands such as
      - water under the lens
      - operation in vacuum
- So solutions from the past are not guaranteed to work in the future

$$\Sigma\left(\frac{1}{\text{failure rate} * \text{impact}}\right)$$

Time proven redundancy can help again – chip/system level

- Examples
  - Memory
    - Parity bit – detect errors
    - Error correction – can repair 1-bit error
  - Disks: RAID configuration
  - CPU: multiple blades, virtual machines
  - Power: dual supplies

- Often cost effective when ‘machine down’ is expensive
  - Requires some form of ‘continuous caretaking’
    - e.g. two switching power supplies ‘made on Monday’ find their way into the same computer cabinet
    - High risk that 2nd supply fails when it has to take over full load

New issues require new cleverness

- Example: wafer clamping
  - Existing design: ‘vacuum’ under wafer clamps it to chuck
  - New EUV light (13.5 nm) requires operation in ‘vacuum’
  - Requires new solution for clamping
  - Requires extra time for chuck maintenance (bring to air, service, pump down)
  - Creates new sources of variability we need to understand
    - E.g wafer flatness changes when clamping is different

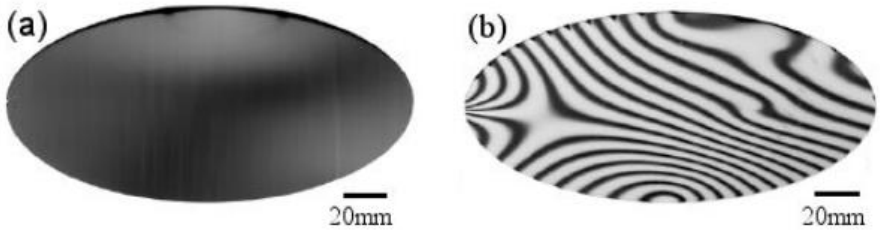

Summary

- Reliablility = 1 / (failure rate * impact)
  so we must work to reduce both
- Moore's law challenges reliablility
  Suppliers and end-users must work together
- Achieving reliablility is a design challenge
  There are no 'proven' designs in lithography
- MBKM write the test specification first

When every reliab step requires innovation, traditional test methods may no longer work

- New, clever, solutions trigger new, unknown reliablility issues.
- Product testing can be too late to find these
  - Manufacturing tests finds issues when machines are almost ready for shipping
  - System tests finds issues when design is already finished and is difficult to change

- So MBKM (My Best Known Method) is to start testing as early as possible

My best known method: create the test specification during design

- Reliability cannot be ‘tested into’ a system
  - Considering testing as integrated part of the design can help identify the unknown reliability issues before the design is finalized
- So I promote to create the test specification as early as possible
  - I ‘sprint’ with the design teams
  - I try to sniff out unknowns early on, look at review results, work through issue lists, talk to developers and feature testers
  - I check out the tests done by the teams
- Come alpha test time I have a good idea of reliability
  - Proof of the pudding is all there is left....
  - This is the basis for my system alpha test specification
  - System testing includes simulation of typical worst case customer use

Final Summary

- Reliablility = 1 / (failure rate * impact)
  so we must work to reduce both
- Moore's law challenges reliablility
  Suppliers and end-users must work together
- Achieving reliablility is a design challenge
  There are no 'proven' designs in lithography
- MBKM write the test specification first.
  My best known method for system engineering:
  run with the developers until it is time to turn alpha.
  Then test on the basis of what you learned from day 1 onwards.