

# Robust Computing Systems

*From Today to the N3XT 1,000×*

Subhasish Mitra





Department of EE & Department of CS  
Stanford University

# World Relies on Computing

Swire, ndirogn es, seotsoes, phrowes, ing data data


Genomics



Security



Smart Cities



Finance



Military




Health Care



Government




Science





# World Relies on Computing

## Genomics



## Smart Cities



## Military




## Health Care



## Government




## Science



## Security



## Finance




- Robust operation
- Performance
- New application horizon

# Research Topics


- Robust operation
  - Bugs, reliability, security
- Revolutionize NanoSystems
  - *1,000*× opportunity
- Program human brain
  - Stanford Big Ideas in Neuroscience

# Outline

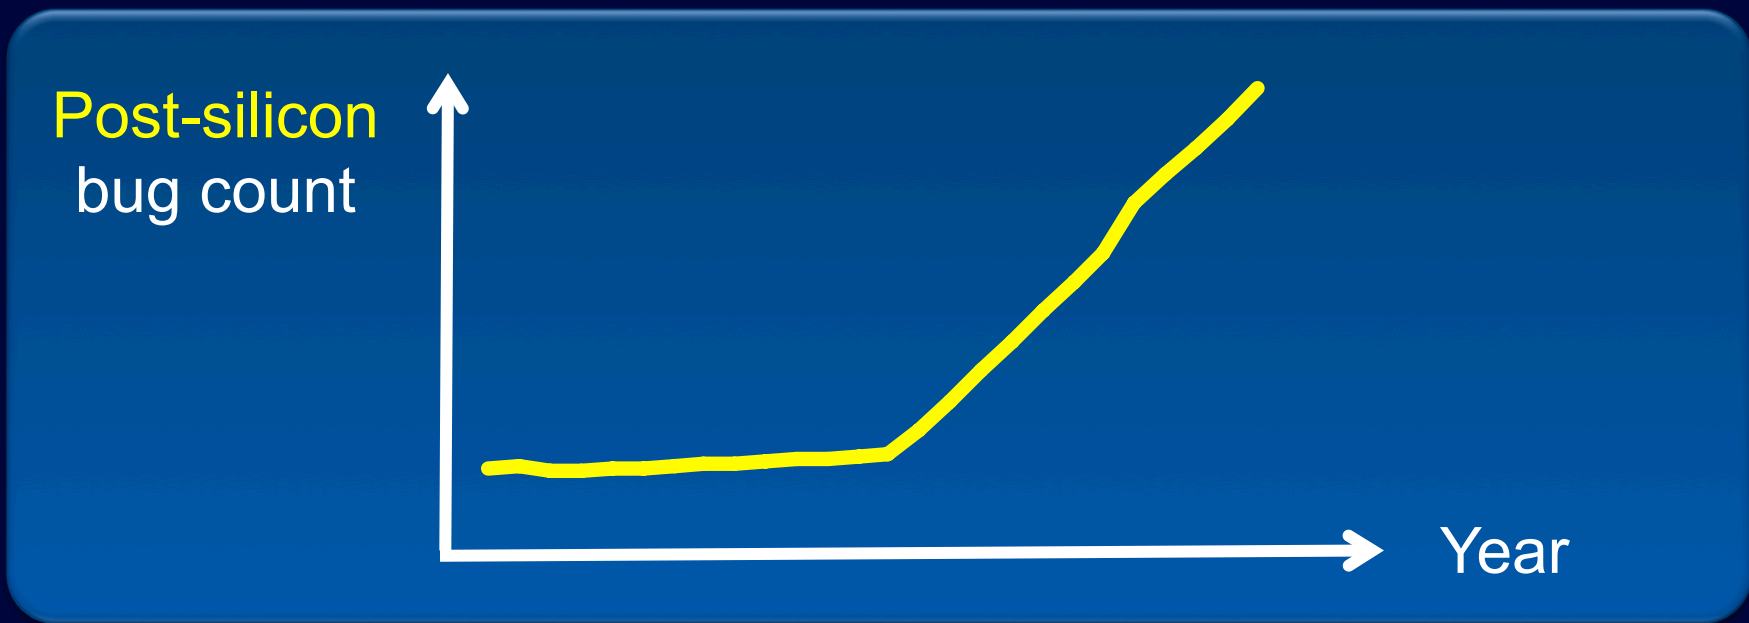
- Design bugs: QED & Symbolic QED



# Pre-Silicon Verification Inadequate




**It's only getting worse: custom hardware**





# Post-Silicon Validation Difficult

System tests




Detect



Weeks or months of manual work



Root-cause & fix



Localize

# Scalability Barriers

- System-level failure reproduction
- Full system simulation



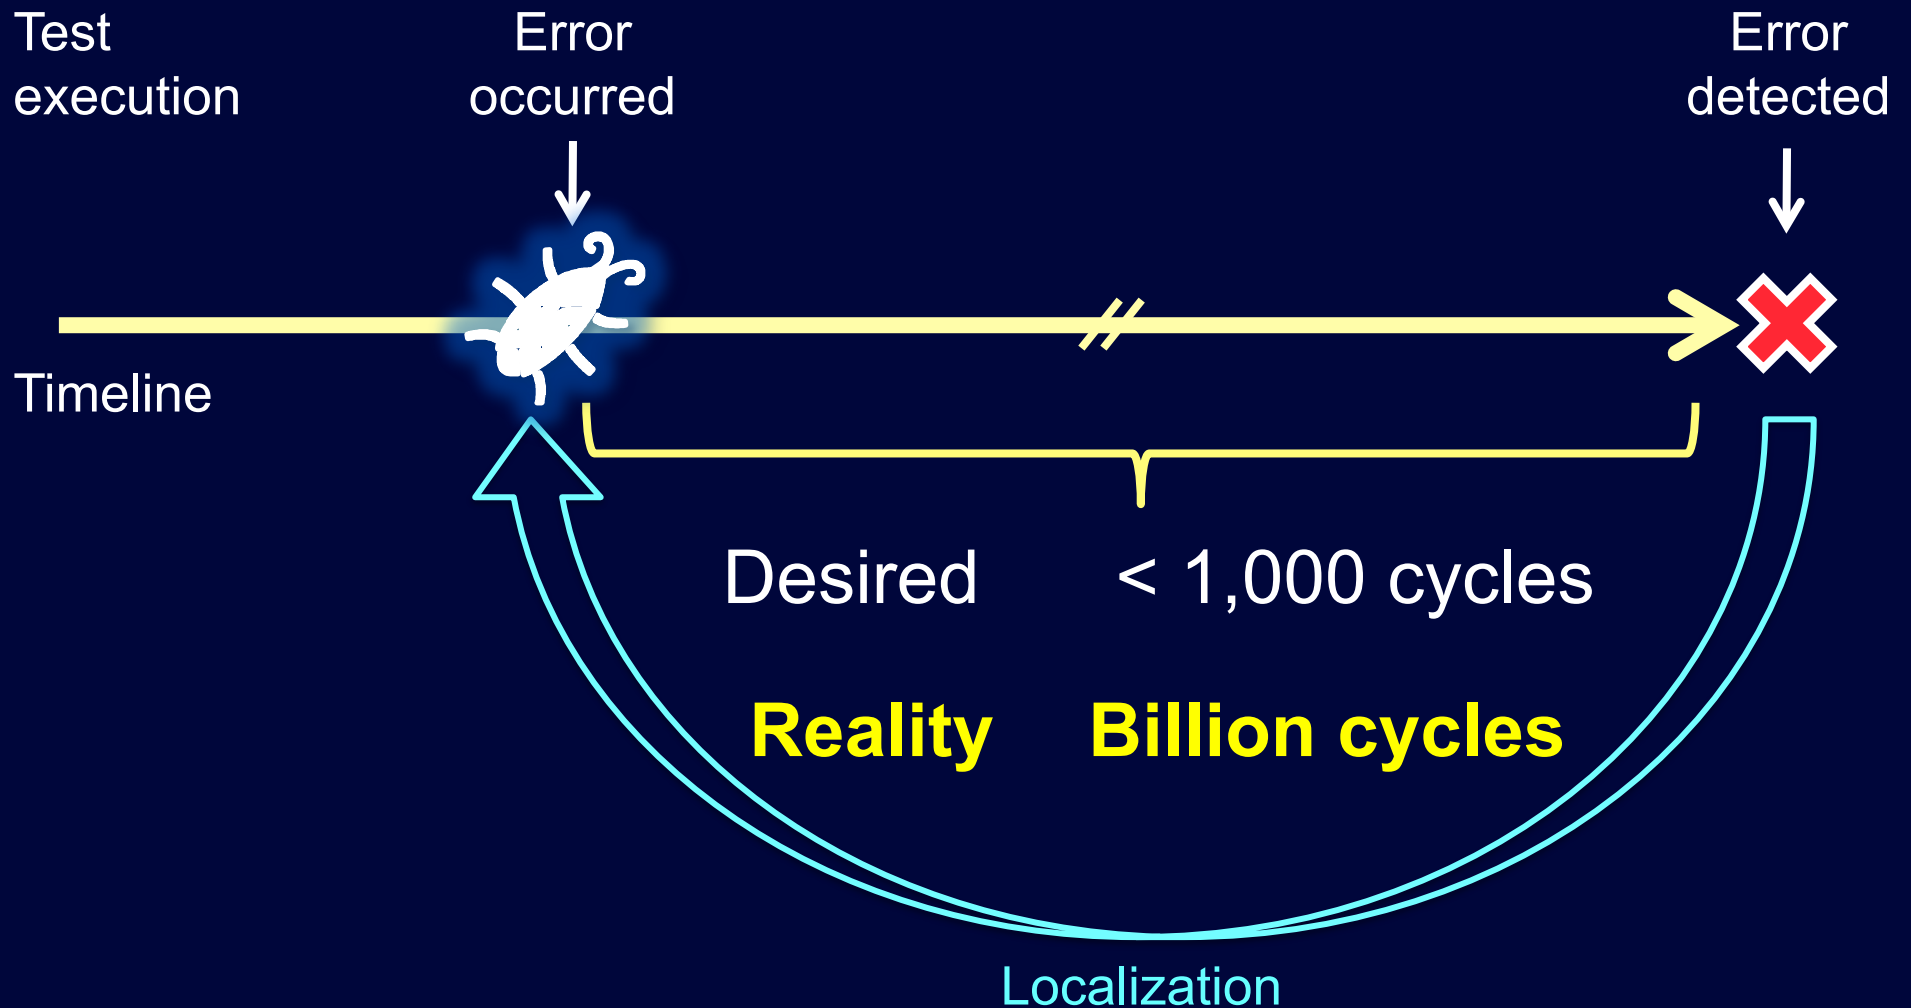
J. Stinson (ex-Intel)

**Post-silicon costs rising faster than design cost**

# QED

- Post-silicon
  - Electrical bugs, logic bugs

# Error Detection Latency





# Quick Error Detection



- Error detection latency: **guaranteed short**
- Coverage: **improved**
- Software-only: **readily applicable**

# Quick Error Detection

- Snippet 1
- Snippet 2
- Snippet 3
- Snippet 4
- Snippet 5
- Snippet 6
- Snippet 7

...

**QED**  
**transforms**

1. Wide variety  
2. Diversity

↑  
Error detection latency target  
Intrusiveness constraints

- Snippet 1
- QED: CFTSS-V 1
- QED: EDDI-V with diversity 1
- QED: PLC 1
- Snippet 2
- QED: CFTSS-V 2
- QED: EDDI-V with diversity 2
- QED: PLC 2
- Snippet 3

...

# Quick Error Detection

- Snippet 1
- Snippet 2
- Snippet 3
- Snippet 4
- Snippet 5
- Snippet 6
- Snippet 7



**QED**  
**transforms**

1. Wide variety  
2. Diversity


↑  
Error detection latency target  
Intrusiveness constraints

- Snippet 1
- QED: CFTSS-V 1
- QED: EDDI-V with diversity 1
- QED: PLC 1
- Snippet 2
- QED: CFTSS-V 2
- QED: EDDI-V with diversity 2
- QED: PLC 2
- Snippet 3



# QED Transforms

## System on Chip



■ Processor cores

■ ■ ■ Uncore, accelerators

**EDDI-V**

PLC

CFTSS-V

Fast QED

CFCSS-V

Hybrid QED



# QED Example: Duplicate & Check

## Validation program

```

void matgen (REAL a[], int lda, int n, REAL b[], REAL *norma)
{
    int i, j;
    int m = 325;
    *norma = 0.0;
    for (i = 0; i < n; i++) {
        for (j = 0; j < n; j++) {
            init = 3125*init % 0x100000000;
            a[lda*i+j] = (init - 32768.0)/65536.0;
            *norma = (a[lda*i+j] > *norma) ? a[lda*i+j] : *norma;
        }
    }
    for (i = 0; i < n; i++) {
        b[i] = 0.0;
    }
    for (i = 0; i < n; i++) {
        for (j = 0; j < n; j++) {
            b[i] = b[i] + a[lda*i+j];
        }
    }
    return;
}

void dgefa(REAL a[], int lda, int n, int ipvt[], int *info)
/*
  internal variables
  REAL t;
  int j,k,kpl,mm1;
  /* gaussian elimination with partial pivoting
  *info = 0;
  mm1 = n - 1;
  if (mm1 <= 0) {
      for (k = 0; k < mm1; k++) {
          kpl = k + 1;
          /* find l = pivot index */
          l = idamax(n-k, &a[lda*k+k], 1) + k;
          ipvt[k] = l;
          /* zero pivot implies this column already
          triangularized */
          if (a[lda*k+k] != ZERO) {
              /* interchange if necessary */
              if (l != k) {
                  t = a[lda*k+k];
                  a[lda*k+k] = a[lda*k+l];
                  a[lda*k+l] = t;
              }
              /* compute multipliers */
              t = -ONE/a[lda*k+k];
              dscal(n-(k+1), t, &a[lda*k+k+1], 1);
          }
          /* row elimination with column indexing */
          for (j = kpl; j < n; j++) {
              t = a[lda*j+k];
              if (l != k) {
                  a[lda*j+k] = a[lda*j+k];
                  a[lda*j+k] = t;
              }
              daxpy(n-(k+1), t, &a[lda*k+k+1], 1,
                  &a[lda*j+k+1], 1);
          }
          else {
              *info = k;
          }
      }
      ipvt[mm1] = mm1;
      if (a[lda*mm1+mm1] == ZERO) *info = mm1;
      checkstop(&a, ipvt, info);
      return;
  }
}

void dgesl(REAL a[], int lda, int n, int ipvt[], REAL b[], int job)
{
    REAL t;
    int k, kb, l, mm1;
    mm1 = n - 1;
    if (job == 0) {
        /* job 0, solve a * x = b
        first solve l'y = b
        if (mm1 >= 1) {
            for (k = 0; k < mm1; k++) {
                l = ipvt[k];
                t = b[l];
                if (l != k) {
                    kb[l] = b[k];
                    b[k] = t;
                }
                daxpy(n-(k+1), t, &a[lda*k+k+1], 1, &b[k+1], 1);
            }
        }
        /* now solve u*u = y */
        for (kb = 0; kb < n; kb++) {
            k = n - (kb + 1);
            b[k] = b[k]/a[lda*k+k];
            t = -b[k];
            daxpy(k, t, &a[lda*k+k], 1, &b[0], 1);
        }
    }
    else {

```

## Trace

...

R1 ← R1 + 5

R2 ← R2 - R1

R3 ← R1 \* R2



Very Long!



Check end\_result

# QED Example: Duplicate & Check

## Validation program

```

void matgen (REAL a[], int lda, int n, REAL b[], REAL *norma)
{
    int i, j;
    int m = 32768;
    *norma = 0.0;
    for (i = 0; i < n; i++) {
        for (j = 0; j < n; j++) {
            init = 3125*init % 0x100000000;
            a[lda*j+i] = (init - 32768.0)/65536.0;
            *norma = (a[lda*j+i] > *norma) ? a[lda*j+i] : *norma;
        }
    }
    for (i = 0; i < n; i++) {
        b[i] = 0.0;
    }
    for (i = 0; i < n; i++) {
        for (j = 0; j < n; j++) {
            b[i] = b[i] + a[lda*j+i];
        }
    }
    return;
}

void dgefa(REAL a[], int lda, int n, int ipvt[], int *info)
/*
  internal variables
*/
REAL t;
int j, k, kpl, l, mml;

/*
  gaussian elimination with partial pivoting
*/
*info = 0;
mml = n - 1;
if (mml <= 0) {
    for (k = 0; k < mml; k++) {
        kpl = k + 1;
        /* find l = pivot index */
        l = idamax(n-k, &a[lda*k+k], 1) + k;
        ipvt[k] = l;
        /* zero pivot implies this column already
        triangularized */
        if (a[lda*k+k] != ZERO) {
            /* interchange if necessary */
            if (l != k) {
                t = a[lda*k+k];
                a[lda*k+k] = a[lda*k+l];
                a[lda*k+l] = t;
            }
            /* compute multipliers */
            t = -ONE/a[lda*k+k];
            dscal(n-(k+1), t, &a[lda*k+k+1], 1);
            /* row elimination with column indexing */
            for (j = kpl; j <= n; j++) {
                t = a[lda*j+k];
                if (l != k) {
                    a[lda*j+k] = a[lda*j+k];
                    a[lda*j+k] = t;
                }
                daxpy(n-(k+1), t, &a[lda*k+k+1], 1,
                &a[lda*j+k+1], 1);
            }
        }
        else {
            *info = k;
        }
    }
}
ipvt[mml] = mml;
if (a[lda*(m-1)*(m-1)] == ZERO) *info = m-1;
checkbsexpected(a, ipvt, info);
return;
}

void dgesl(REAL a[], int lda, int n, int ipvt[], REAL b[], int job)
{
    REAL t;
    int k, kb, l, mml;
    mml = n - 1;
    if (job == 0) {
        /* job == 0, solve a * x = b
        first solve l'y = b
        */
        if (mml == -1) {
            for (k = 0; k < mml; k++) {
                l = ipvt[k];
                t = b[l];
                if (l != k) {
                    b[l] = b[k];
                    b[k] = t;
                }
                daxpy(n-(k+1), t, &a[lda*k+k+1], 1, &b[k+1], 1);
            }
        }
        /* now solve u*x = y */
        for (kb = 0; kb < n; kb++) {
            k = n - (kb + 1);
            b[k] = b[k]/a[lda*k+k];
            t = -b[k];
            daxpy(k, t, &a[lda*k+k], 1, &b[0], 1);
        }
    }
    else {

```

## QED Trace

R16 ← R1; R18 ← R2; R19 ← R3

• • •

R1 ← R1 + 5

R16 ← R16 + 5

R1 == R16

R2 ← R2 - R1

R18 ← R18 - R16

R2 == R18

R3 ← R1 \* R2

R19 ← R16 \* R18

R3 == R19

• • •

Check end\_result



Quick!

# QED Improves Coverage

## Validation program

```

void matgen (REAL a[], int lda, int n, REAL b[], REAL *norma)
{
    int i, j;
    int i1, j1;
    *norma = 0.0;
    for (i = 0; i < n; i++) {
        for (j = 0; j < n; j++) {
            init = 3125*init % 0x100000000;
            a[lda*i+j] = (init - 32768.0)/65536.0;
            *norma = (a[lda*i+j] > *norma) ? a[lda*i+j] : *norma;
        }
    }
    for (i = 0; i < n; i++) {
        b[i] = 0.0;
    }
    for (i = 0; i < n; i++) {
        for (j = 0; j < n; j++) {
            b[i] = b[i] + a[lda*j+i];
        }
    }
    return;
}

void dgefa(REAL a[], int lda, int n, int ipvt[], int *info)
/*
   internal variables
*/
REAL t;
int j,k,kl,lm1;
/*
   gaussian elimination with partial pivoting
*/
*info = 0;
lm1 = n - 1;
if (lm1 <= 0) {
    for (k = 0; k < n; k++) {
        kl = k + 1;
        /* find l = pivot index */
        l = idamax(n-k, &a[lda*k+k], 1) + k;
        ipvt[k] = l;
        /* zero pivot implies this column already
           triangularized */
        if (a[lda*k+k] != ZERO) {
            /* interchange if necessary */
            if (l != k) {
                t = a[lda*k+k];
                a[lda*k+k] = a[lda*k+l];
                a[lda*k+l] = t;
            }
            /* compute multipliers */
            t = -ONE/a[lda*k+k];
            dscal(n-k+1, t, &a[lda*k+k+1], 1);
            /* row elimination with column indexing */
            for (j = kl; j < n; j++) {
                t = a[lda*j+k];
                if (l != k) {
                    a[lda*j+k] = a[lda*j+l];
                    a[lda*j+l] = t;
                }
                daxpy(n-k+1, t, &a[lda*k+k+1], 1,
                    &a[lda*j+k+1], 1);
            }
        }
        else {
            *info = k;
        }
    }
}
ipvt[lm1] = n-1;
if (a[lda*(n-1)*(n-1)] == ZERO) *info = n-1;
checkio(gefa, ipvt, info);
return;
}

void dgesl(REAL a[], int lda, int n, int ipvt[], REAL b[], int job)
{
    REAL t;
    int k, kb, lm1;
    lm1 = n - 1;
    if (job == 0) {
        /* job 0: solve a * x = b
           first solve l * y = b
           if (lm1 == -1) {
               for (k = 0; k < n; k++) {
                   l = ipvt[k];
                   t = b[l];
                   if (l != k) {
                       b[k] = b[l];
                       b[l] = t;
                   }
                   daxpy(n-k+1, t, &a[lda*k+k+1], 1, &b[k+1], 1);
               }
           }
           /* now solve u * v = y */
           for (kb = 0; kb < n; kb++) {
               k = n - (kb + 1);
               b[k] = b[k]/a[lda*k+k];
               t = -b[k];
               daxpy(k, t, &a[lda*k+0], 1, &b[0], 1);
           }
        }
    }
    else {

```

## Trace


• • •  
 $R1 \leftarrow R1 + 1$



• • •  
 $R1 \leftarrow 0$

• • •  
 Check  $R1 == 0$  **Error masked!**

R1



# QED Improves Coverage

## Validation program

```

void matgen (REAL a[], int lda, int n, REAL b[], REAL *norma)
{
  int i, j;
  int i1, j1;
  *norma = 0.0;
  for (i = 0; i < n; i++) {
    for (j = 0; j < n; j++) {
      init = 3125*init % 0x100000000;
      a[lda*i+j] = (init - 32768.0)/65536.0;
      *norma = (a[lda*i+j] > *norma) ? a[lda*i+j] : *norma;
    }
  }
  for (i = 0; i < n; i++) {
    b[i] = 0.0;
  }
  for (i = 0; i < n; i++) {
    for (j = 0; j < n; j++) {
      b[i] = b[i] + a[lda*i+j];
    }
  }
  return;
}

void dgefa(REAL a[], int lda, int n, int ipvt[], int *info)
/*
  internal variables
  REAL t;
  int j,k,kpl,mm1;
  /* gaussian elimination with partial pivoting
  *info = 0;
  mm1 = n - 1;
  if (mm1 <= 0) {
    for (k = 0; k < mm1; k++) {
      kpl = k + 1;
      /* find l = pivot index */
      l = idamax(n-k, &a[lda*k+k], 1) + k;
      ipvt[k] = l;
      /* zero pivot implies this column already
      triangularized */
      if (a[lda*k+k] != ZERO) {
        /* interchange if necessary */
        if (l != k) {
          t = a[lda*k+k];
          a[lda*k+k] = a[lda*k+l];
          a[lda*k+l] = t;
        }
        /* compute multipliers */
        t = -ONE/a[lda*k+k];
        dscal(n-(k+1), t, &a[lda*k+k+1], 1);
        /* row elimination with column indexing */
        for (j = kpl; j <= n; j++) {
          t = a[lda*j+k];
          if (l != k) {
            a[lda*j+k] = a[lda*j+k];
            a[lda*j+k] = t;
          }
          daxpy(n-(k+1), t, &a[lda*k+k+1], 1,
            &a[lda*j+k+1]);
        }
      }
      else {
        *info = k;
      }
    }
  }
  ipvt[mm1] = mm1;
  if (a[lda*(mm1-1)] == ZERO) *info = mm1-1;
  checkiopefa(a, ipvt, info);
  return;
}

void dgesl(REAL a[], int lda, int n, int ipvt[], REAL b[], int job )
{
  REAL t;
  int k, kb, l, mm1;
  mm1 = n - 1;
  if (job == 0) {
    /* solve a * x = b
    first solve l * y = b
    */
    if (mm1 >= 1) {
      for (k = 0; k < mm1; k++) {
        l = ipvt[k];
        t = b[l];
        if (l != k) {
          b[l] = b[k];
          b[k] = t;
        }
        daxpy(n-(k+1), t, &a[lda*k+k+1], 1, &b[k+1]);
      }
    }
    /* now solve a * x = y */
    for (kb = 0; kb < n; kb++) {
      k = n - (kb + 1);
      b[k] = b[k]/a[lda*k+k];
      t = -b[k];
      daxpy(k, t, &a[lda*k+0], 1, &b[0]);
    }
  }
  else {

```

## QED Trace

• • •

R1 ← R1 + 1

R18 ← R18 + 1

R1 == R18



• • •

R1 ← 0

• • •

Check R1 == 0 masked!

R1 R18

3	2	2
---	---	---

No QED: bug escape

QED: quick detection



# Diversity-Enhanced QED

## Validation program

```

void matgen (REAL a[], int lda, int n, REAL b[], REAL *norma)
{
    int i, j;
    int m = 32768;
    *norma = 0.0;
    for (i = 0; i < n; i++) {
        for (j = 0; j < n; j++) {
            init = 3125*init % 0x100000000;
            a[lda*i+j] = (init - 32768.0)/65536.0;
            *norma = (a[lda*i+j] > *norma) ? a[lda*i+j] : *norma;
        }
    }
    for (i = 0; i < n; i++) {
        b[i] = 0.0;
    }
    for (i = 0; i < n; i++) {
        for (j = 0; j < n; j++) {
            b[i] = b[i] + a[lda*i+j];
        }
    }
    return;
}

void dgefa(REAL a[], int lda, int n, int ipvt[], int *info)
/*
   internal variables
*/
REAL t;
int j,k,kl,lm1;
/*
   gaussian elimination with partial pivoting
*/
*info = 0;
m1 = n - 1;
if (m1 <= 0) {
    for (k = 0; k <= m1; k++) {
        kl = k + 1;
        /* find l = pivot index */
        l = idamax(n-k, &a[lda*k+k], 1) + k;
        ipvt[k] = l;
        /* zero pivot implies this column already
           triangularized */
        if (a[lda*k+k] != ZERO) {
            /* interchange if necessary */
            if (l != k) {
                t = a[lda*k+k];
                a[lda*k+k] = a[lda*k+l];
                a[lda*k+l] = t;
            }
            /* compute multipliers */
            t = ONE/a[lda*k+k];
            dscal(n-(k+1), t, &a[lda*k+k+1], 1);
            /* row elimination with column indexing */
            for (j = kl; j <= n; j++) {
                t = a[lda*j+k];
                if (l != k) {
                    a[lda*j+k] = a[lda*j+l];
                    a[lda*j+l] = t;
                }
                saxpy(n-(k+1), t, &a[lda*k+k+1], 1,
                    &a[lda*j+k+1], 1);
            }
        }
        else {
            *info = k;
        }
    }
}
ipvt[m1] = m1;
if (a[lda*(m1-1)] == ZERO) *info = m1;
checkstoped(a, ipvt, info);
return;
}

void dgesl(REAL a[], int lda, int n, int ipvt[], REAL b[], int job)
{
    REAL t;
    int k, kb, l, m1;
    m1 = n - 1;
    if (job == 0) {
        /* job == 0, solve a * x = b
           first solve l*y = b
        */
        if (m1 == -1) {
            for (k = 0; k <= m1; k++) {
                l = ipvt[k];
                t = b[l];
                if (l != k) {
                    b[k] = b[l];
                    b[l] = t;
                }
                daxpy(n-(k+1), t, &a[lda*k+k+1], 1, &b[k+1], 1);
            }
        }
        /* now solve u*u = y */
        for (kb = 0; kb < n; kb++) {
            k = n - (kb + 1);
            b[k] = b[k]/a[lda*k+k];
            t = -b[k];
            daxpy(k, t, &a[lda*k+0], 1, &b[0], 1);
        }
    }
    else {

```

## QED Trace

$$a \leftarrow 1; a' \leftarrow a; b \leftarrow 1, b' \leftarrow b$$

$$R1 \leftarrow a + b$$

$$R18 \leftarrow 5a' + 5b'$$

$$5 \times R1 == R18 \quad \times$$



$\begin{bmatrix} 6 & 2 \\ 14 & 10 \end{bmatrix}$	R1
$\begin{bmatrix} 6 & 2 \\ 14 & 10 \end{bmatrix}$	R18



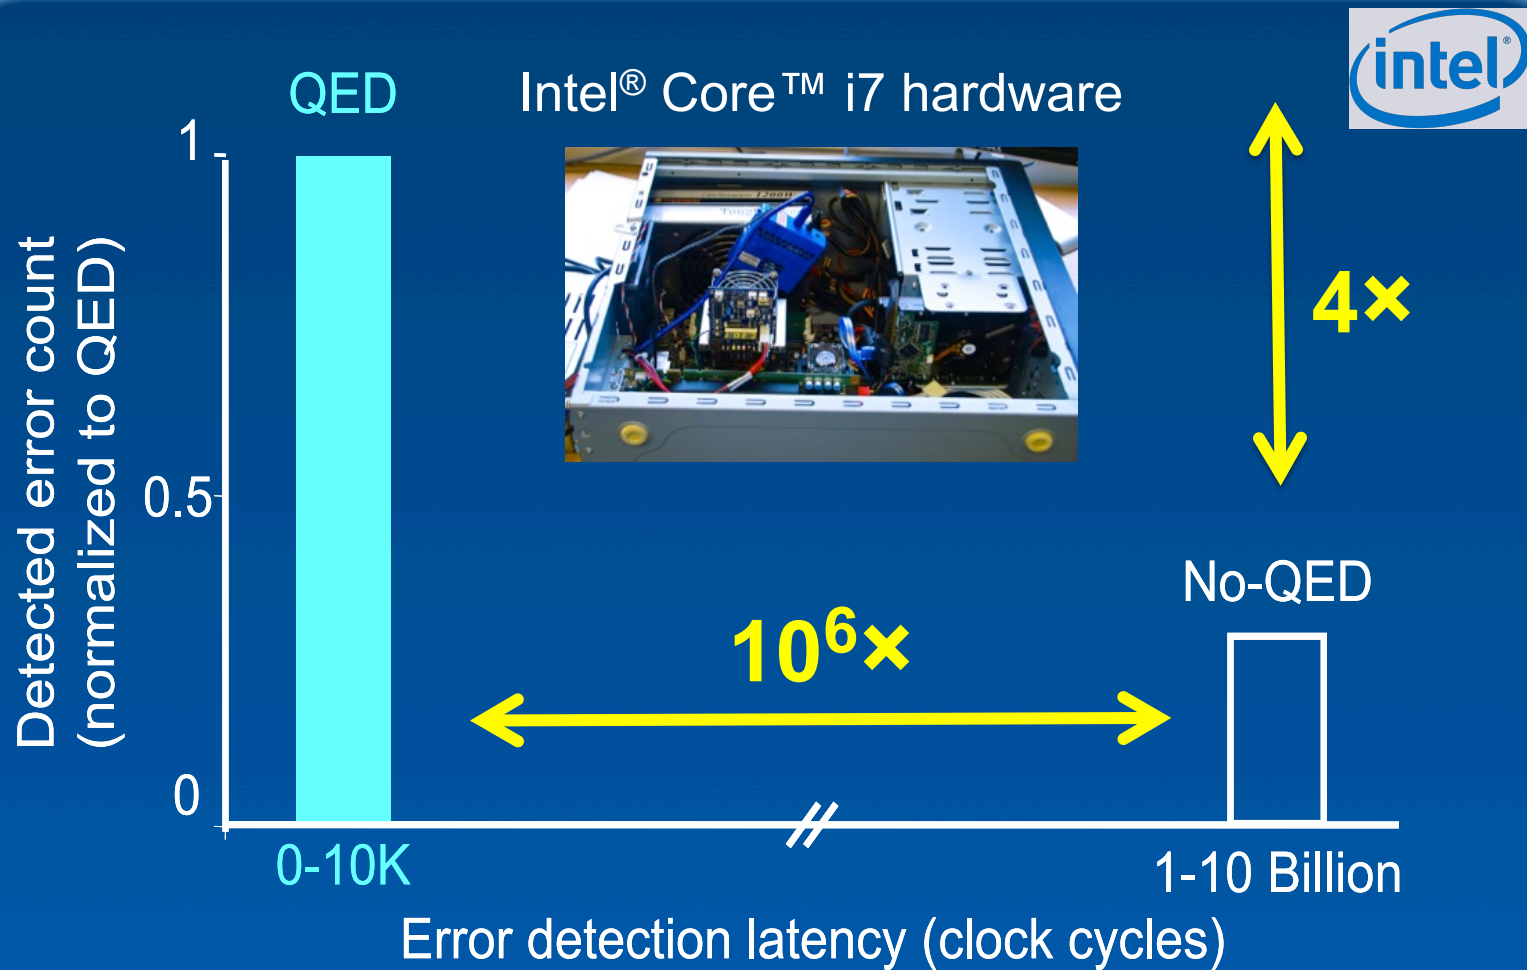
Many diversity techniques

e.g., ED<sup>4</sup>I [Oh IEEE Trans. Comp. 02]

# QED Coverage Considerations

- Challenge
  - Intrusiveness: coverage impact ?
- Systematic solutions
  - QED family tests
  - Hardware-enhanced QED

# QED Effective for Electrical Bugs



# E-QED

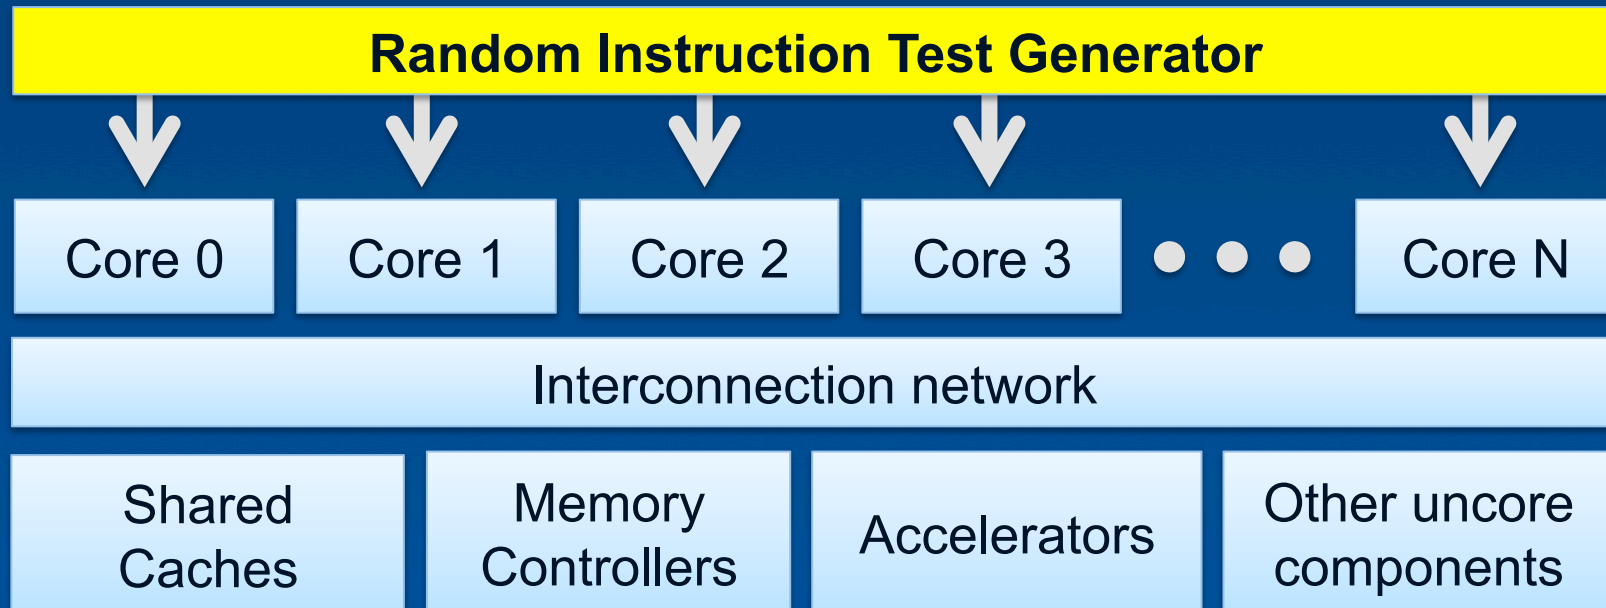
- **Electrical bug localization**

	<b>E-QED</b>
Flip-flop candidates	18 (1 million total) <b>50,000× localization</b>
Area impact	2.5% (actually 0%)
Debug effort	Automatic
Runtime	~ 9 hours

OpenSPARC T2 SoC (500M transistors)


# QED Effective for Logic Bugs

Freescale SoC hardware



Error detection latency (cycles)

No QED	QED
<b>15 Billion</b>	<b>9</b>





# Symbolic QED

- Pre-silicon
  - Logic bugs

# Traditional Bounded Model Checking

Property

Design



If property violated



Counter-example = bug trace

# Traditional BMC vs. Symbolic QED

	Traditional BMC	Symbolic QED
Properties	Manual	<b>Automatic QED checks (Universal property)</b>
Design size	Small blocks	<b>Large SoCs</b>
False fails	Possible	<b>None</b>

# BMC Symbolic QED

“Universal” Property +  
QED-consistent starting state

Design + QED Module  
(no hardware overhead)



If property violated



Counter-example = bug trace

# Symbolic QED Results

- Many designs: processors, accelerators, ...

## Infineon case study



Several automotive microcontroller cores

**All recorded bugs detected (+ more)**

20 seconds or less per bug


**Verification productivity: 8× – 60× improved**

# BUT...

- Big designs ?
- Solution: QED checks compositional
  - Preserved across partial instances

# Symbolic QED: Billion-Transistor SoCs

20 mins. to 7 hours



92 “difficult” bug scenarios (industrial bug databases)

OpenSPARC T2 SoC (500M transistors): difficult bugs inserted

# More Opportunities

- Hardware security
  - Derive new vulnerabilities
    - Beyond Spectre, Meltdown
  - Detect Trojans
- Firmware
- Large-scale systems



# Outline

- NanoSystems: *N3XT* 1,000×



# Improve Computing Performance

Design  
techniques



Energy  $\times$  Execution time





Device  
technologies



# Improve Computing Performance

Design techniques

Multi-cores



**Target:**  
**1,000× performance**

New innovations required

Device technologies


# Solution: NanoSystems

*Transform new nanotech*


*into new systems*

*enable new applications*


New devices




New fabrication



New sensors



New architectures




# Abundant-Data Applications


## Memory wall: processors, accelerators

Wide range of domains


Genomics  
classification



Natural language  
processing



Deep  
learning



■ Compute ■ Memory

# Nano-Engineered Computing Systems Technology



COVER FEATURE **REBOOTING COMPUTING**

## Energy-Efficient Abundant-Data Computing: The N3XT 1,000x


Mohamed M. Sabry Aly, Mingyu Gao, Gogo Hills, Chi-Shuen Lee, Greg Pitner, Max M. Shulaker, Tony F. Wu, and Mehdi Azregzain, Stanford University  
 Jeff Bokor, University of California, Berkeley  
 Franz Franchetti, Carnegie Mellon University  
 Kenneth E. Goodson and Christos Kozyrakis, Stanford University  
 Igor Markov, University of Michigan, Ann Arbor  
 Kuntia Oskotou, Stanford University  
 Larry Pileggi, Carnegie Mellon University  
 Eric Pop, Stanford University  
 Jan Rabney, University of California, Berkeley  
 Christopher Ré, H.-S. Philip Wong, and Subhasish Mitra, Stanford University

Next-generation information technologies will process unprecedented amounts of loosely structured data that overwhelm existing computing systems. N3XT improves the energy efficiency of abundant-data applications 1,000-fold by using new logic and memory technologies, 3D integration with fine-grained connectivity, and new architectures for computation immersed in memory.



# N3XT NanoSystems


Computation immersed in memory



# N3XT NanoSystems

Computation immersed in memory

Increased functionality



***Impossible with today's technologies***



# N3XT Computation Immersed in Memory

3D Resistive RAM  
Massive storage




1D CNFET, 2D FET  
Compute, RAM access

STT MRAM  
Quick access

1D CNFET, 2D FET  
Compute, RAM access

1D CNFET, 2D FET  
Compute, Power, Clock




No TSV

Ultra-dense,  
fine-grained  
vias



Silicon  
compatible

# Carbon Nanotube FET (CNFET)

CNT:  $d = 1.2\text{nm}$



CNFET



## Energy Delay Product

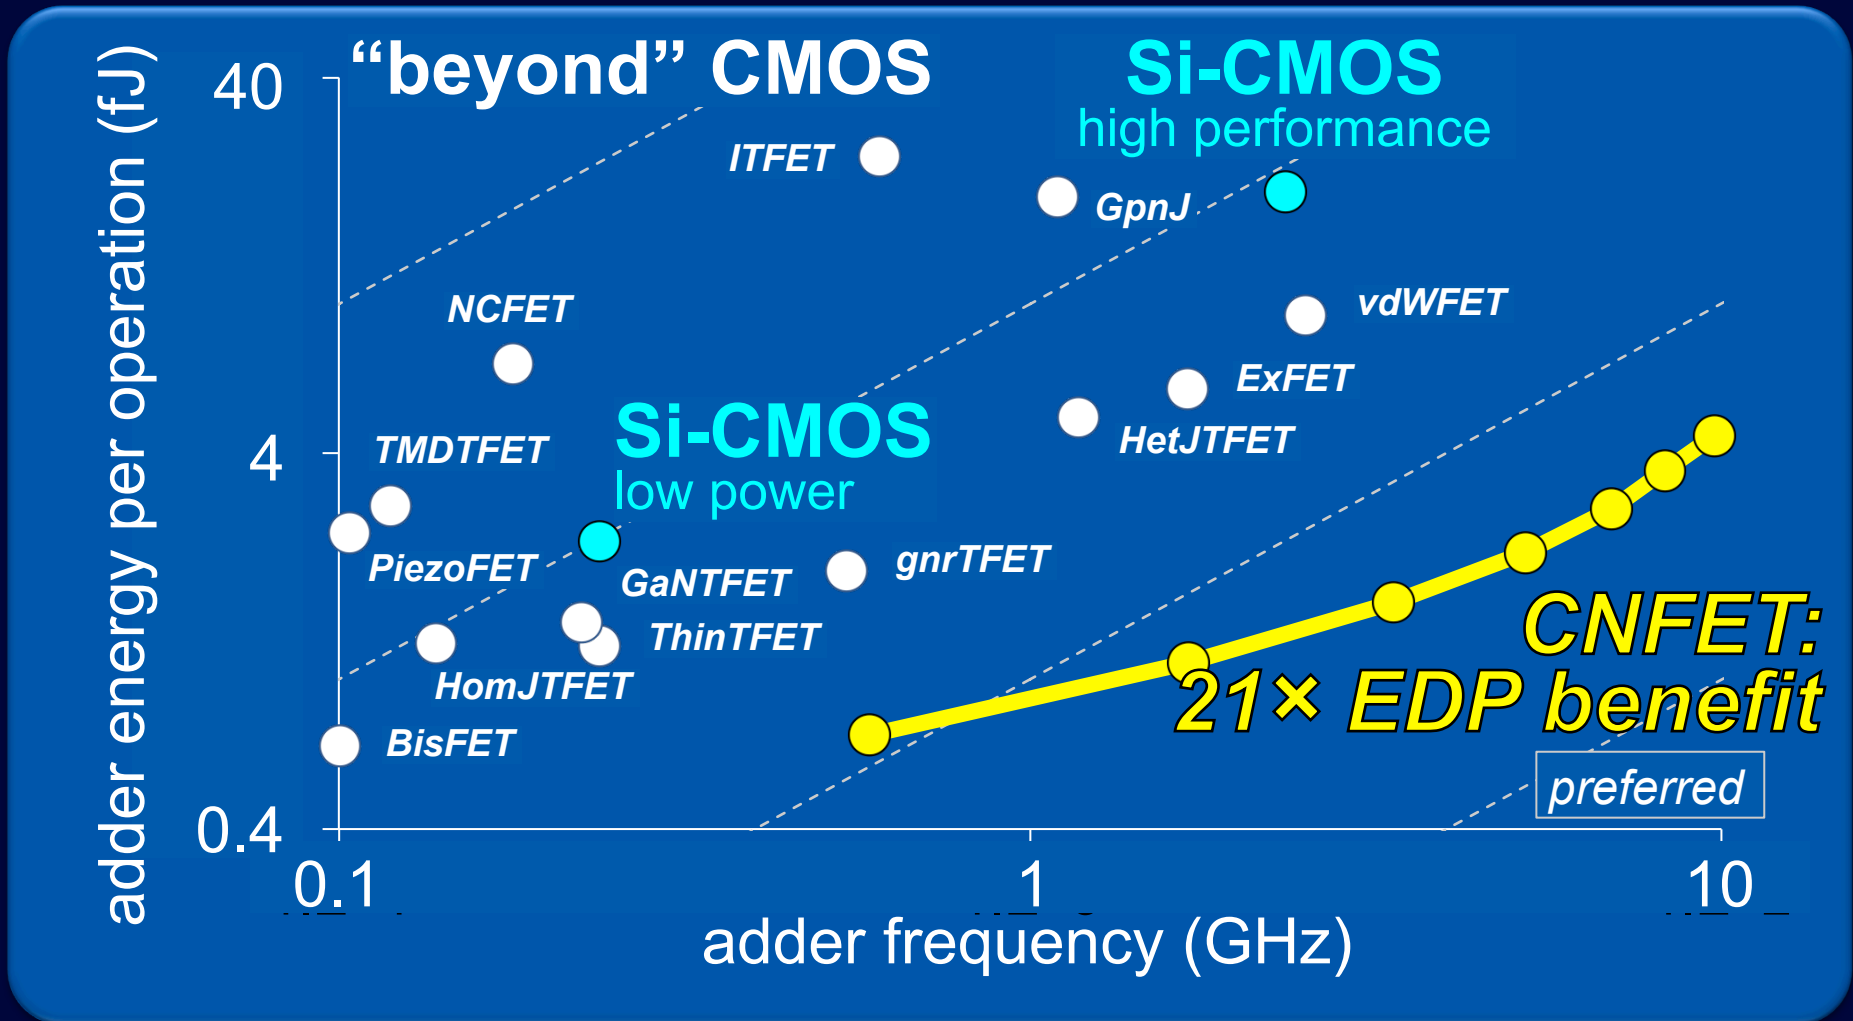
- $\sim 10\times$  benefit

## Processor case studies

[Stanford + IMEC + TSMC,  
IBM, other commercial]

# Putting into Perspective


- Existing technology benchmarking + **CNFETs**




# Big Promise, Major (Past) Obstacles

- Process advances alone inadequate

## Mis-positioned CNTs




## Metallic CNTs



**Imperfection-immune paradigm**


# Imperfection-Immune VLSI

## Mis-positioned CNTs



- Arbitrary logic functions

## Metallic CNTs





- Scalable m-CNT Removal

# Most Importantly

- VLSI processing
  - No per-unit customization
- VLSI design
  - Immune CNT library



# CNT Computer






**10X EDP, BUT...**

**How can we do better ?**

# 3D Integration


- Massive ILV density  $\gg$  TSV density

TSV (chip stacking)



Through silicon via  
(TSV)


Dense, e.g., monolithic



Nano-scale  
inter-layer vias (ILVs)


# Device + Architecture Benefits

Naturally enabled:  $< 400\text{ }^{\circ}\text{C}$  fabrication



# 3D NanoSystem


2 Million CNFETs, 1 Mbit RRAM







# 3D NanoSystem

Abundant data: Terabytes / second




Millions of sensors



Ultra-dense vertical connections  
x100,000

Memory

1 Megabit RRAM




CNT computing logic

classification accelerator

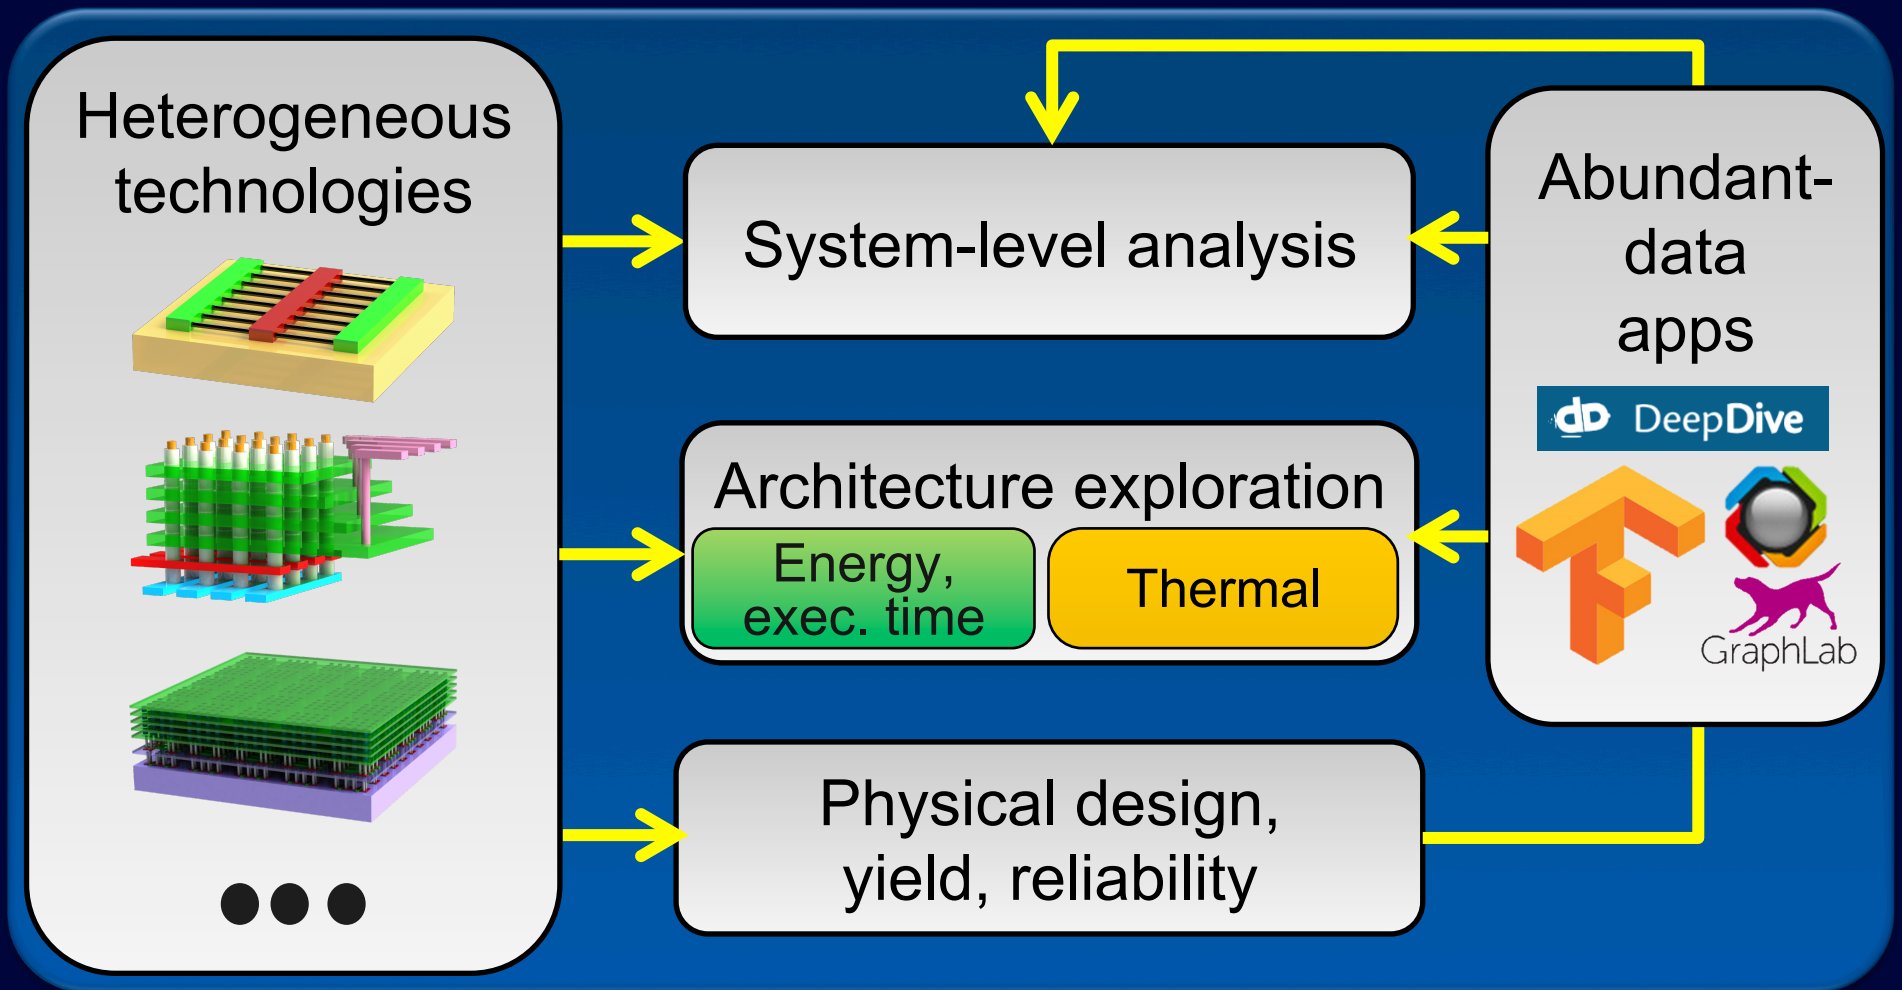
***In-situ classification***

***Extensive, accurate***



# N3XT Simulation Framework

Joint technology, design & app. exploration




# Massive Benefits: Deep Learning, Graph Analytics, ...



IBM graph analytics



**~1,000× benefits, existing software**



# Many NanoSystem Opportunities

## Cross-layer solutions

Tech. + design + app.

### 1. Multiple bits / RRAM cell

Special RRAM programming +  
weight encoding

### 2. New RRAM resilience


10-year continuous inference

[Le IEEE TED 19, Wu ISSCC 19]  
Stanford + CEA LETI

## Brain-inspired $\supset$ neural nets

e.g., Hyperdimensional (HD)

1. One-shot learning
2. Classify language: 98% accuracy



Live ISSCC demo


[Wu ISSCC 18, IEEE JSSC 18]  
Stanford + Berkeley




# DARPA 3DSoC Program



Max Shulaker  
Anantha Chandrakasan



Subhasish Mitra  
H.-S. Philip Wong  
Simon S. Wong



Brad Ferguson  
Mark Nelson



Jefford Humes

# Outline

- Conclusion

# Neural Interfaces: Natural Resolution

1. Treat disorders, augment capabilities
2. Understand brain
3. New neuro-inspired computing ?

New abstractions, new ICs, closed-loop brain experiments



Collaborator: Prof. E.J. Chichilnisky, Stanford

# Conclusion

- Robust operation, performance, new applications
  - Big opportunities
  - New solutions: elegantly simple, effective

## QED & Symbolic QED

Pre-silicon & post-silicon

Automatic, overnight

Billion-transistor scale

## N3XT

NanoSystems

Ultra-dense & fine-grained 3D

1,000× opportunity