



# Raven/Hurricane/Resiliency Tape-outs


Alon Amid, Rimas Avizienis, Stevo Bailey, Milovan Blagojevic, Chris Celio, Po-Hung Chen, Pi-Feng Chiu, Martin Cochet, Henry Cook, Palmer Dabbelt, Jessica Iwamoto, Vighnesh Iyer, Ruzica Jevtic, Ben Keller, Jaehwa Kwak, Hanh-Phuc Le, Yunsup Lee, Howie Mao, Nandish Mehta, Alberto Puggelli, Brian Richards, Colin Schmidt, Keertana Settaluri, Nicholas Sutardja, Andrew Waterman, John Wright, *Brian Zimmer*

Elad Alon, Krste Asanovic, Borivoje Nikolic

12/5/2017




# Motivation



- How can circuit designers and architects improve energy efficiency without transistor scaling?
- Fine-grained adaptive voltage scaling (AVS) maximizes energy efficiency
  - Fine-grained both spatially and temporally
- Looks beautiful on paper, but *many* issues

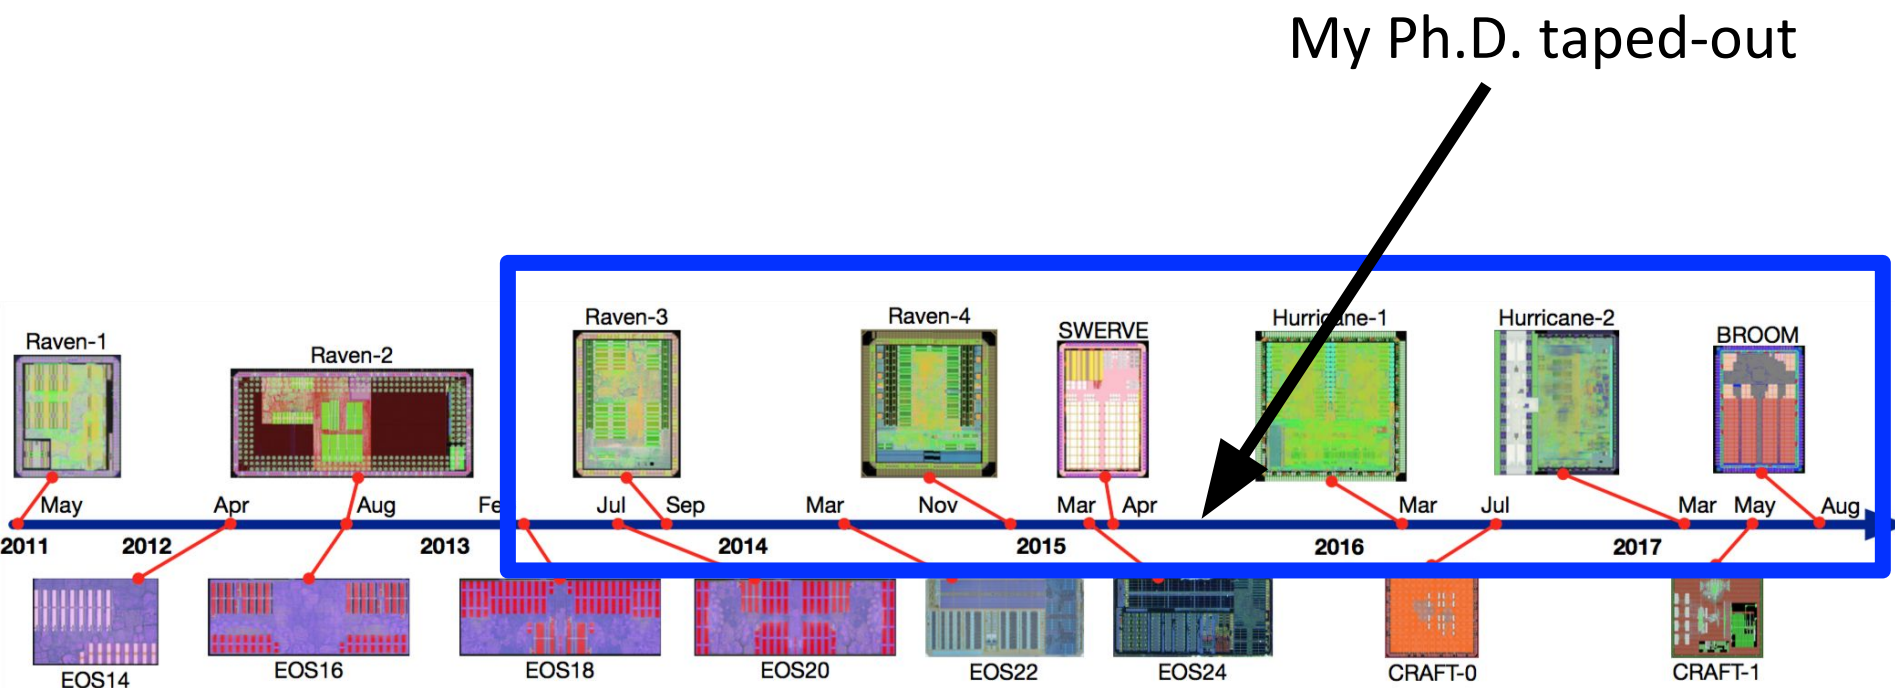
# Key problems

- **Efficiency:** How do you actually reduce voltage while maintaining conversion efficiency?
- **Granularity:** How do you supply independent voltages to small spatial domains?
- **Control:** How do you decide when to change the voltage?
- **Reliability:** How do you keep the circuits working at low voltage?
- **Practicality:** Can you actually build it?




# Why tape-out chips?

- Get papers into circuits conferences
- Bora already reserved X mm<sup>2</sup>, so we better not embarrass him
- **Ground research in reality**

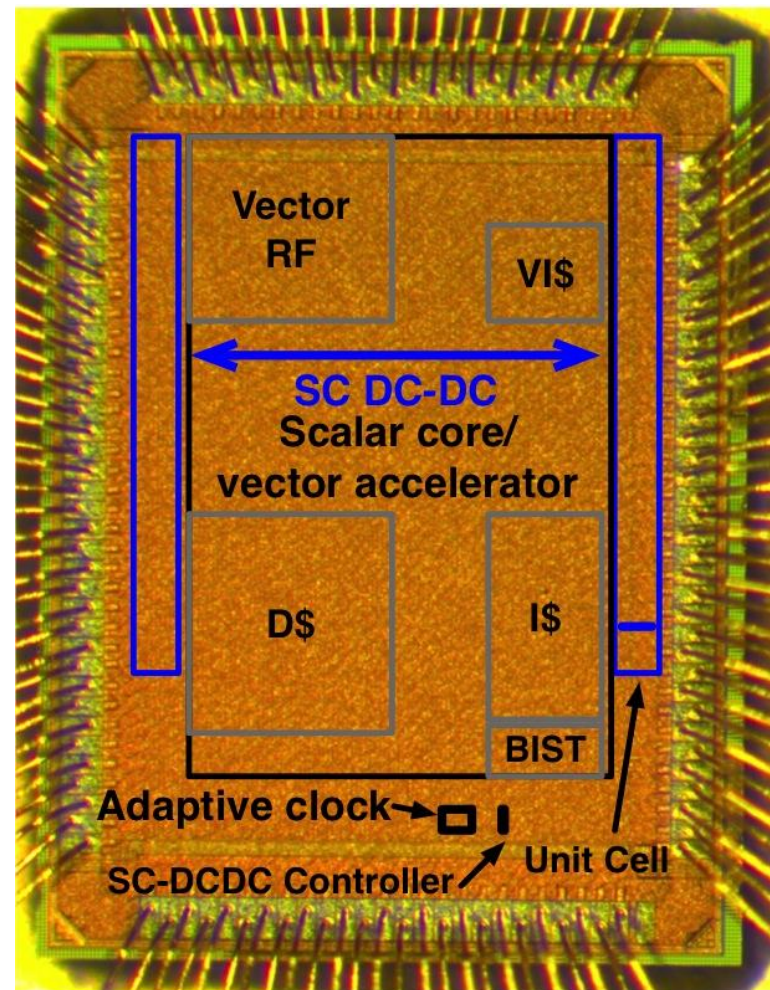


# Simultaneous-Switching Switched-Capacitor DCDC

- Division of input voltage possible using on-chip capacitors and switches
- Adaptive clock tracks resulting ripple





- Novelty: 1<sup>st</sup> system with simultaneous-switching DCDC
- ST 28nm FDSOI, 2.5mm<sup>2</sup>
- RISC-V processor
- 1V, 0.9V, 0.67V, 0.5V modes 2nF of on-chip capacitance
- 34 double-precision floating point GFLOPS/W (26 with DCDC)





# Results



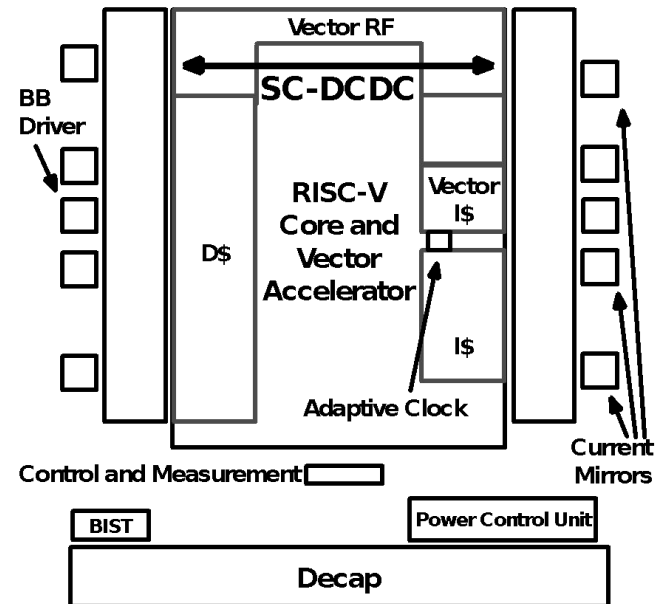
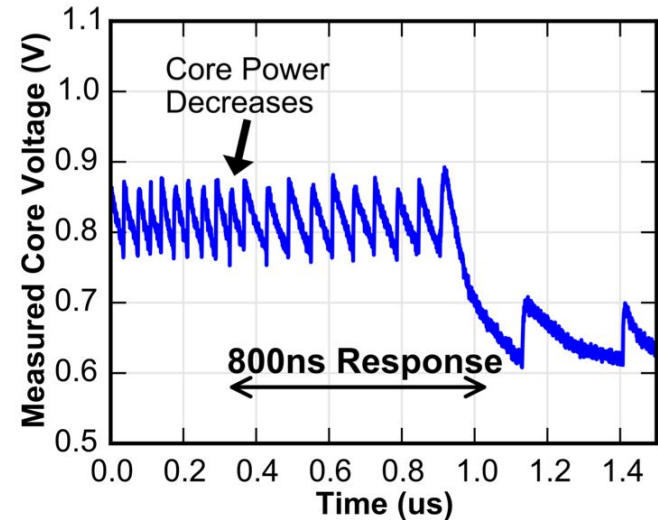
## **Strong foundation, start building a real system**

- Could change voltage to track activity...but we are just changing it once and measuring
- Could do this for multiple cores...but we only have 1 core
- Could build an efficient processor...but fake memory system



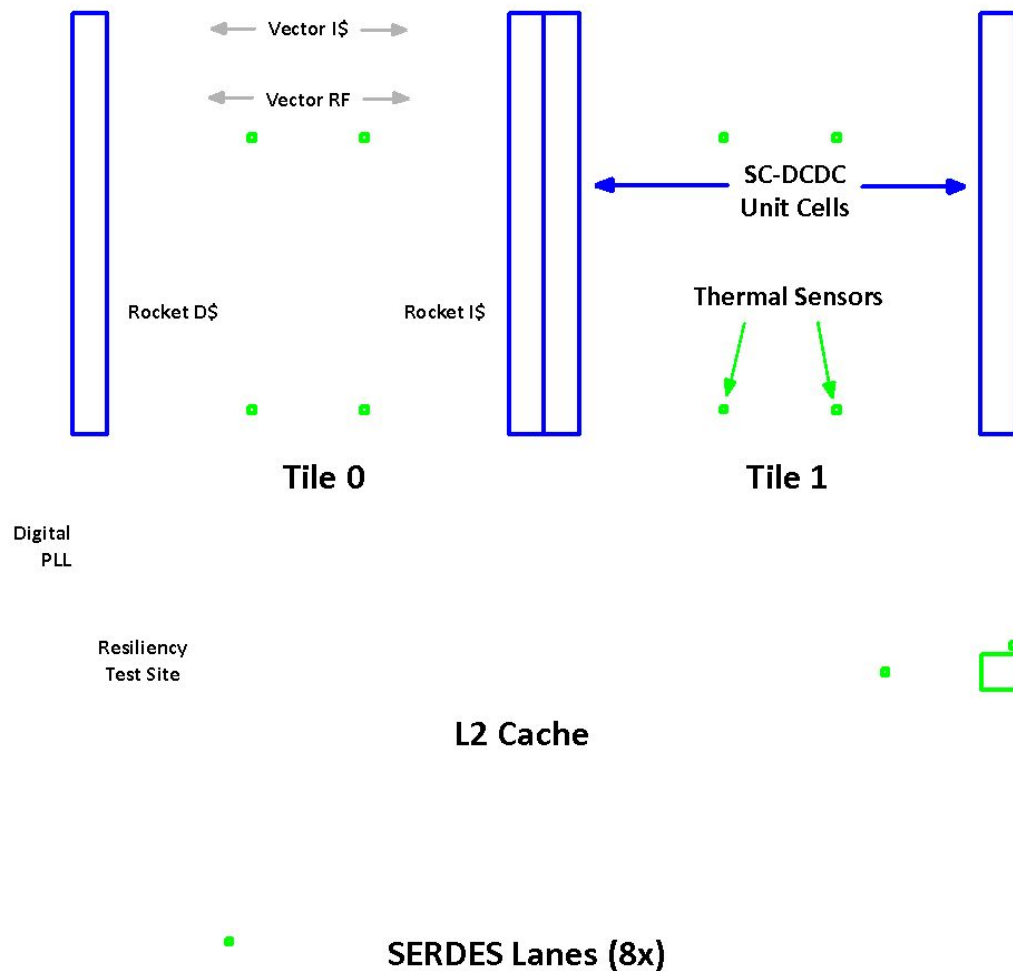
# Raven-4

- Novelty: Measure instantaneous power by counting the DCDC toggle frequency
- First Zscale processor as power management
- 28nm FD-SOI
- 3.03mm<sup>2</sup> die area
- 54 GFLOPS/W




# Hurricane-1

- Novelty: Two cores, faster memory system, counter-based power management
- 28nm FD-SOI
- 7.98mm<sup>2</sup> die area



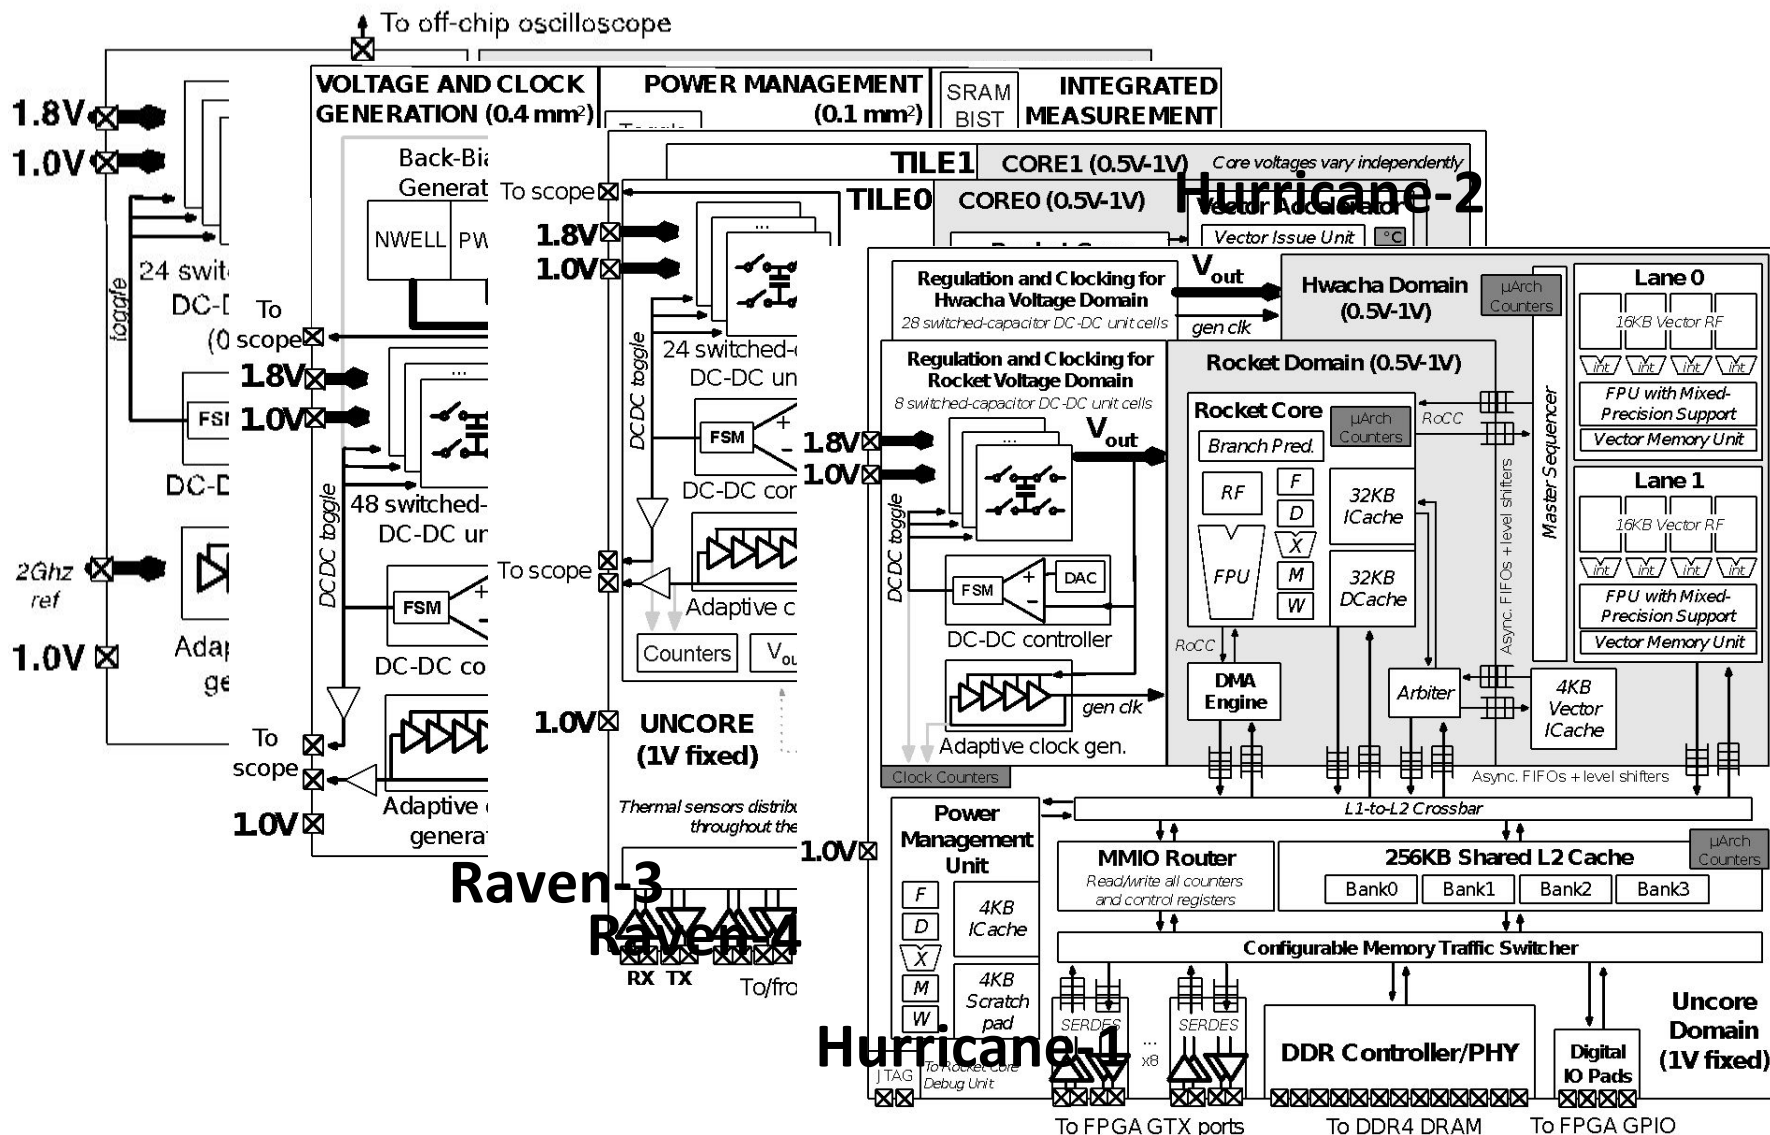
# Hurricane-2





- Novelty: Separate voltage domains for Rocket and Hwacha
- Many micro-architectural counters
- Actual DDR PHY
- 28nm FD-SOI
- 17.30mm<sup>2</sup> die area

# Evolution of chips

## Omnigraffle Figures




- Novelty: Avoid failing SRAM bitcells at low voltage
- RISC-V Rocket in-order processor+1MB L2 cache
- 2mm x 3mm TSMC 28nm





- Novelty: First BOOM core tape-out and line recycling for further  $V_{min}$  reduction in L2
- TSMC 28nm





# Come a long way in 6 years



- Raven-1:
  - May 2011
  - My first tape-out
  - Codename: Trainwreck
  - ‘Working’ RISC-V processor:

Yunsup Lee [yunsup@eecs.berkeley.edu](mailto:yunsup@eecs.berkeley.edu)

Good news. After setting the various voltages going to the SRAM, raven1 finally runs a very simple program. It can run the following program:

```

        addi $x1, $x0, 1
        addi $x1, $x0, 1
        addi $x1, $x0, 1
        addi $x1, $x0, 1
        addi $x1, $x0, 1
        mtpcr $x1, $scr16
1:      beq $x0, $x0, 1b
  
```



# **Lessons learned (from Git commits)**

# Importance of typing accuracy

```
commit 34dec1e9de7903a243995070eadded248022cde6
Author: Ben Keller <bkeller@eecs.berkeley.edu>
Date:   Tue Sep 16 17:20:59 2014 -0700
```

```
    Fixing Brian's stupid typos
```

```
commit 946dfce2f7f913d9a7878115485a8ff0b8ad4f20
Author: Stevo Bailey <stevo.bailey@eecs.berkeley.edu>
Date:   Tue Oct 14 09:53:57 2014 -0700
```

```
    [stevo]: fixing carriage return in mikis block
```

# Encourage good habits

```
commit 2b47309c0eedeee71c7c1d506091e5e0815b40f4
Author: Ben Keller <bkeller@eecs.berkeley.edu>
Date:   Tue Mar 7 21:50:05 2017 -0800
```

```
    Comment this out til I figure out what it does
```

```
commit 6d941d903cd4b67c24a2a30cebd83a83f6c01bb3
Author: John Wright <johnwright@eecs.berkeley.edu>
Date:   Sun Mar 6 00:13:33 2016 -0800
```

```
    Don't do that, it's bad. Don't be bad, be good.
```



# Nothing is ever fixed on the first try

```
commit eb328afb16aef097b152f9869b891574f2433506
Author: John Wright <johnwright@eecs.berkeley.edu>
Date:   Wed Mar 2 16:28:36 2016 -0800
```

really for real fix things this time

```
commit c2f4ff14fb0969607679945e6fdd884215058f14
Author: Yunsup Lee <yunsup@cs.berkeley.edu>
Date:   Wed Nov 14 13:15:59 2012 -0800
```

now it's actually fixed

```
commit b3a5e25f7c8295d6140cc1cb0bc0d8bc75388685
Author: Andrew Waterman <waterman@eecs.berkeley.edu>
Date:   Sun Nov 25 19:46:48 2012 -0800
```

fix D\$ writeback bug

I swear I did this last week... perhaps I am finally losing it!

```
commit 7fedd518515f1a3ee630fb44ba047eb11855889d
Author: Keertana Settaluri <ksettaluri6@berkeley.edu>
Date:   Sun Nov 20 13:51:08 2016 -0800
```

Dis pad frame. Last edit. Last commit. Plzzz

# Always preemptively apologize

```
commit e8ed04afec35f56b4cdc6bd201951032a6fc1a84
Author: John Wright <johnwright@eecs.berkeley.edu>
Date:   Wed Mar 1 20:33:55 2017 -0800
```

Wow this is such a hack

```
commit c2fbed4b6b528e29f4e33ba84f5c25abad5c7af4
Author: Brian Zimmer <bmzimmer@eecs.berkeley.edu>
Date:   Thu Oct 23 15:35:46 2014 -0700
```

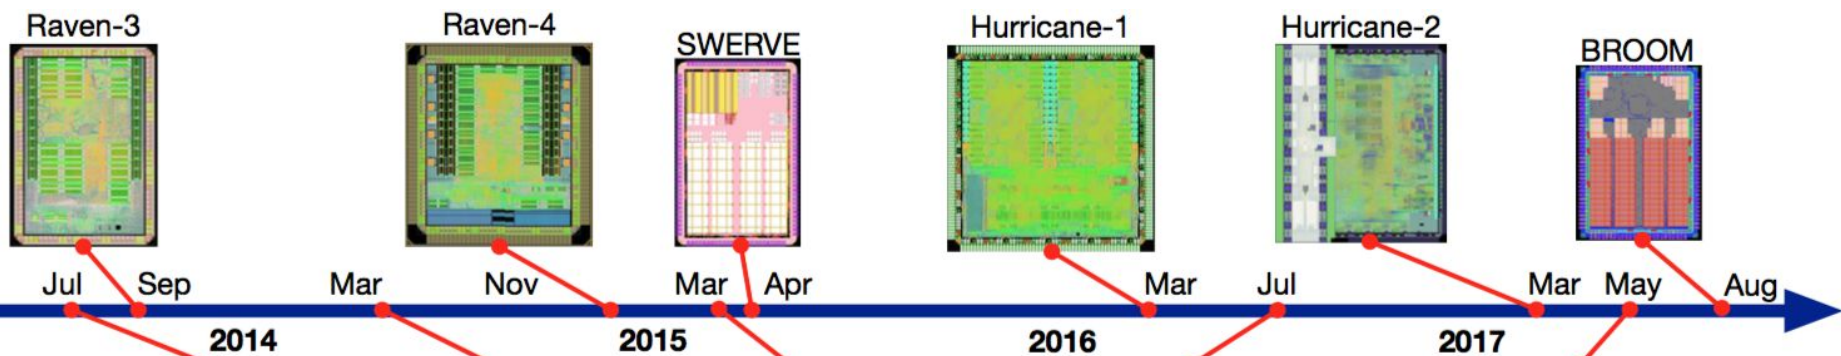
added dcdc comparator clock tree...if I broke something I'll buy you a beer

```
commit cc2e49f7d425fe3a9e8ce064921d2b04461c7d99
Author: Jaehwa Kwak <jhkwak@bwrc.eecs.berkeley.edu>
Date:   Tue Nov 11 20:33:17 2014 -0800
```

.... sorry...

# Conclusion

- **Efficiency:** >85% conversion efficiency for 0.5V to 1V voltage range
- **Granularity:** Smaller than per-core DVFS
- **Control:** 100ns responses to change in activity
- **Reliability:** 25% Vmin reduction with < 2% area overhead
- **Practicality:** Agile hardware booting Linux with small team of graduate students



# Acknowledgements

- Fabrication donation by STMicroelectronics and TSMC
- ASPIRE Members