What I learned building a parallel processor company from scratch

by Andreas Olofsson (HIPEAC-2017)

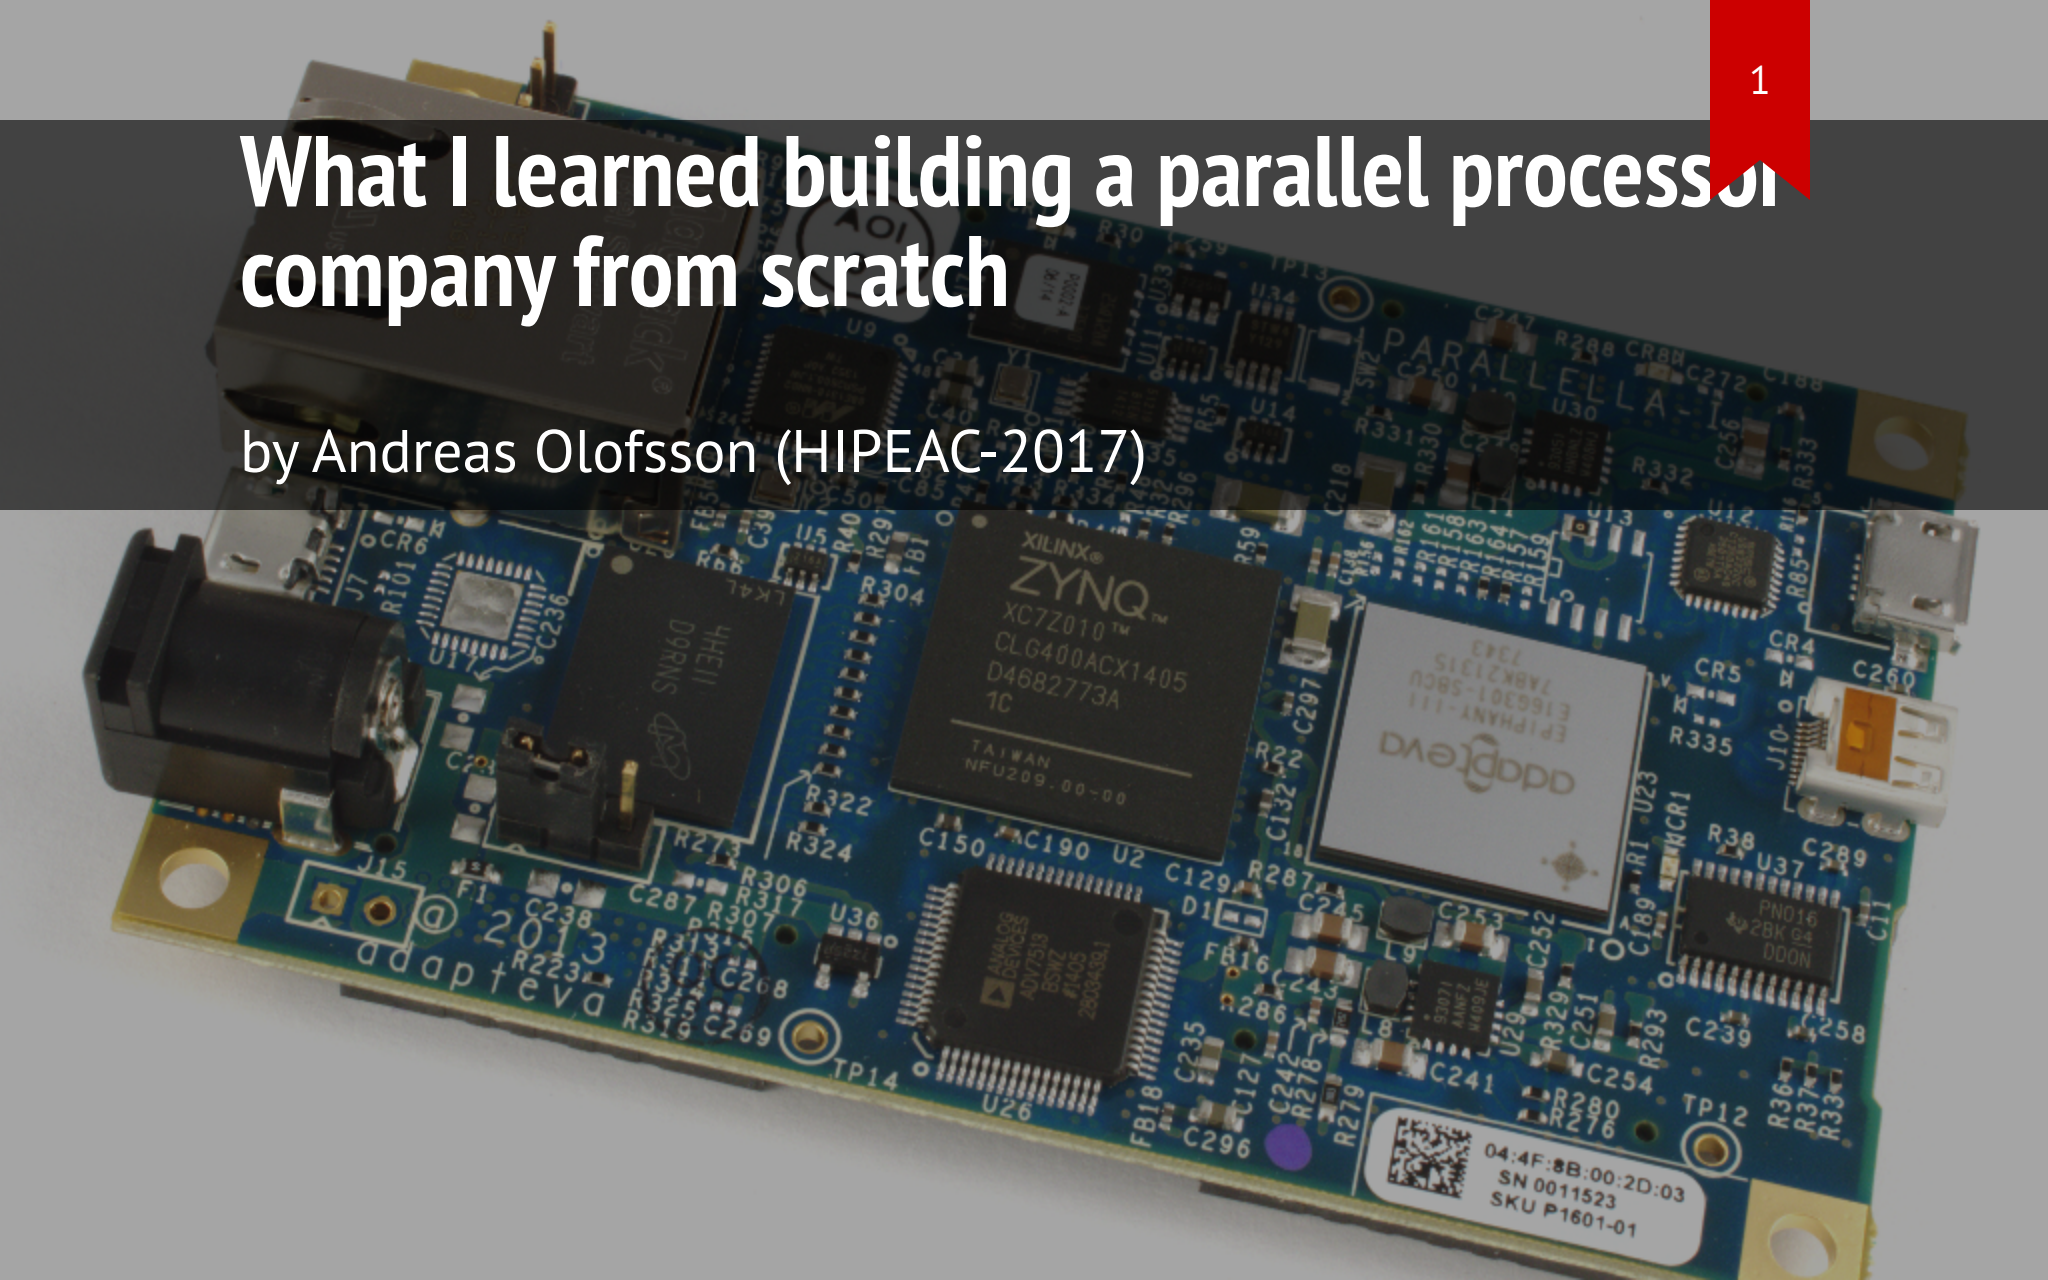

Disclaimer

> *This presentation summarizes lessons learned at Adapteva from 2008-2016. Statements and opinions are my own and do not represent my current employer.*

Adapteva Intro

- Founded in 2008
- 5 chip tapeouts to date with $7M raised
- #1 in Processor energy efficiency (100 GFLOPS/W at 16nm)
- First company to tape out a 1024 core 64-bit processor
- Only product that scales to 1 Billion cores
- Over 100 academic publications validating approach
- 10,000 customers

The beginning (1973)

Finding my passion (1998-2005)

- TigerSHARC DSP
- Designed in Israel
- We worked VERY hard
- We spent $100M!
- **Lessons:**
  - KISS
  - Don't fight technology

Finding your recipe (2006-2007)

- Camera CCD readout ICs
- Invented ISA in 2 weeks
- 3-4 person design team
- ">>" $100M in revenue
- **Lessons:**
  - Control architecture
  - Keep teams small
  - Profit from derivatives

Starting a company (2008)

- Sole founder at 35
- I KNEW my idea would work!
- Self funded with life savings
- Peak optimism/efficiency
- **Lessons:**
  - Check assumptions
  - Read
  - Find advisers

Creating technology (2008)

- Basic arch complete in May 2008
- All grant proposals rejected
- All publications rejected
- Feedback: impossible, not novel
- **Lessons:**
  - Publishing is broken!
  - Grant system is broken!
  - …or I am broken…

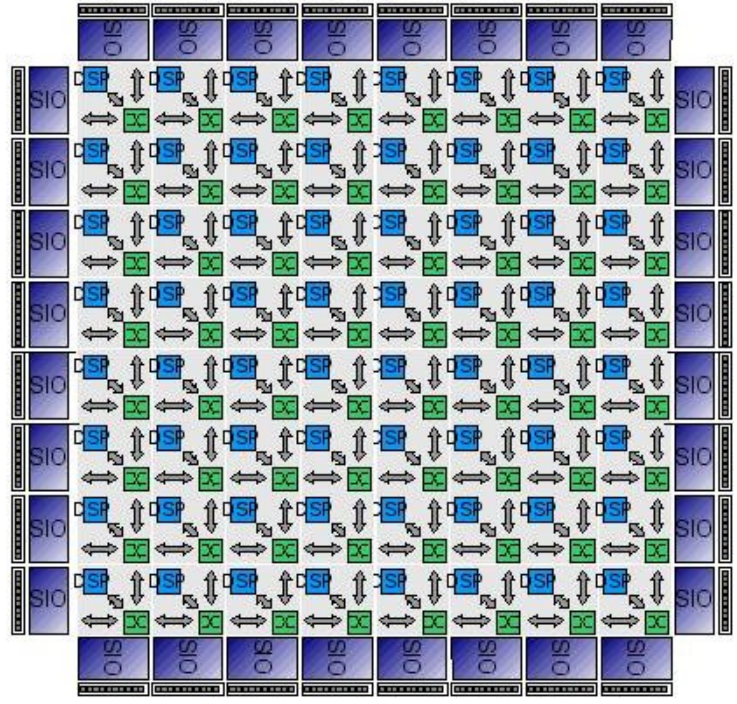

Validating technology (2008)

- Layout completed in Oct 2008
- 1GHz operation and 50 GFLOPS/W
- Gave me courage to ask for $$
- **Lessons:**
  - Trust your instincts
  - Leverage contacts
  - Get to MVP fast!
  - VCs need $$ validation

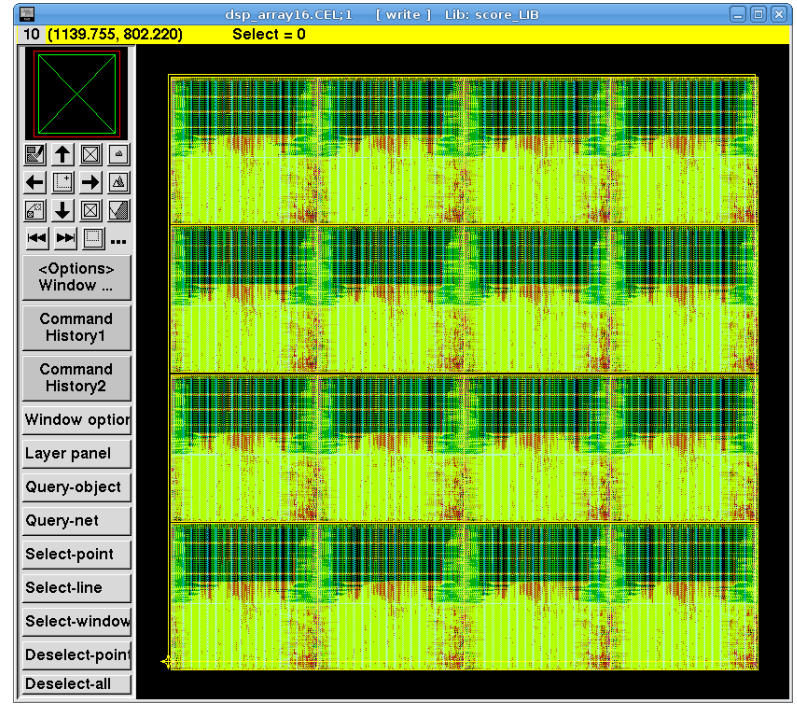

First prototype (2009)

- Epiphany-I prototype
- 16-cores/65nm/12mm^2
- "Solo effort"
- 6 weeks to tapeout
- **Lessons:**
  - Fight the fear
  - Sab Kuch Malega
  - Foundry + EDA were key

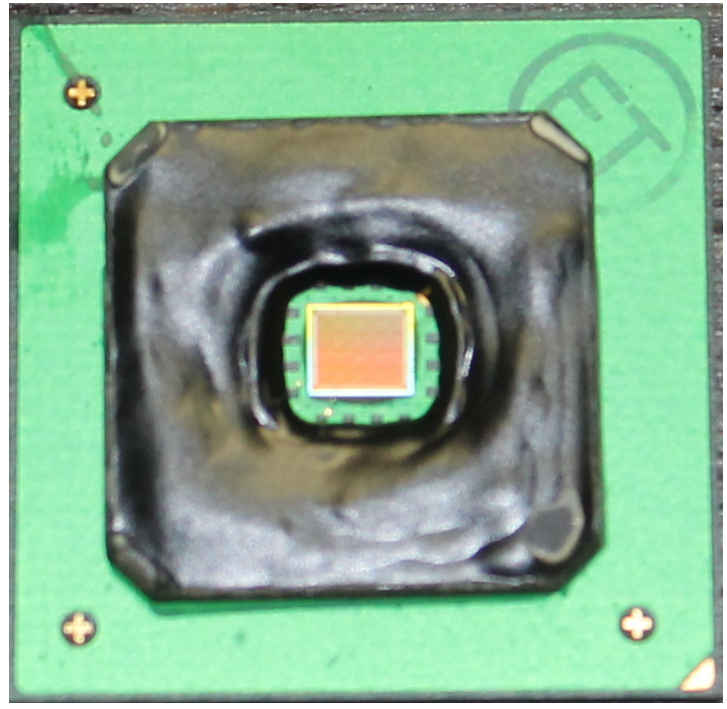

Team building (2010)

- Hired Roman, Oleg
- Best people I knew
- DV+design compliment
- Arch+DV+Design
- **Lessons:**
  - Hire people you know
  - Hire diversity
  - Hire technical first

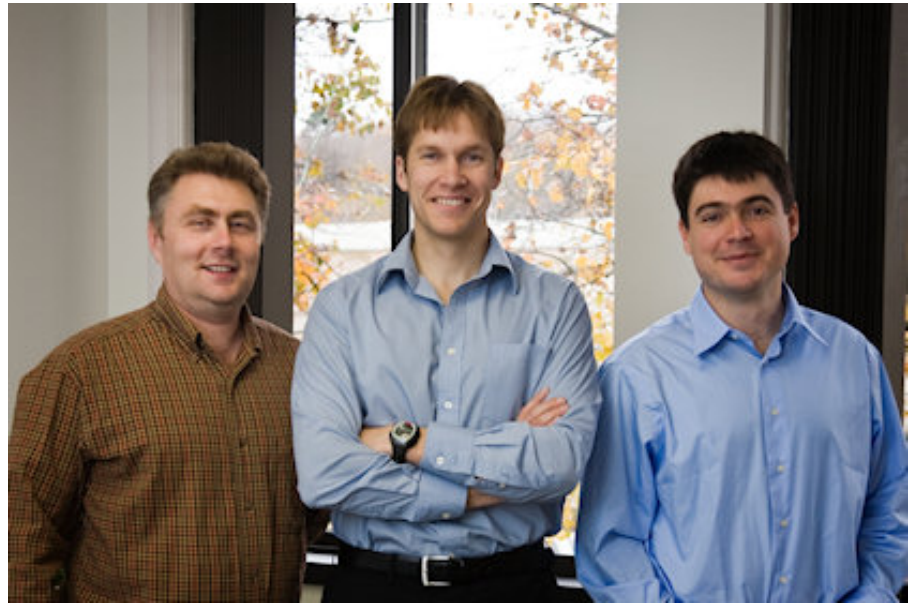

Build a product (2010)

- Dec 2010 tapeout
- Epiphany-III
- 16 cores, 65nm
- 50,000 built to date
- **Lessons:**
  - Get it right
  - But never stop refining
  - Be flexible and fast

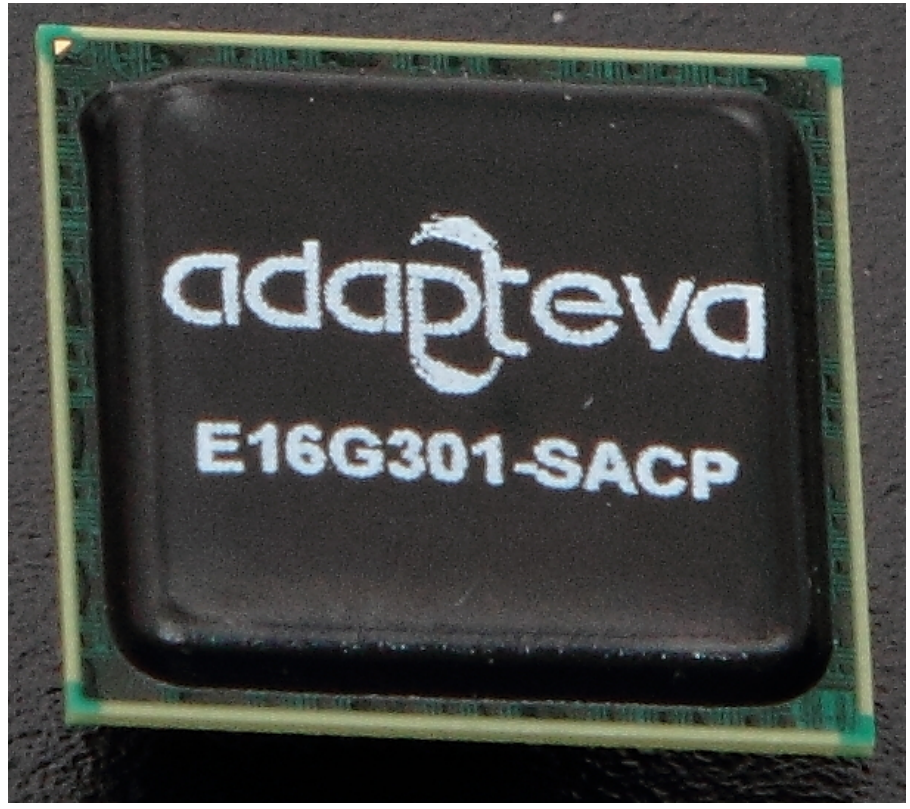

Nirvana (2011)

- Chips arrived may 2011
- Worked perfectly!
- Delivered to Bittware
- I was "done", felt at peace
- **Lessons:**
  - Get the right partner
  - Tech is the easy part!

Epiphany-IV (2011)

- 64-cores/28nm/11mm^2
- 70 GFLOPS/W, world #1
- Also created 1024 core layout
- Lots of tire kickers!
- **Lessons:**
  - Create COMPLETE solutions!
  - $$ talks, BS walks
  - Be opportunistic

Marketing Lessons (2011-2016)

- Lots of press!
- But no market pull…
- **Lessons:**
  - Press != traction
  - Semi sold door to door
  - Money talks, BS walks
  - Beware time parasites

Selling Lessons (2011-2016)

- I failed
- Bittware failed
- Super sales guys failed
- **Lessons:**
  - Beware the sales guy myth
  - Can't force sales
  - Get honest feedback
  - Sell solutions, not components

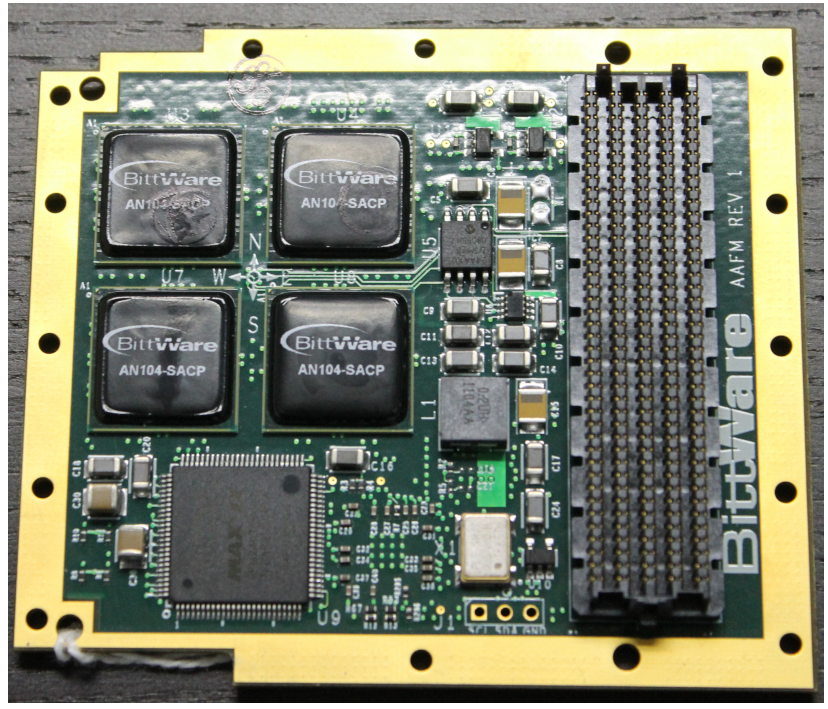

Kickstarting Parallel Computing (2012)

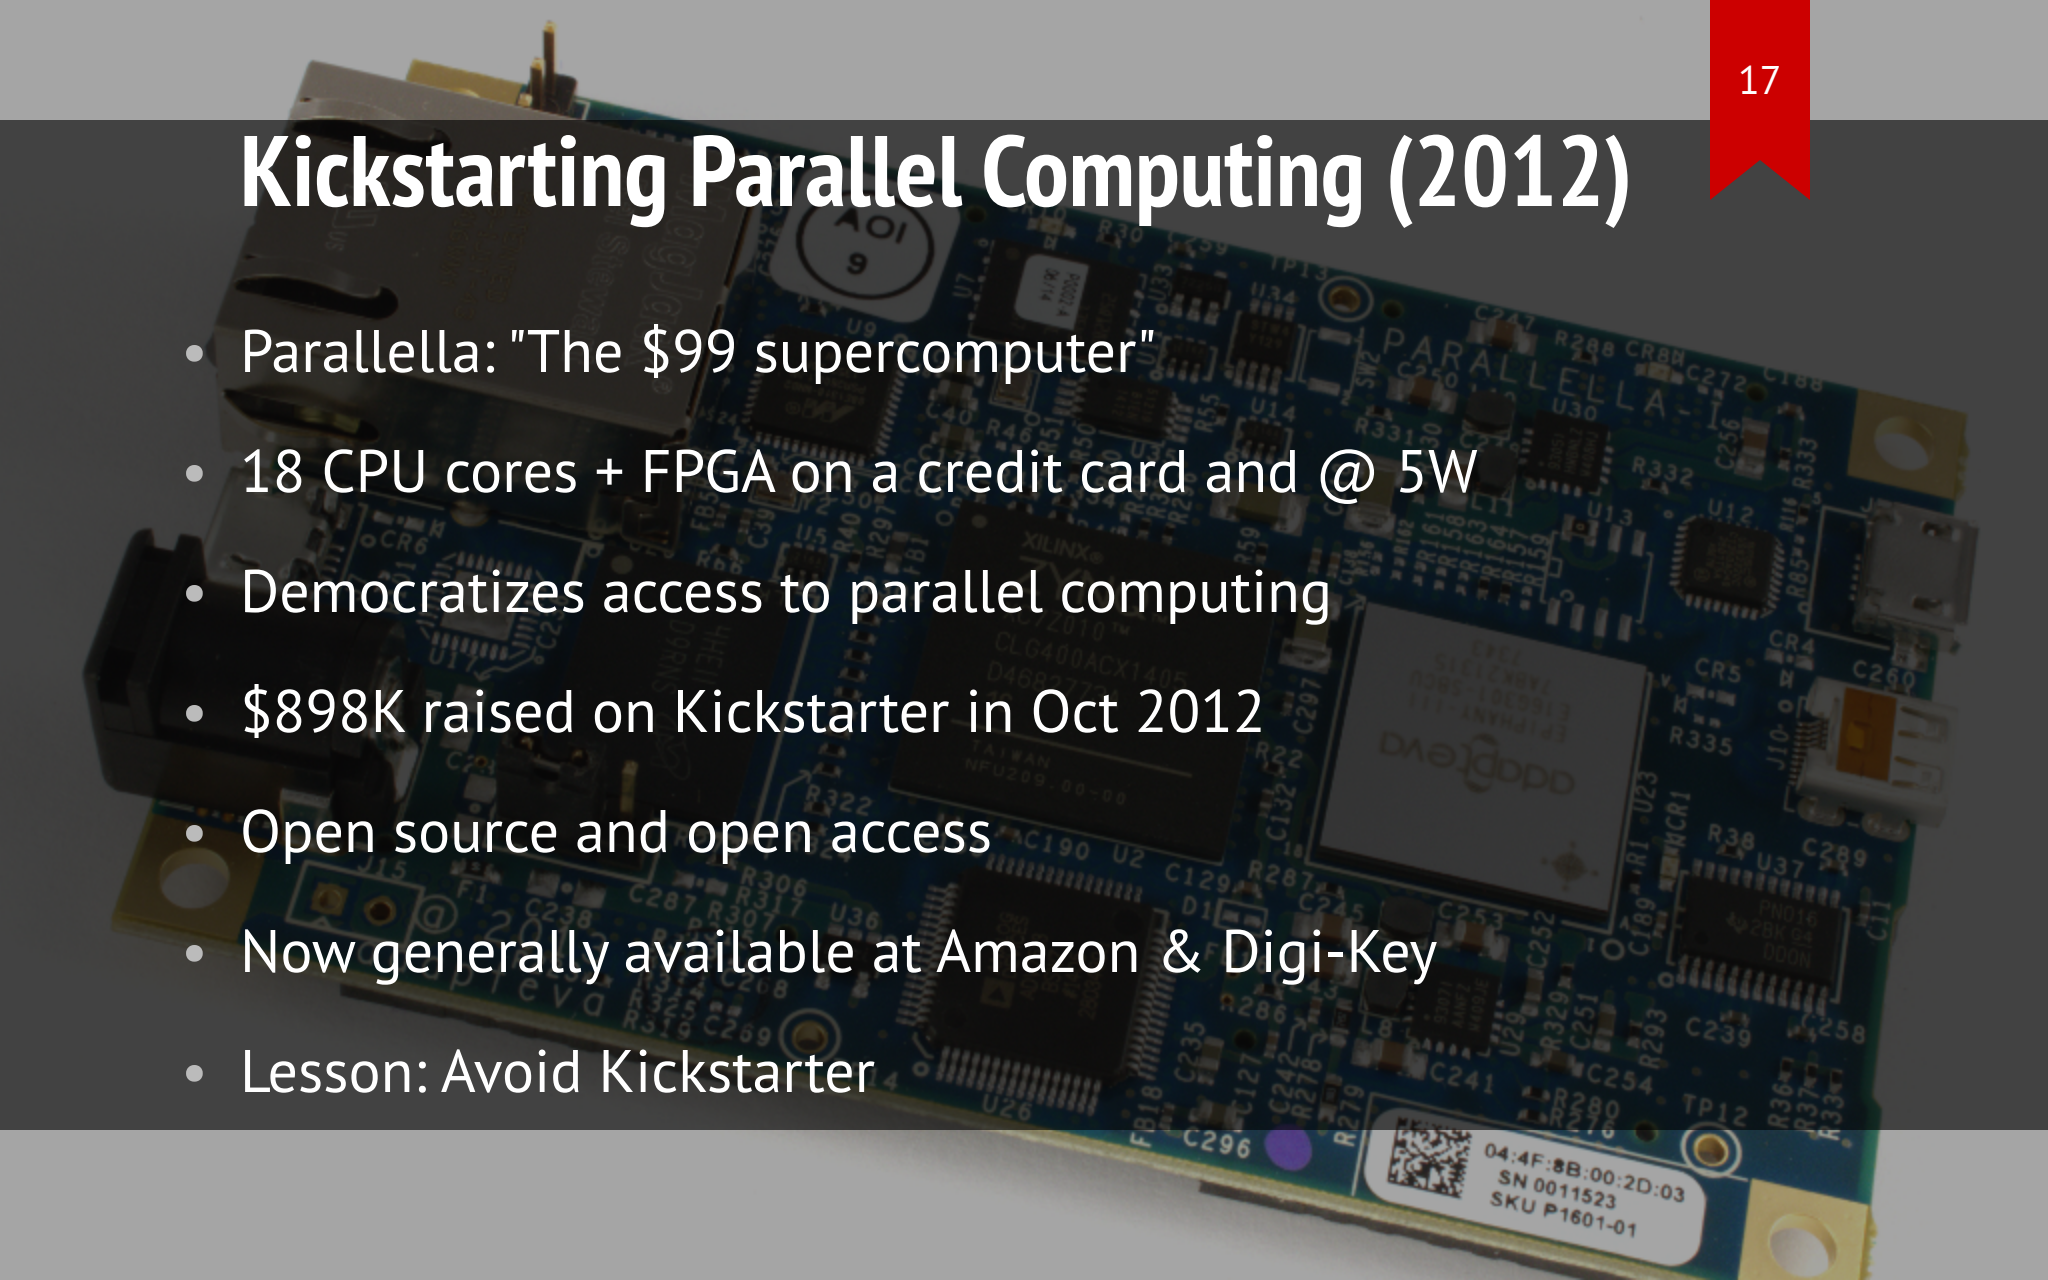

- Parallella: "The $99 supercomputer"
- 18 CPU cores + FPGA on a credit card and @ 5W
- Democratizes access to parallel computing
- $898K raised on Kickstarter in Oct 2012
- Open source and open access
- Now generally available at Amazon & Digi-Key
- Lesson: Avoid Kickstarter

Technology Lessons (2008-2016)

- Creating an ISA is not hard, but a lot of work
- 2D Meshes are fantastic
- Tiled layout are magical
- Heterogeneous is the present/future
- Parallel architectures are still not mainstream
- Moore's law still going pretty strong
- Chip design costs wildly exaggerated

Importance of Reducing Friction

Typical Skill	Unicorn	Researcher	Maker	Consumer
Board design	Yes	No	No	No
Software plumbing	Yes	No	No	No
System administration	Yes	No	Yes	No
Soldering, assembly	Yes	No	Yes	No
Building application	Yes	Yes	Yes	No
RTFM	Yes	Yes	Yes	?
Total Reach	10	10K	10M	1B

A/B Community Testing (2012-2016)

Effort	Result
PAL bounty program	30 functions of dubious quality
Github OSS efforts	Partial fail, limited contributors
supercomputer.io	Failed, not enough interest
Camera bounty	Failed, not completed
Workshop	1 successful, 1 failed
University program	Failed
Conference talks	Failed

Partnering (2012-2016)

- Magma (Synopsys) : EDA
- BittWare : channel/customer/investor
- GlobalFoundries: manufacturing
- Ericsson: investment, evaluation
- ARL: software development
- Evilcorp (beware): predatory

Fundraising (2008-2016)

- Total ($7.2M) : $200K, $900K (debt), $1.5M, $3.6M, $1M
- I raised self/grants/kickstarter/angels/ff/VC/corporate
- I Spent 1,000 hours to close deal with 3 major investors!
- All deals closed due to personal relationship
- Don't raise VC money without momentum
- Understand how VCs and corporate investors think

Manufacturing

- Need Tier-1 manufacturing partners to compete
- Beware of the valley of death (1,000-->10,000 units)
- Building complex <$100 products in US is not recommended
- Beware EOLs and non-standard parts!
- Only use standard parts available at Digikey
- Always negotiate pricing directly with manufacturer
- Getting above 95% yield requires engineering effort

Epiphany-V Status

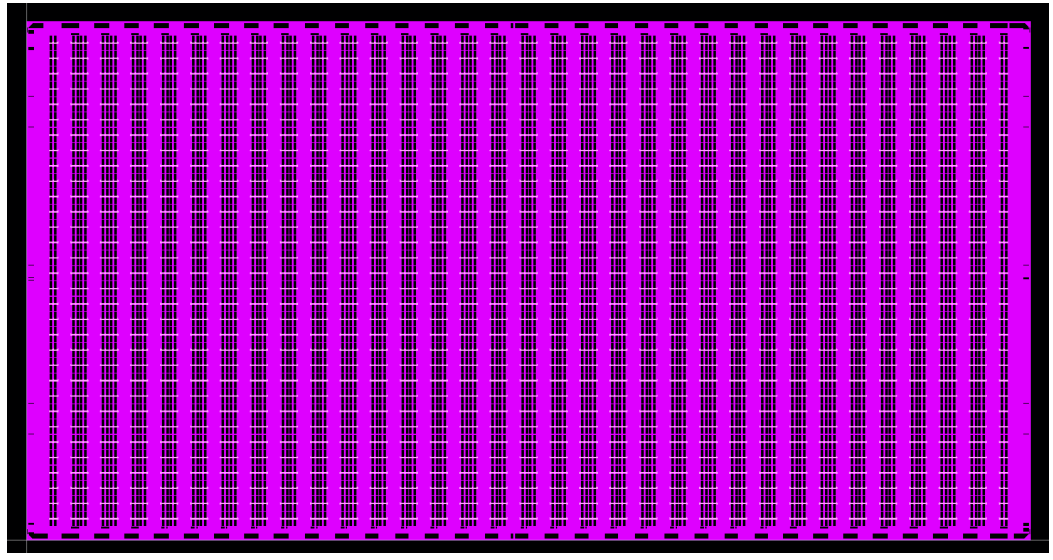

- 1024 64bit cores
- 16nm, 117mm^2
- 4.5B transistors
- 64 MB SRAM
- 1024 GPIO signals
- One full-time designer
- Extended ISA for deep learning, comms, crypto
- Dies are back, silicon bringup starting March 2017

Lesson Summary

- **Team:** Need team lined up from day one
- **Marketing:** Emotional. Sell with pictures and stories.
- **Selling:** Validated market pull. Be persistent and patient
- **Fundraising:** Avoid if possible
- **Pricing:** Use low/high pricing model
- **Manufacturing:** One wrong partner can kill your company
- **Product:** Create solutions
- **Software:** Hardware without software is useless