



Reducing I/O Bottlenecks with Numerical Compression

Al Wegener, CTO
Simplify Systems
awegener@simplify.com



Simplify Systems, Inc.
Santa Clara, CA 95051
www.simplify.com
(888) LESS-BITS
(408) 249-1500

Samplify Investors

MOOG

2009



Computed Tomography

- Prism CT
- Exclusive licensee

IDT

2010




Wireless

- Prism IQ
- Strategic investor
- Acquired wireless infrastructure patent portfolio

Schlumberger

2011



High Performance Computing

- APAX
- Strategic investor
- Exclusive licensee for oil & gas exploration

mamiya

2011



Imaging

- Prism 3
- Strategic investor

Intel VPFPU – ISSCC Paper (Feb 2012)

ISSCC 2012 / SESSION 10 / HIGH-PERFORMANCE DIGITAL / 10.3


10.3 A 1.45GHz 52-to-162GFLOPS/W Variable-Precision Floating-Point Fused Multiply-Add Unit with Certainty Tracking in 32nm CMOS

Himanshu Kaul, Mark Anders, Sanu Mathew, Steven Hsu, Amit Agarwal, Farhana Sheikh, Ram Krishnamurthy, Shekhar Borkar

Intel, Hillsboro, OR

controlled by up to
overflow carries (c
ing signals at 24b.
Overall support fo
area by 12% for 4:

To realize perform:
sion accuracy, cer
cuits to calculate i



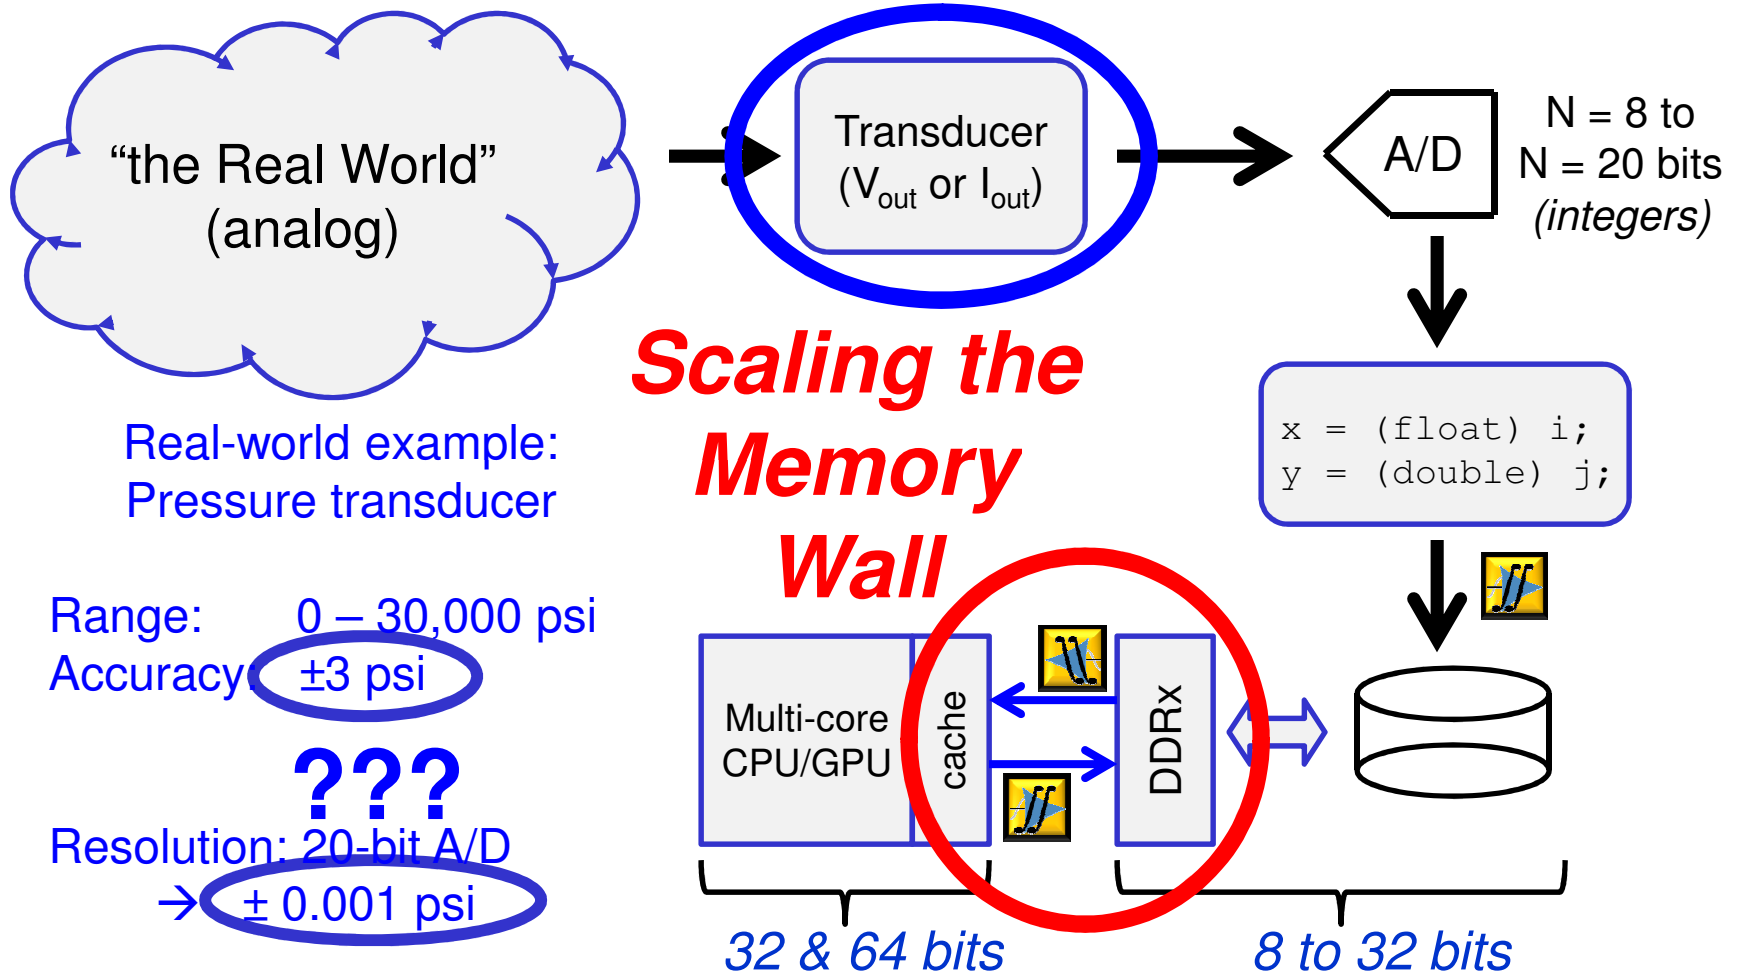
Variable Precision Floating Point Unit

(10.3) A 1.45GHz 52-to-162GFLOPS/W Variable-Precision Floating-Point Fused Multiply-Add Unit with Certainty Tracking in 32nm CMOS


- 1st reported variable-precision floating point unit with accuracy tracking for multiply-add
- Today's floating-point math wastes energy, time, and storage by using worst-case precision everywhere
- Using variable precision (24-bit→12-bit→6-bit) as needed can cut energy by 50%
- Uses NTV circuits for up to 7x further efficiency gain

...perhaps IEEE-754 is sometimes overkill ?


APAX's "Just Right" Numerical Resolution



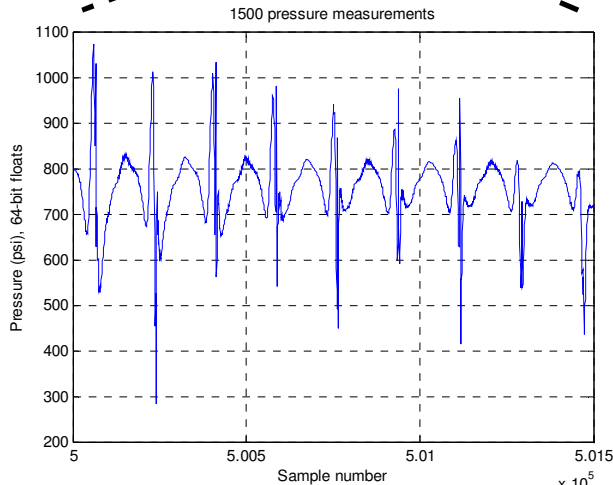
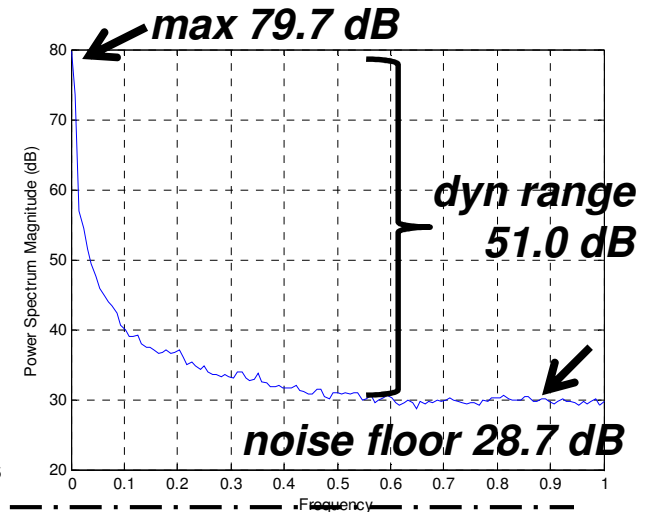
Example: How Much Compression?



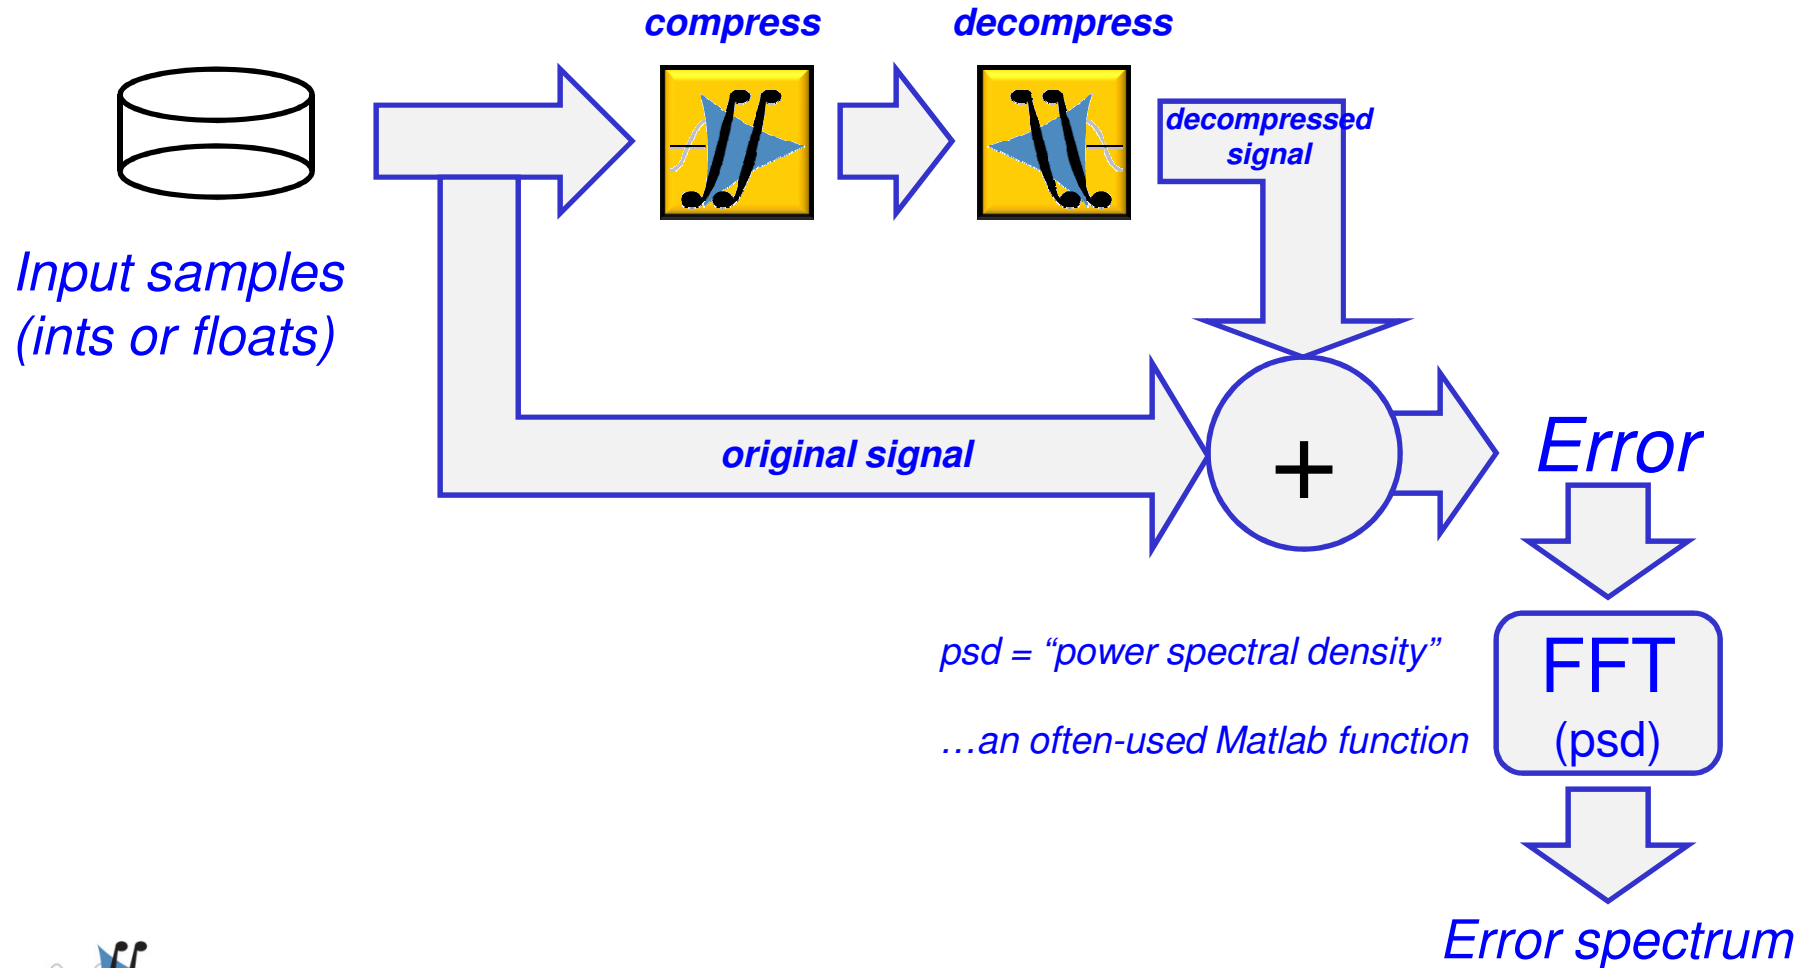
HISTOGRAM




SPECTRUM




Lossy Error = Original - Decompressed




Example: Error at 2:1 and 3:1 Compression



Range: {281.66, 1071.13 psi}
Blue: Original
Red: Decompressed
(do you see any blue ?)

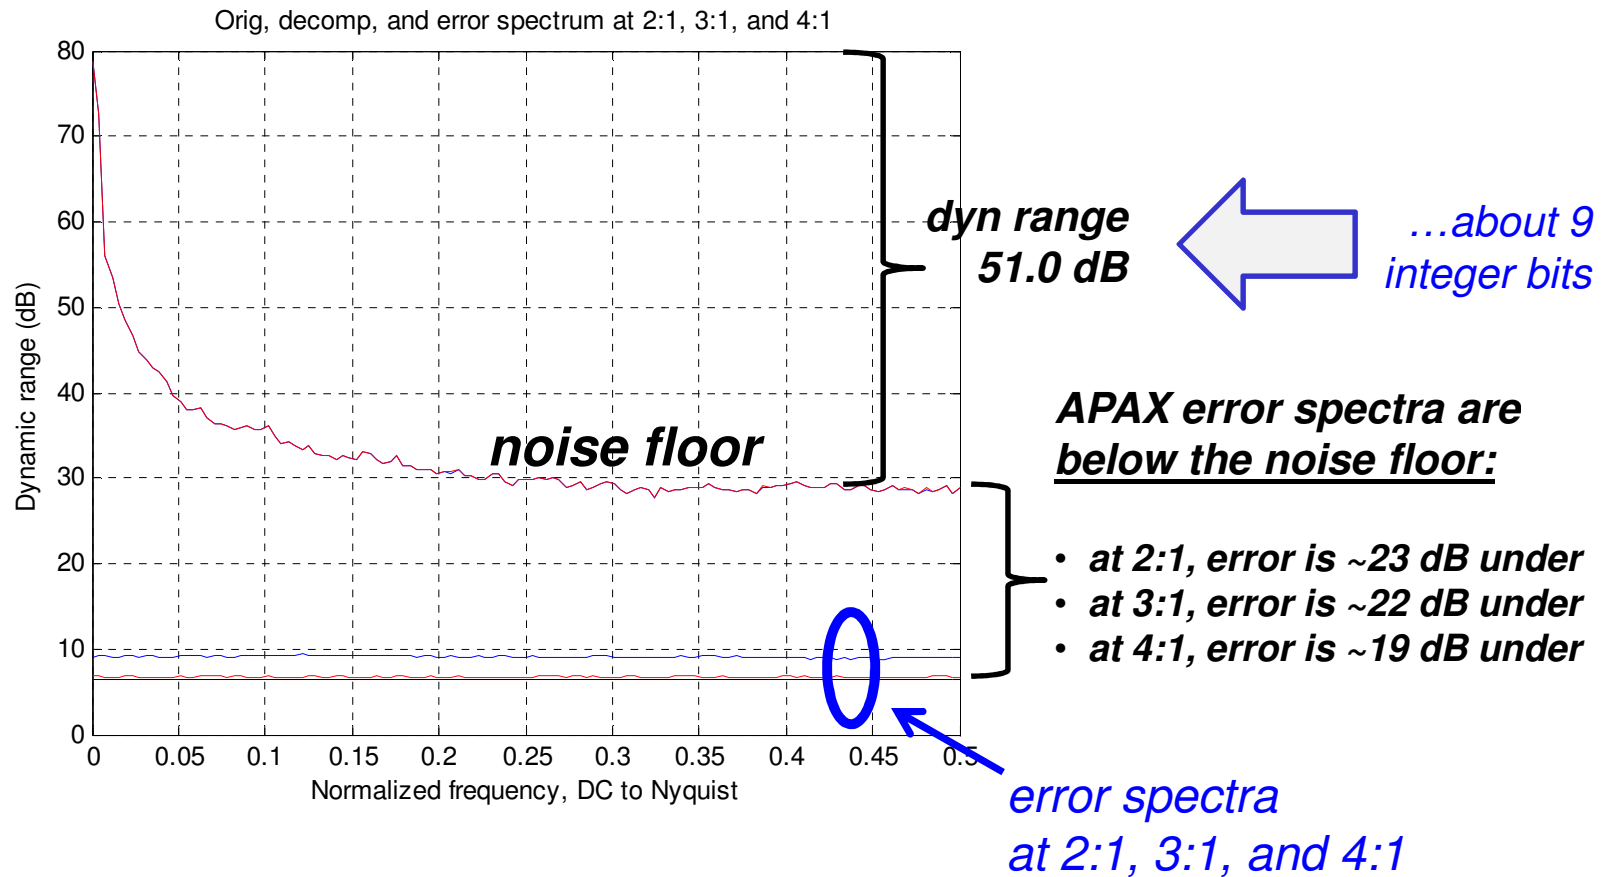


Error after 2:1 compression:
{-0.164, +0.189}, or
{-0.02%, +0.02%}
...vs. accuracy of ± 3 psi ?



Error after 3:1 compression:
{-1.099, +1.113}, or
{-0.13%, +0.14%}
...vs. accuracy of ± 3 psi ?

Example: Spectra at 2:1, 3:1, and 4:1




Lossless: OK. But LOSSY ? Yes ... 2:1 to 6:1

INTEGERS

Signal Type	Lossless C. R.	Fixed rate C. R. & quality metrics
Wireless baseband (3G, LTE)	1.2:1 – 1.5:1	1.6:1 – 2.3:1 EVM, PCDE, ACLR
Wireless RF (3G, LTE)	2:1 – 3:1	3:1 – 5:1 EVM, PCDE, ACLR
Computed tomography	1.6:1 – 2.7:1	3:1 – 4.5:1 Radiologists & SSIM
Ultrasound (ADC)	1.5:1 – 2:1	2:1 – 3:1 Sonographers & SSIM
Ultrasound (beamformer)	2:1 – 3:1	3:1 – 4:1 Sonographers & SSIM
Images & video	1.5:1 – 2:1	2:1 – 3:1 viewers, PSNR, SSIM
Oscilloscope (SerDes & LVDS)	1.3:1 – 2:1	2:1 – 4:1 BER, rise/fall time
Radar	2:1 – 3:1	3:1 – 5:1 p_d, p_{fa}
High-speed Imaging	1.4:1 – 2:1	2:1 – 4:1 Photographers, SSIM
Motion est. frames	1.6:1 – 2.5:1	2:1 – 6:1 Viewers, SSIM
Graphics textures	1.8:1 – 2.8:1	3:1 – 8:1 Gamers, SSIM


FLOATS




Aerospace/Defense



Automotive



Weather Forecasting



Medical Imaging



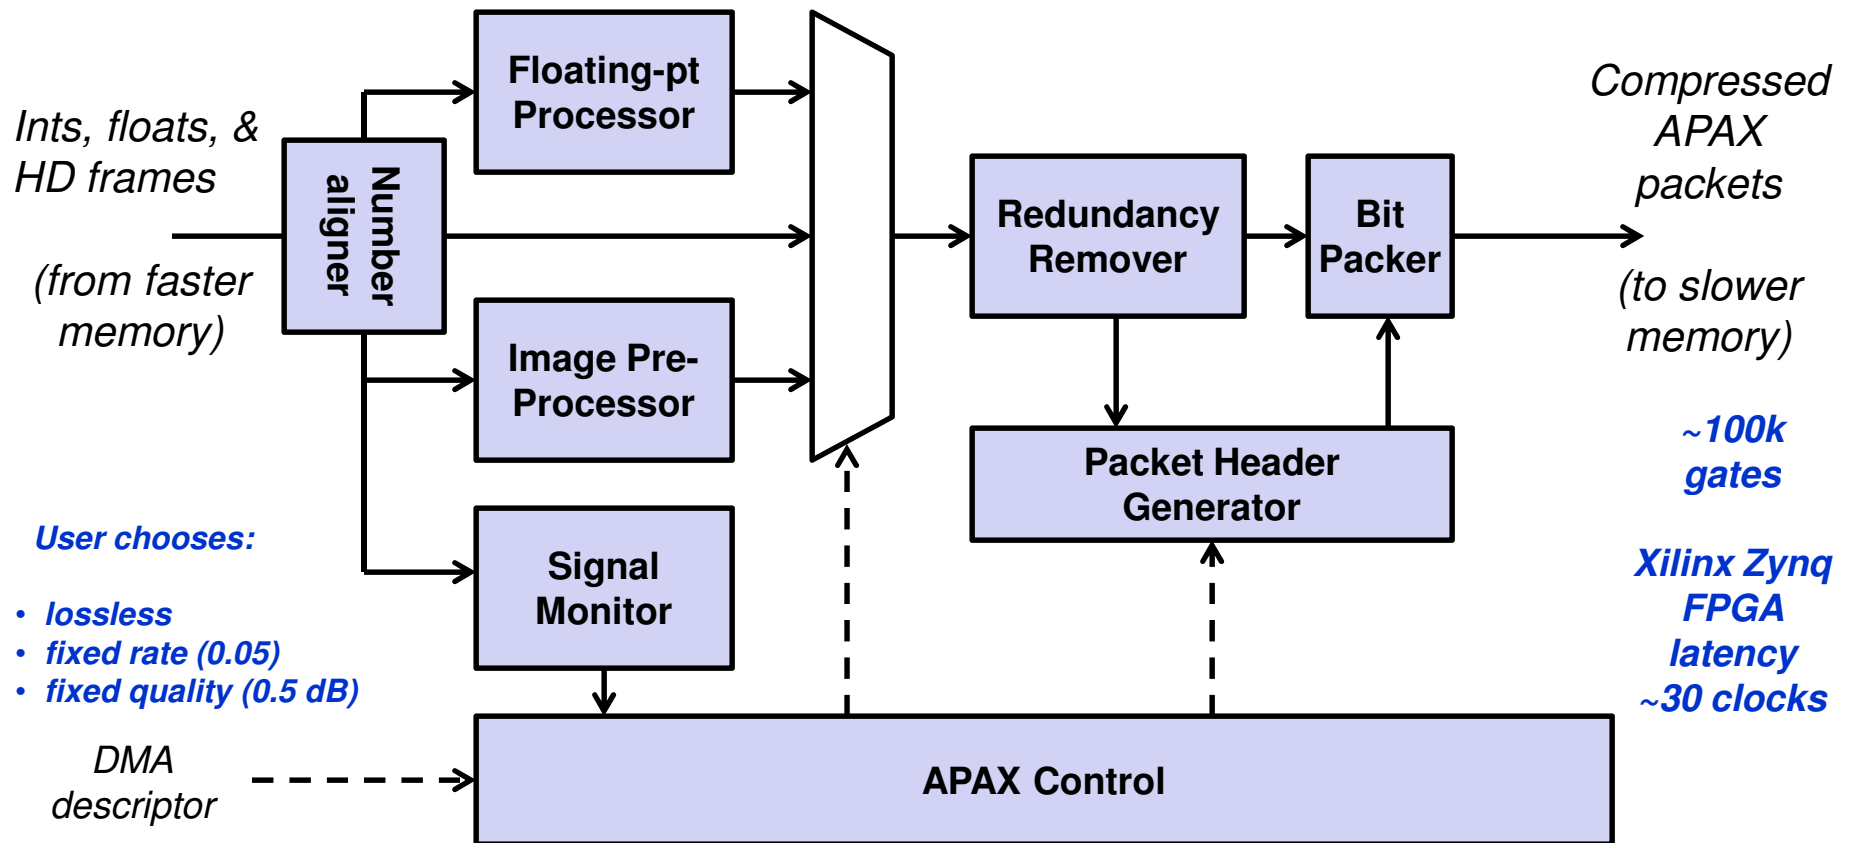
Financial Analysis



Genomics

GPU Application	Fixed Rate Comp Ratio
Hydrodynamics	3.3:1
Car crash simulation	1.5:1
Computational fluid dynamics	1.5:1
Financial models	1.45:1
Image recognition	4.2:1

APAX: a Numerical Compressor for I/O





Memory Wall Take-away:

Compression is easier than designing faster DDR

APAX Software Products

*60 Msps/sec
to 400 Msamp/sec
per x86 core*

Software Product	Details
APAX library (.dll, .so) 	APAX_profile() APAX_compress() APAX_decomp()
APAX MPI extensions	MPI_WRITE_APAX_C() MPI_READ_APAX_D()
APAX file system extensions	fwrite_apax_c() fread_apax_d()
APAX for CUDA & OpenCL	cudaDMAcustom()



To try APAX on your datasets...



Al Wegener

Samplify Systems
Santa Clara, CA

awegener@samplify.com
408-260-4000

Samplify Systems, Inc: Intellectual Property Solutions for Data Acquisition and Data Processing Bottlenecks

Samplify Systems, Inc. is a provider of intellectual property solutions which solve data acquisition and data processing bottlenecks. Samplify's Prism Compression addresses data acquisition bottlenecks in industrial, scientific, medical markets, while Samplify's APAX technology accelerates software applications which are performance-limited by I/O, storage, or memory bandwidth bottlenecks in the computing, consumer, and mobile device markets.

Samplify is a private company based in Silicon Valley, with backing from leading venture capitalists, Charles River Ventures and Formative Ventures and strategic partners Schlumberger, IDT, and Cosmo.

www.samplify.com



Samplify's APAX Technology Featured on EngineeringTV

From DesignWest 2012, Bill Wong interviews Samplify's Founder and CTO, Al Wegener, on the company's APAX technology for High Performance Computing, Big Data, and consumer electronics applications.

