LDAS Prototyping & Testing

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

Kent Blackburn & Stuart Anderson, LIGO Laboratory

- Data Distribution and Testing
  - Frame Compression
  - Client / Server Frame I/O using Sockets
  - High Speed Network Testing
  - Quick Look Analysis with PAW
- Single Point Benchmarks
  - FFT Benchmarks
  - 40 meter data flow *(GRASP)* Benchmarks
- Scalable Template Analysis Modeling
- Periodic Source Searches

Data Distribution and Testing

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Current of Frame I/O Library supports data compression:

GZIP Level	Differentiation	Translated Frame Size	Frame Size vs. Raw Data Size	Time (cpu) to Translate
None	No	1282532 KB	97.67%	975s (7.4%)
1	Yes	667693 KB	50.85%	1461s (75%)
1	No	726269 KB	55.31%	1494s (72%)
3	Yes	640549 KB	48.78%	1799s (78%)
3	No	706373 KB	53.80%	1863s (77%)
6	Yes	621157 KB	47.31%	3951s (91%)
6	No	697533 KB	53.12%	3187s (83%)
9	Yes	619965 KB	47.21%	4940s (91%)
9	No	696613 KB	53.05%	4401s (87%)

  - results from 200 MHz, single CPU Sun Ultra2 workstation
  - nearly 50% reduction using 10x CPU increase & 1.5x increase clock time
  - translates to direct savings in media!

Data Distribution and Testing

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Implemented Client-Server Frame I/O support to DAQS
  - » Design uses TCP/IP protocol over Unix Sockets (portable & flexible)
- Tunable parameters allow optimal performance
- Three hardware/OS configurations tested
  - » Between Sun workstations via 10baseT ethernet
    - – Peak performance 1 MB/sec
  - » Baja MIPS processor to Sun Sparc10 via 10baseT ethernet
    - – Peak performance 1.1 MB/sec
  - » Baja MIPS processor to Sun Ultra via 100baseT ethernet
    - – Peak performance 5.9 MB/sec (satisfies LIGO bandwidth needs!)

Data Distribution and Testing

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- High Speed Network Tests at CACR
  - ›› 100GB on Sun transferred over ATM to HPSS at bandwidth of 4.4MB/sec
- Bandwidth between protocols tested

	HIPPI (FP) Paragon-HPSS server	HIPPI (TCP/IP)	ATM (TCP/IP)
HIPPI (FP) Paragon-HPSS server	Mem-Mem: 50MB/sec Parallel Filesystem-Mem: 30MB/sec	-	-
HIPPI (TCP/IP)	-	HPSS Server-HPSS Server: 6MB/sec Paragon-HPSS Server: 2.4MB/sec	Sun-HPSS Server: 1MB/sec HPSS Server-Sun: 0.2MB/sec Sun-SP2: 2MB/sec
ATM (TCP/IP)	-	Sun-HPSS Server: 1MB/sec HPSS Server-Sun: 0.2MB/sec Sun-SP2: 2MB/sec	Sun-Sun: 9-14MB/sec SP2-SP2: 10-15MB/sec Sun-HPSS Server: 7-9MB/sec

Data Distribution and Testing

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Quick-Look of 40m data with Frame I/O and PAW *(n-tuples)*

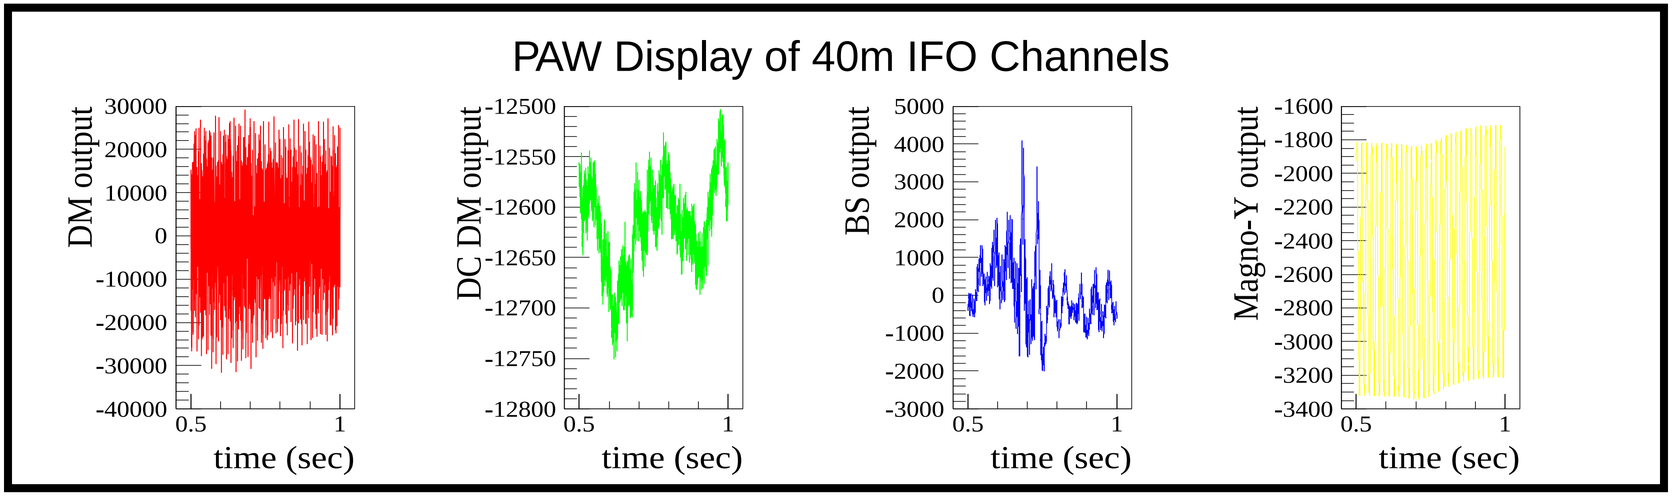

  - Multiple channels of data viewed simultaneously
  - Provides interactive (command line) interface
  - Future: Time series signal processing (FFTs, spectral cross correlations...)

Single “Point” Benchmarks

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Fast Fourier Transform Benchmarks
  - MFLOPS based on $5N\log_2 N$ floating point operations
  - Various processors tested
  - Random complex values used in simulated signal

		MFlops using 2^{16} complex points					MFlops using 2^{21} complex points				
CPU	MHz	NRv2	FFTpack	FFTW	ESSL	MLIB	NRv2	FFTpack	FFTW	ESSL	MLIB
Sun Ultra	*168*	*21*	*35*	*74*	-	-	*10*	*17*	*61*	-	-
RS6000	*66*	*12*	*22*	*52*	*108*	-	*10*	*12*	*45*	*105*	-
RS6000	*135*	*13*	*20*	*84*	*65*	-	*9*	*10*	*62*	*153*	-
PA8000	*180*	*104*	*74*	*160*	-	*227*	*5*	*6*	*23*	-	*89*
Pentium Pro	***200***	***15***	***34***	***54***	-	-	***13***	***29***	***51***	-	-

Single “Point” Benchmarks

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- 40m Data Flow *(GRASP)* Benchmarks
  - Binary inspiral searches using optimal (Weiner) filtering
  - Results representative of generic short template searches
  - Benchmarks from 5 different systems:
    - *LIGO’s Sun Ultra cluster; SGI’s Origin 2000; CACR’s IBM SP2, Paragon & Beowulf*
  - Demonstrates feasibility of implementing LIGO LDAS on-line systems

	Intel Paragon	SGI Origin 2000	Sun Ultra2 Workstation Cluster	**Beowulf PentiumPro Cluster**
Full LIGO (extrapolation)	*205*	*90*	*116*	***96***

Number of CPUs needed to keep up with LIGO data stream (extrapolation) using 1997 available systems
Beowulf is the name given to a cluster of high performance Pentium Pro PCs running Linux

Scalable Template Analysis Modeling

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Significant component of on-line detection algorithms based on optimal *(Weiner)* filtering *(linear filtering)*
  - other filtering techniques: model independent - wavelets, adaptive,...
- Signal waveforms *(templates)* accurately known
- Uses FFT *(inverse)* to convolve data with templates
- Physics provides foundation for two classes of waveforms:
  - Inspiral of binary system of neutron stars and black holes *(~10^4 templates)*
  - Quasi-normal mode ringdown of excited Kerr black holes *(~10^2 templates)*
    - Instrumental signitures also modeled by damped oscillators *(small class of these)*
- Analysis of data using templates *embarrassingly* parallel**!**

LDAS Modeling - Data Flow

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

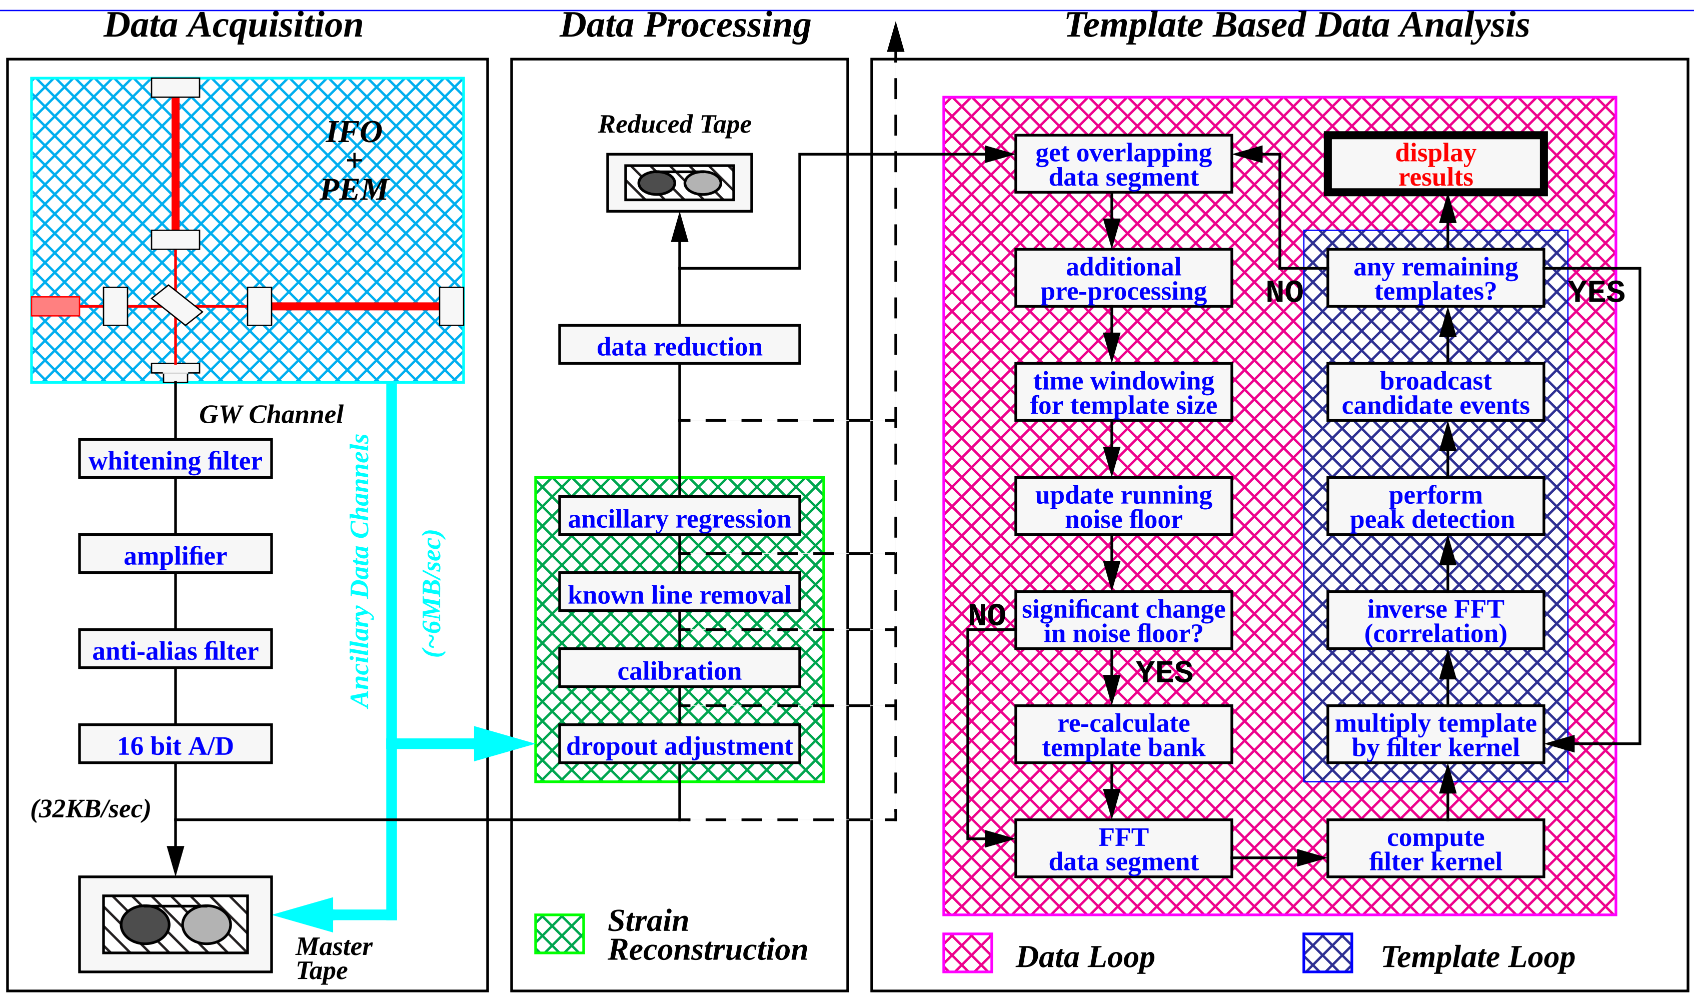

LIGO-G970261-00-E

LDAS Modeling - Excel Spreadsheet

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- LDAS data flow diagram cast into Excel Spreadsheet
- Spreadsheet model include macro to micro details:
  - IFO data rates, algorithms, number nodes, ram, storage,...
  - clock cycles for floating point operations, memory copies, I/O,...
- Broken down into: source parameters, data conditioning, inspiral templates, ringdown templates & costs estimates
- Used to scope out performance and cost estimates for 3 different classes of computer systems: supercomputers, clusters of unix workstations, clusters of PCs *(Beowulf)*
- Model found to be in line with 40m data flow benchmarks!

LDAS Modeling - Data Conditioning

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Data Conditioning includes:
  - ›› calibration to best estimate of strain
    - – linear regression of ancillary signals
    - – narrow line removal
    - – data dropout corrections
  - ›› data reduction *(bandwidth & simplification)*
- Much of the data conditioning carried out in frequency domain (FFT's and linear algebra involved)
- Equivalent to analysis associated with IFO diagnostics
- Can be carried out on single high end workstation!

LDAS Modeling - Data Conditioning

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

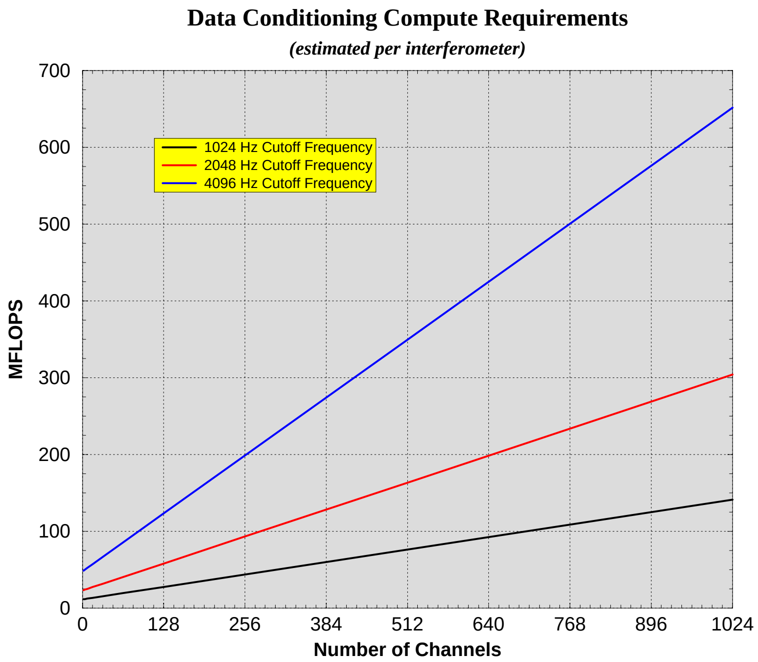

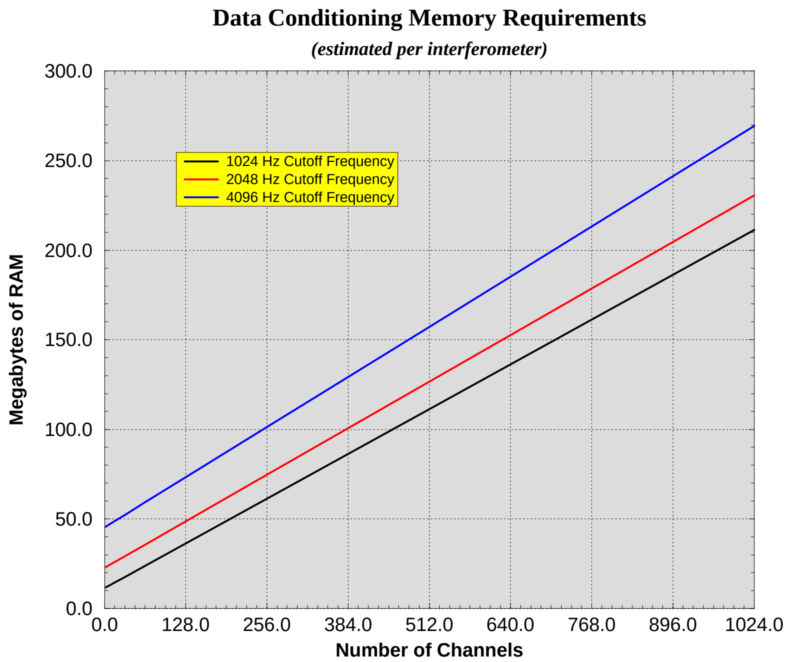

LDAS Modeling - Binary Inspiral

- Number of Templates: $\sim 6.8 \times 10^5$ (0.1 / Loss) $(0.2 / M_{solar})^{(5/3)}$ $(145\text{Hz} / f_{best})^{(8/3)}$
  - » Parameterized by
    - – loss rate; desired fractional loss of events L
    - – depth of search; minimum mass in binary system M
    - – shape of instrument sensitivity; frequency of lowest noise f_{best}
- Long templates: for one solar mass this is ~90 second
- Frequency content of signal weighted up to ~1024Hz
  - » reducing bandwidth reduces computation as $\sim N \log_2 N$
- Searches down to one solar mass achievable with ~20 GFLOPS making on-line analysis possible!

LDAS Modeling - Binary Inspiral

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

Binary Inspiral Templates

(estimated per interferometer)

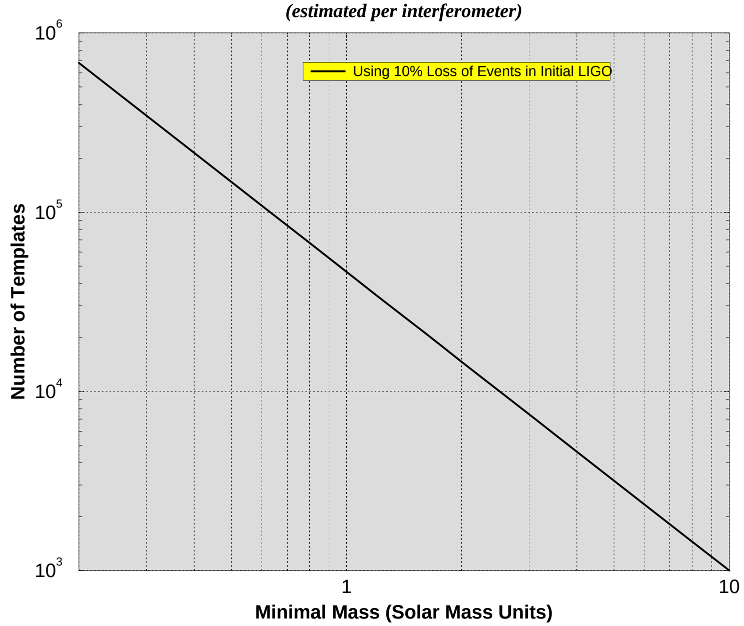

Binary Inspiral Template Data Bank Size

(estimated per interferometer with 8x overlap)

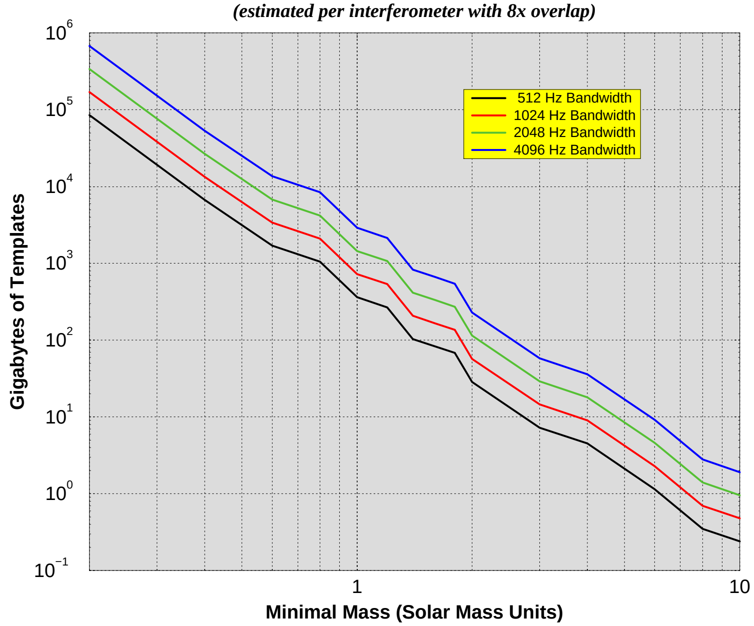

LDAS Modeling - Binary Inspiral

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

Binary Inspiral Template Compute Requirements

(estimated per interferometer with 8x overlap)

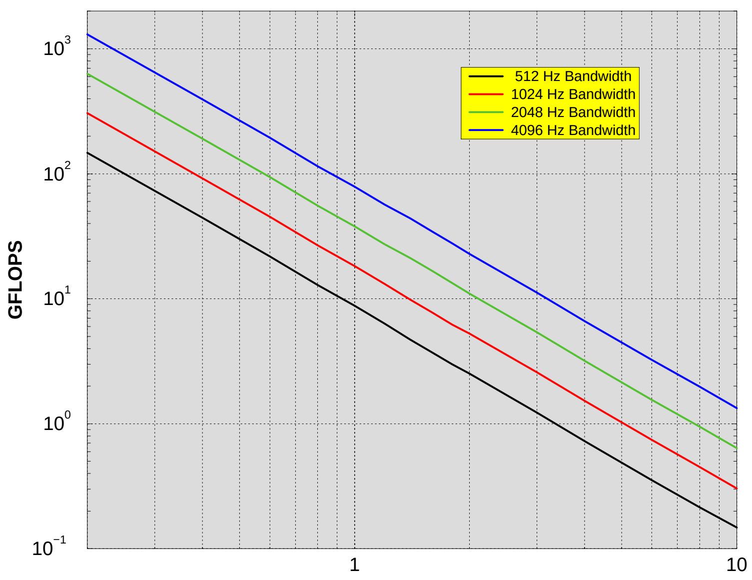

Binary Inspiral Template Memory Size

(estimated per node with 8x overlap)

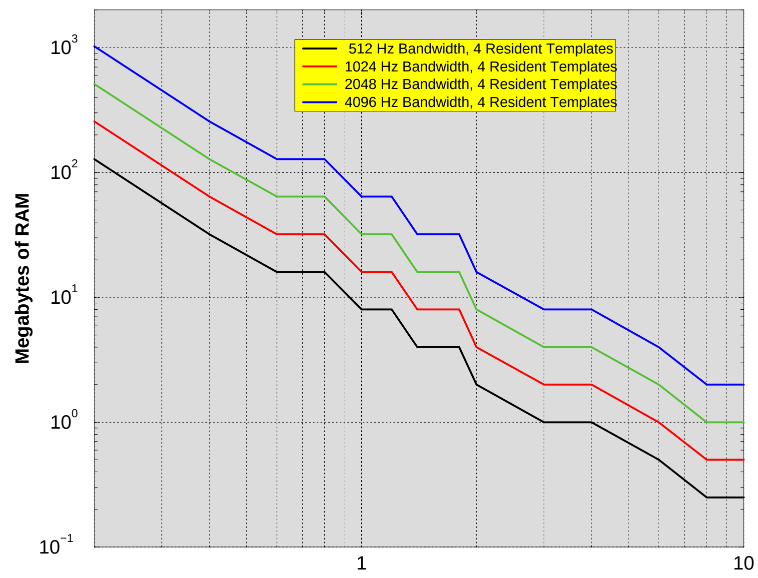

LDAS Modeling - Ringdown

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Number of Templates: ~2700 (0.1 / Loss)(Q/100)(1+(log(f_{max}/10^3)-log(f_{min}/10^2))/2)
  - ›› Parameterized by
    - – loss rate; desired fractional loss of events L
    - – quality factor Q *(astrophysically < ~20; instrumentally much higher)*
    - – maximum frequency in search f_{max} *(shot noise limits this to few kHz)*
    - – minimum frequency in search f_{min} *(seismic wall in noise sets this at ~30Hz)*
- Instrumental signatures will strongly mimic this waveform
  - ›› useful for characterizing instruments non-Gaussian noise
  - ›› binary inspiral of several tens of solar mass system followed by black hole ringdown *(precursor)* in LIGO band
- Requires roughly 1/100th the compute performance(cost)!

LDAS Modeling - Ringdown

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

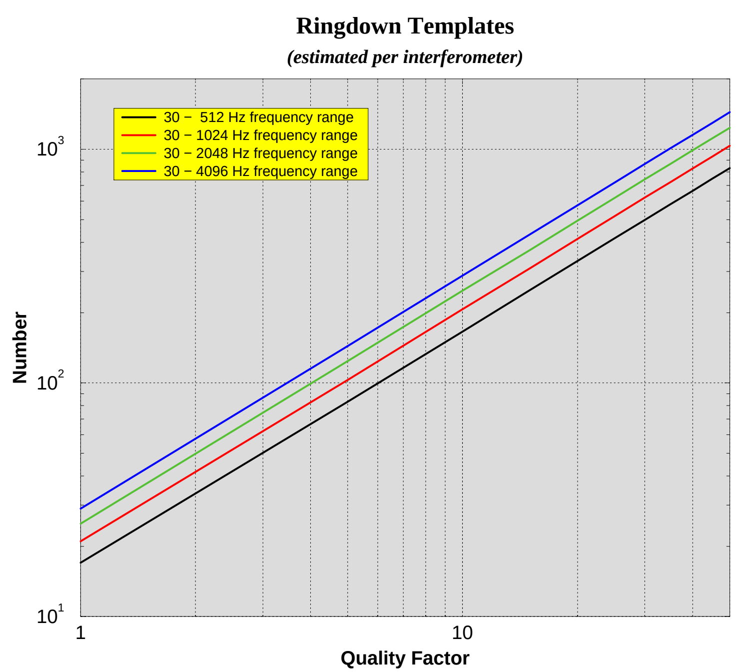

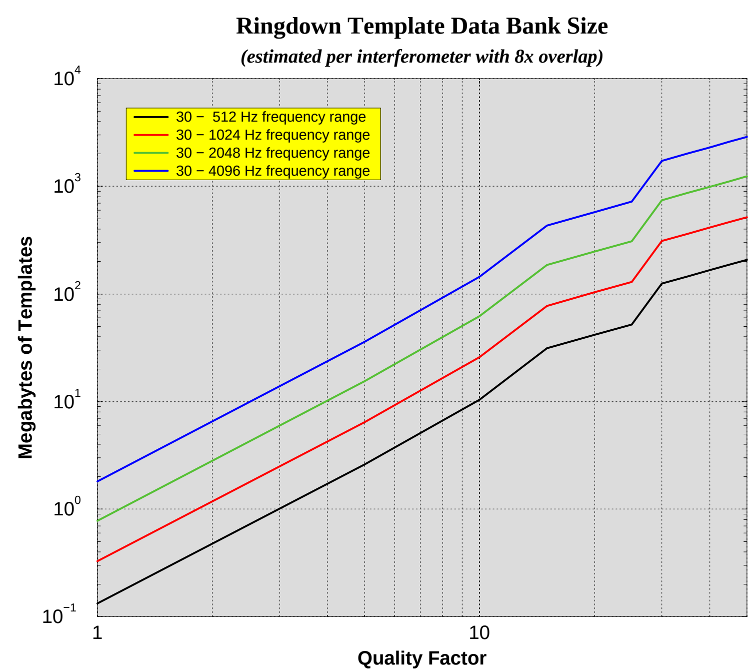

LDAS Modeling - Ringdown

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

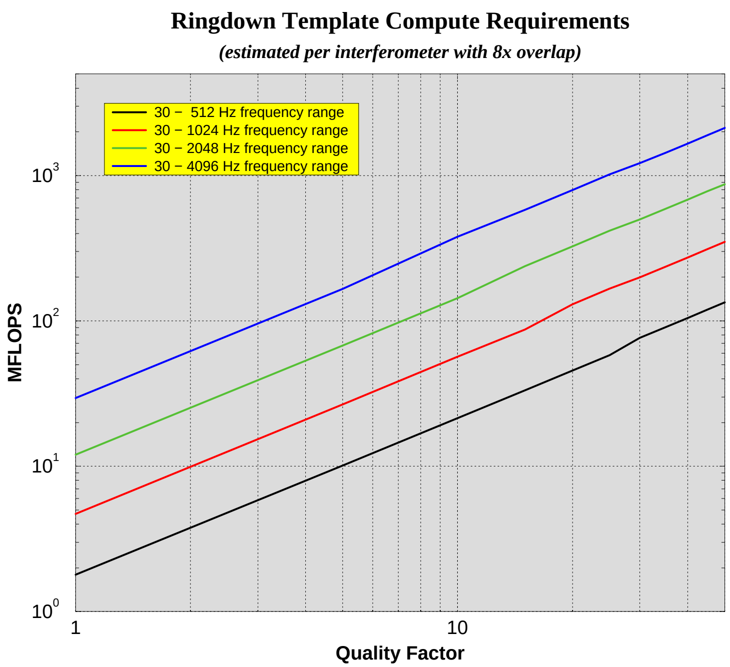

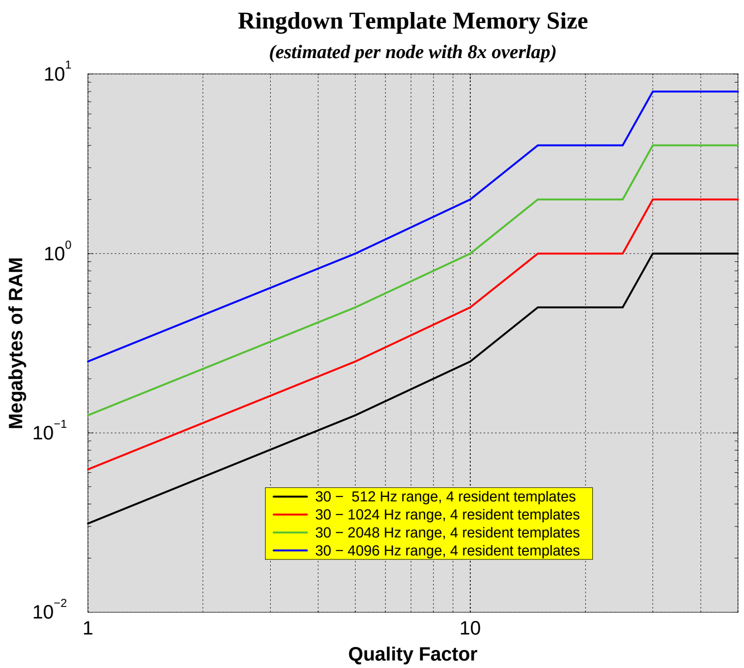

LDAS Modeling - Cost Estimates

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Cost estimates made for 3 distinct classes of computing
  - ›› Cost/performance for LDAS analysis on these systems spanned roughly one order of magnitude
    - Supercomputers *(fastest processors networks & memory)*
    - Clusters of Unix workstations connected by fast network
    - Beowulf *(Cluster of fast PCs using free software and fast network)*
  - ›› consistent with 40 meter data flow *(GRASP)* benchmarks
- All 3 classes support ANSI compilers, MPI and POSIX
  - ›› Same code can be run on all three systems
- Beowulf most cost effective system today! *(~$2000/node)*

LDAS Modeling - Supercomputers

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

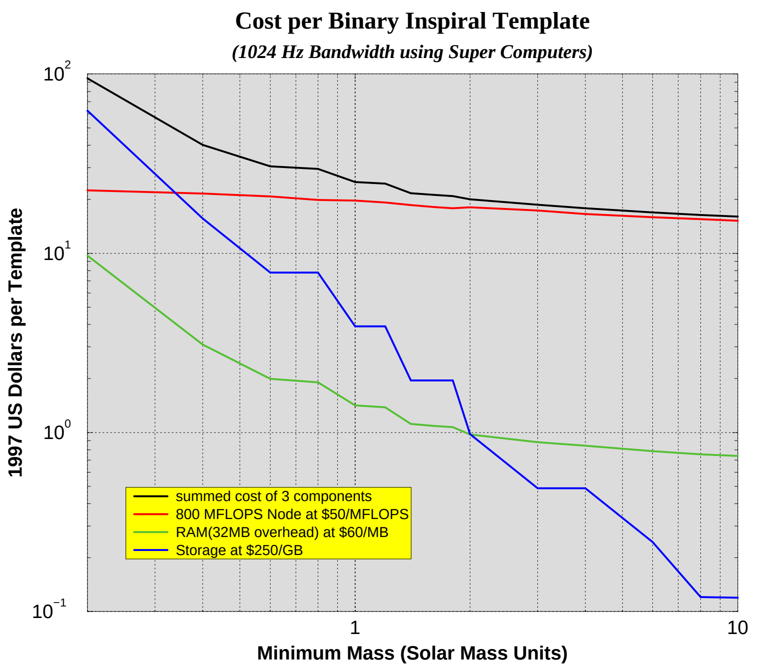

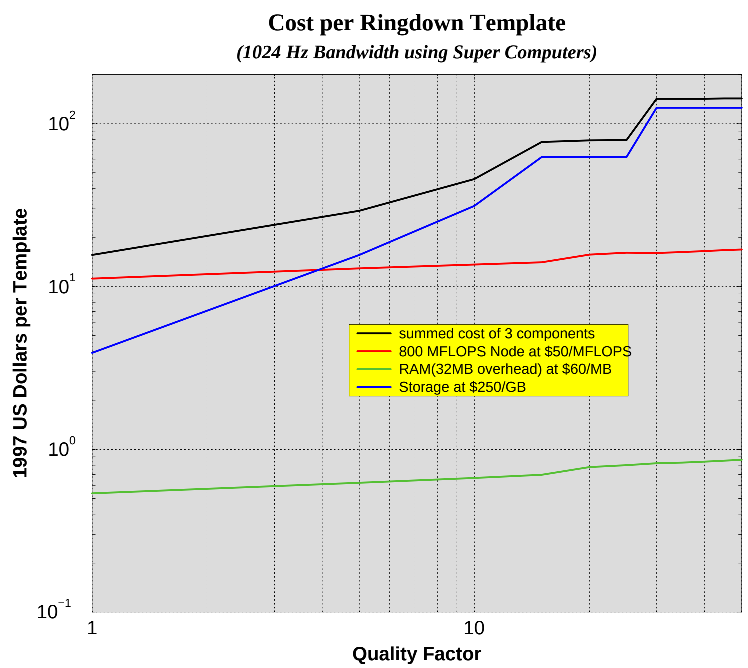

LDAS Modeling - Workstations

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

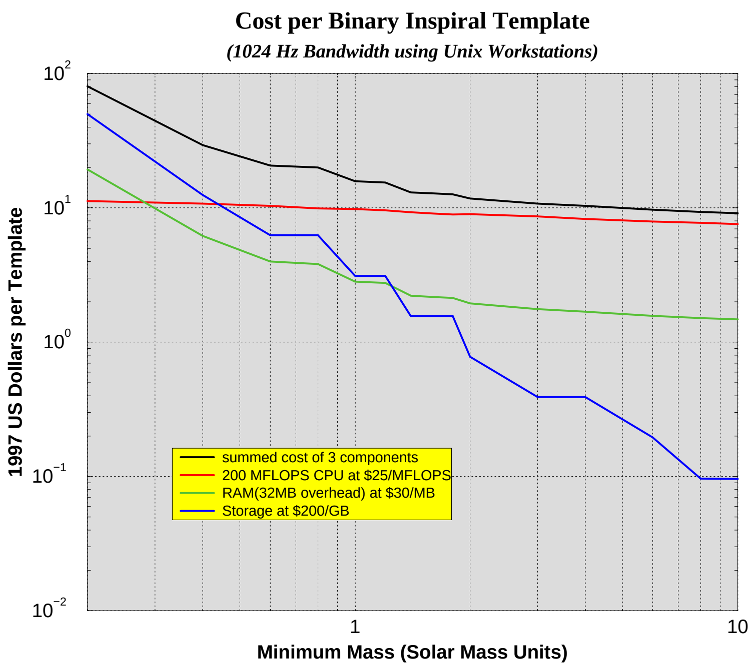

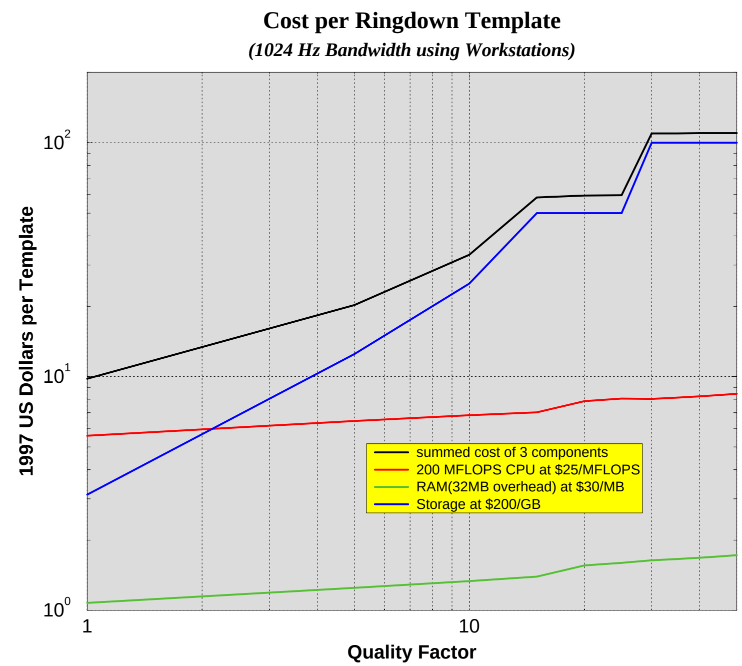

LDAS Modeling - PCs

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

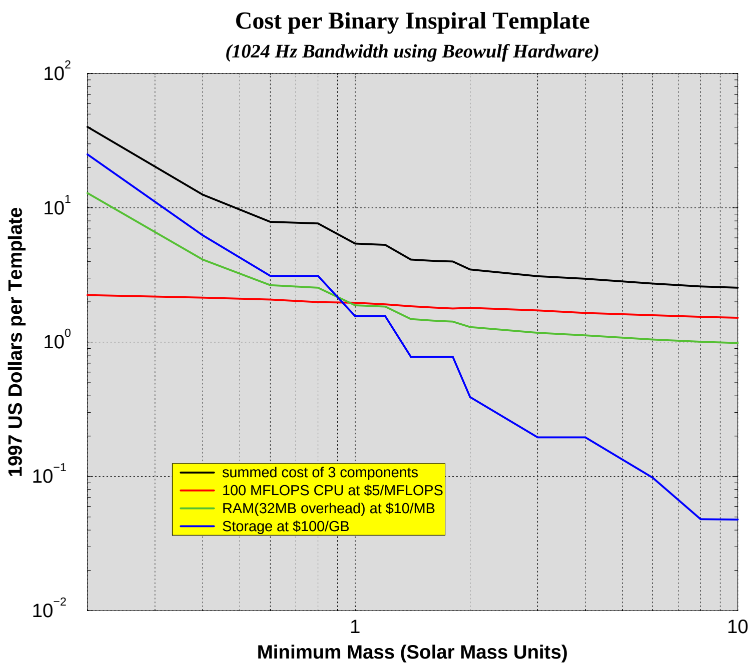

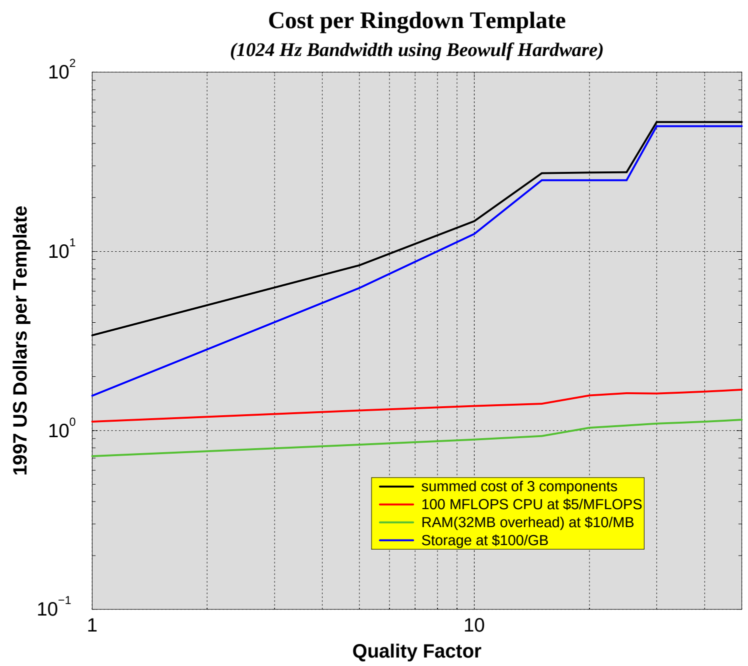

LDAS Modeling - Comparison

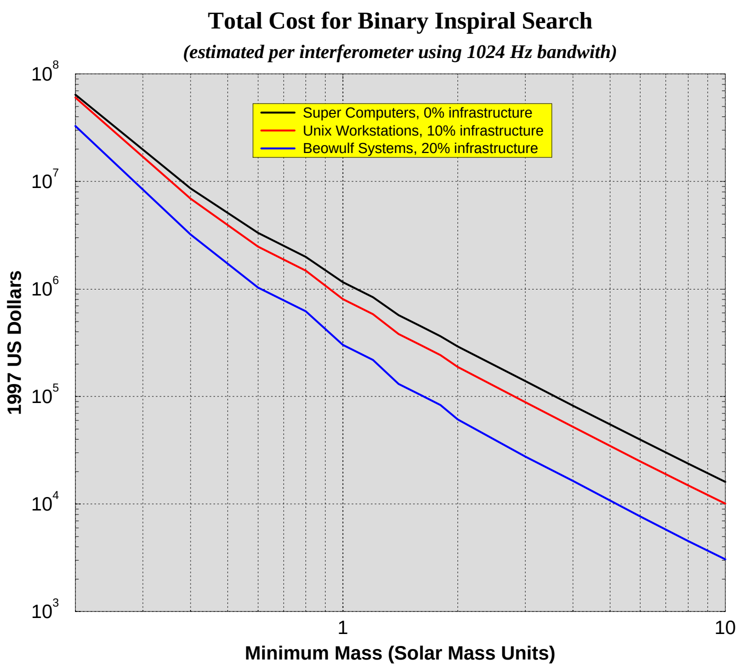

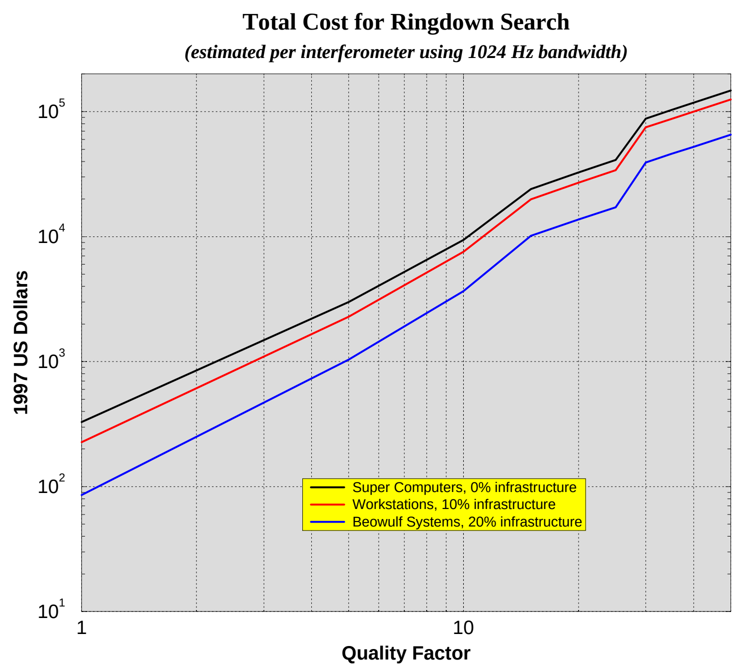

Periodic Source Searches

PAC3 Meeting - Hanford WA, Nov. 6-7, 1997

- Modeling effort underway to scale search algorithms

›› **Periodic search represents different class of problem *(>TFLOPS)***

›› **Communication fabric is heavily stressed by giga-point FFTs**

›› **Best sensitivity found with 1024 stacking: *(~ weeks of data)***

›› ***NOT* an on-line search, however**

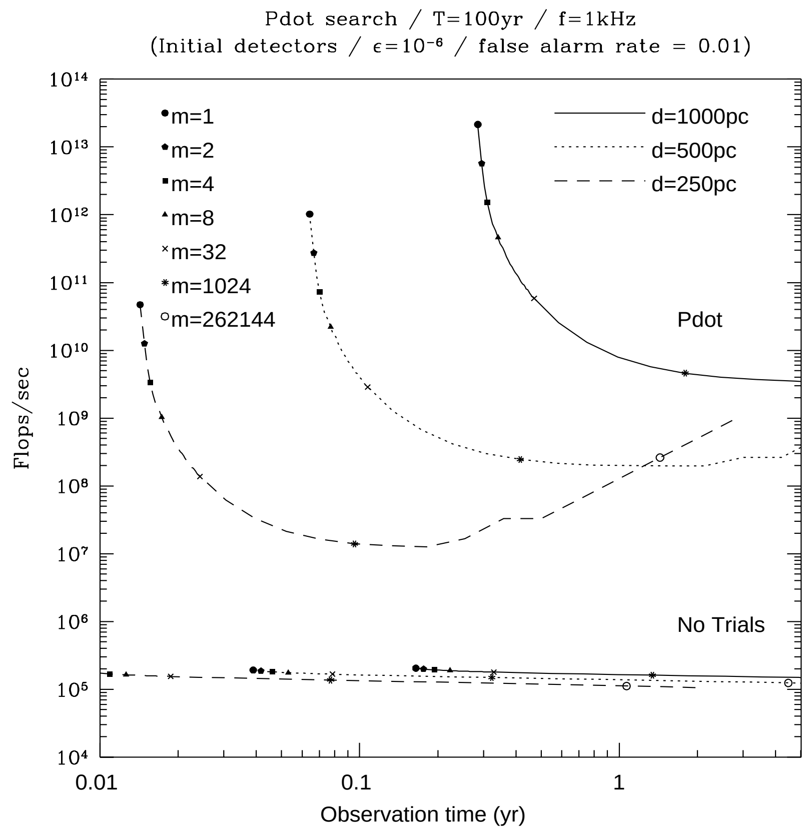

›› **Incoherent stacking of data optimizes computation for given sensitivity**

›› **All curves involve search over period; top 3 include search for first derivative of period in time**