

# **Energy Usage for Parallel Haskell Programs**

### **Kevin Hammond**

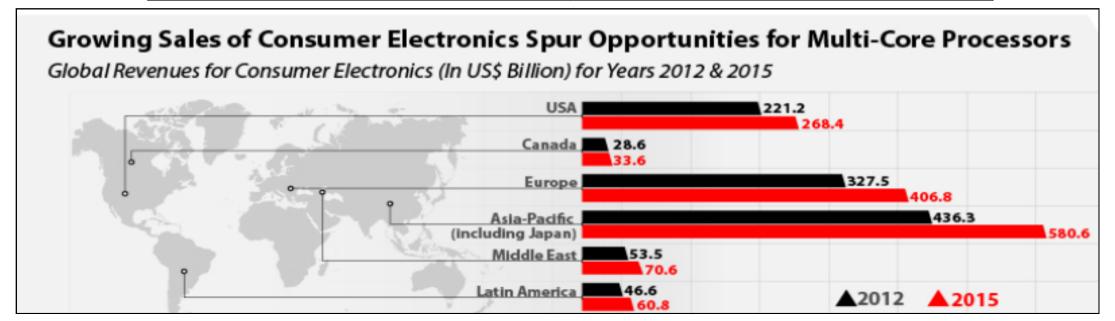
**Hylomorph Solutions** 

WG2.11, Salem, May 2nd, 2019

Work undertaken by:
Yasir Alguwaifli
University of St Andrews
In: Proc. Lambda Days 2019
Submitted to CC&PE

### Multicore Momentum

Cloud Computing	IoT					
Performance	Embedded Systems					
Multimedia						





# From Multicore to Manycore



Intel Xeon Phi 7290: 72 x86 cores (at 1.5GHz), 245W Aimed at supercomputing applications



# It's not just about large systems

- Even mobile phones are multicore
  - Samsung Exynos 5 Octa has 8 cores, 4 of which are Dark
- Performance/energy tradeoffs mean systems will be increasingly parallel
- Even embedded systems are becoming multicore and heterogeneous
  - NVidia Tegra TX2 has 6 ARM CPU cores and 256core Pascal GPU
- If we don't solve the multicore/manycore challenge, then no other computing advances will matter!



ALL Future Programming will be Parallel!

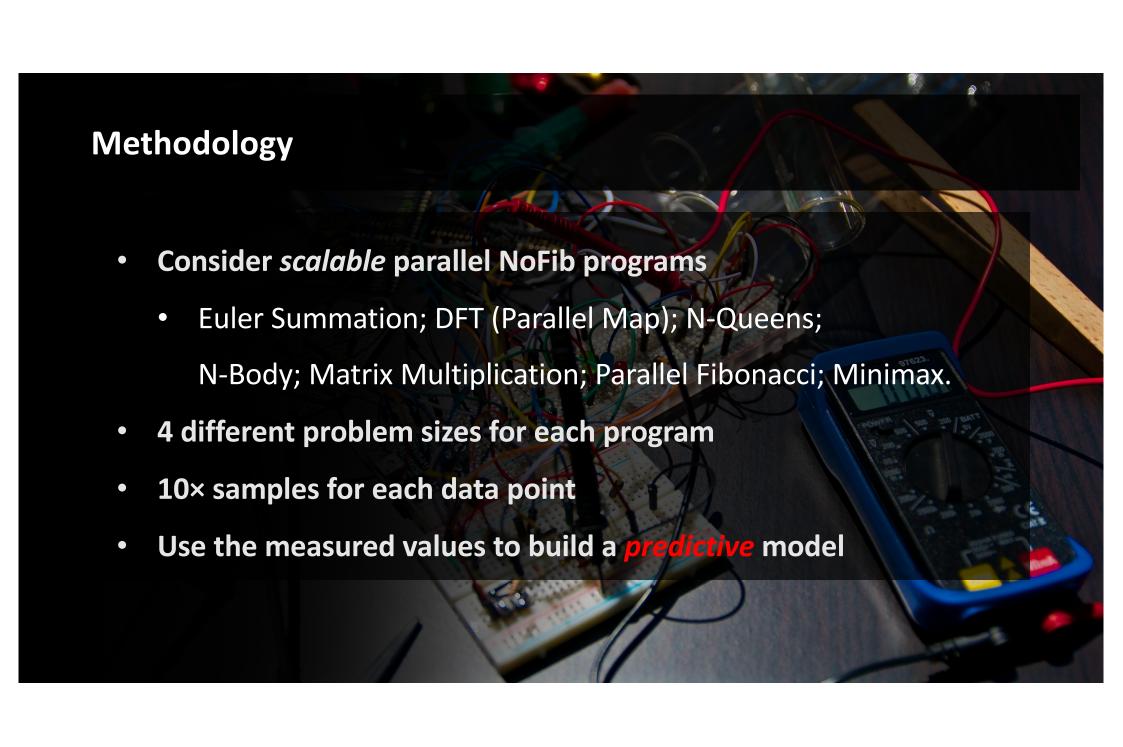


### Main Research Questions

Q1: Can we relate energy usage to easily obtained metrics?

Q2: Is it possible to build good energy models for parallel Haskell programs?

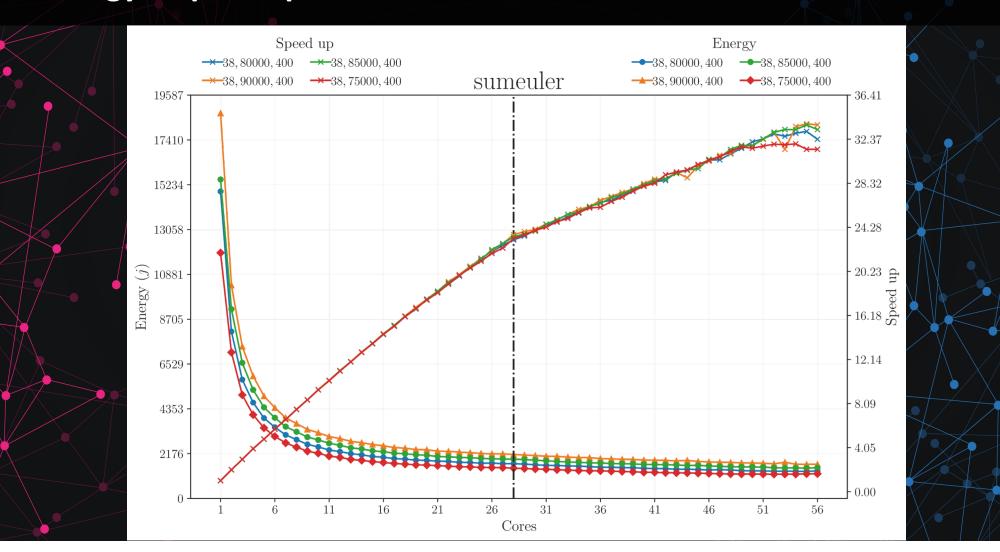




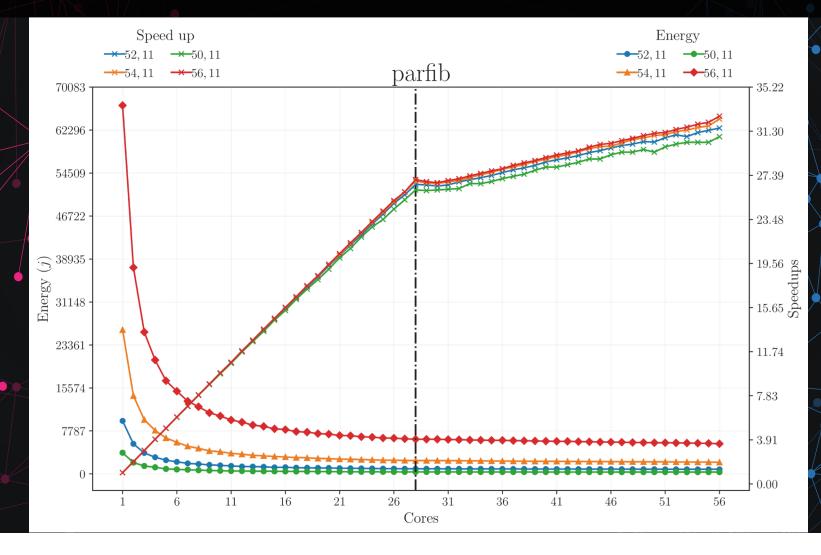
### **Experimental Setup**

- Intel's Running Average Power Limit (RAPL) for energy sampling
- GHC 7.6.3 with -O2 & -feager-blackholing
- Red Hat 4.8.5-11 with kernel 3.10.0-514.21.2.el7.x86\_64
- 2×Intel Xeon E5-2690 @ 2.601 GHz, 28 Physical Cores
- Fixed Clock Frequency (DVFS disabled)

### **Energy v Speedup: Euler Summation**



# Energy v Speedup: Parallel Fibonacci



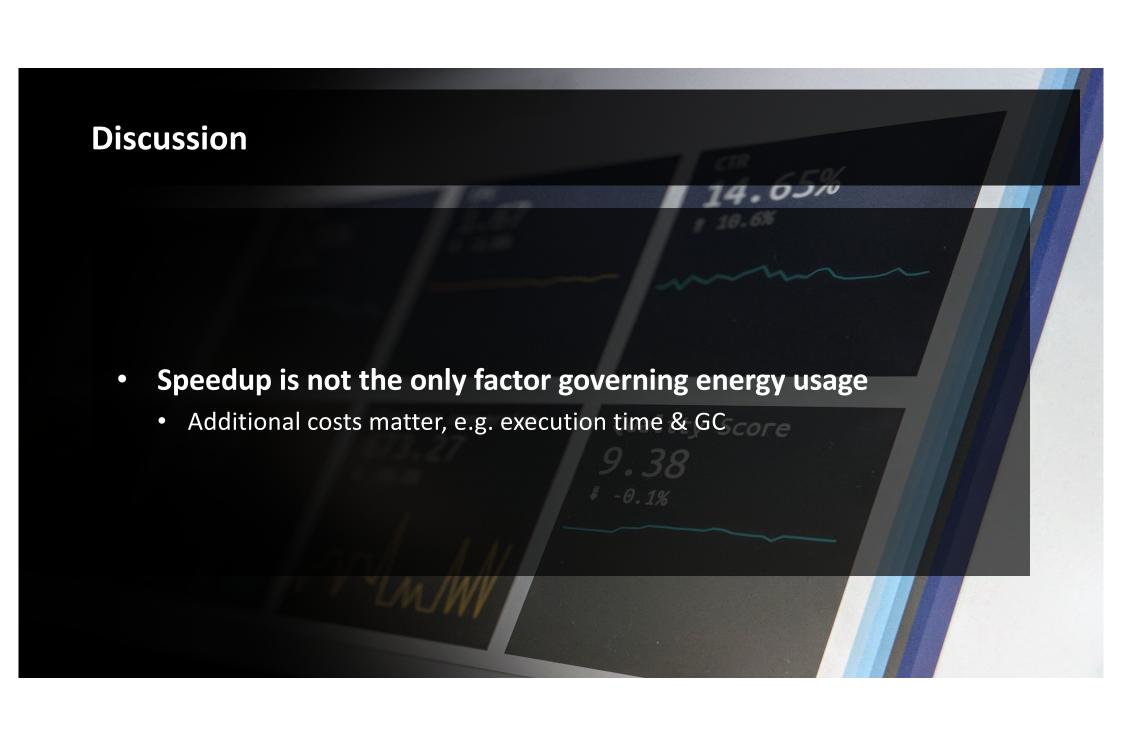
### **Energy v Speedup: N-Body** Speed up Energy nbody **→**50000 <del>~~</del>150000 <del>×</del>50000 13167 35.4611704 31.52 10241 27.58 8778 23.64 Energy (j)731519.70 g - 15.76 S 4389 11.82 2926 7.88 1463 3.94 0.00 26 21 31 51 56 11 Cores



 Energy consumption is strongly correlated with speedup (i.e. execution time)

- Hyper-threading has a significant effect
  - There is such a thing as a "free lunch"

### **Energy v Speedup: Minimax** Speed up Energy $-\times$ 4,10 $-\times$ 4,12 minimax <del>-×</del>4,6 25880 10.53 23004 9.36 20129 8.19 172537.02 Energy (1) 14378 (1) 11502 4.68 8627 3.51 5751 2.34 2876 1.170.00 26 21 31 51 56 11 Cores



### **Constructing an Energy Model**

- Apply Multi-variate regression
  - Linear model:  $Y_i = \beta_0 + \beta_1 x_{i1} + ... + \beta_n x_{in}$
- Following features considered:
  - Productivity (Roughly = Speedup)
  - Garbage Collection (GC) Time
  - Execution ("Mutation") Time
  - Average Power Usage

### **Constructing an Energy Model**

- Apply Multi-variate regression
  - Linear model:  $Y_i = \beta_0 + \beta_1 x_{i1} + ... + \beta_n x_{in}$
- Following features considered:
  - Productivity (Roughly = Speedup)
  - Garbage Collection (GC) Time
  - Execution ("Mutation") Time
  - Average Power Usage
- $E = m c^2$

### **Constructing an Energy Model**

- Apply Multi-variate regression
  - Linear model:  $Y_i = \beta_0 + \beta_1 x_{i1} + ... + \beta_n x_{in}$
- Following features considered:
  - Productivity (Parallel Utilisation)
  - Garbage Collection (GC) Time
  - Execution ("Mutation") Time
  - Average Power Usage
- Energy =  $\beta_0 + \beta_1$  Prod +  $\beta_2$  GC +  $\beta_3$  Mut +  $\beta_4$  Power

# Application-Specific Energy Model

Program	Parfib	Sumeuler	Nbody	DFT	Queens	MatMult	Minimax	
Coefficient								
$\hat{\beta}_0$ (Intercept)	-1.953e+03	333.2	-3.966e+02	567.54	-287.7	-287.7	152.59	
$\hat{\beta}_1$ (productivity)	-4.944e-01	-0.30394	-1.232e-01	0.111	-0.347	-0.347	570.06	
$\hat{\beta}_2$ (Mutation)	-1.983e+03	13.57	-3.128e+02	448.08	-85.84	-85.8	666.77	
$\hat{\beta}_2$ (GC)	285.09	7.378e+03	-7.251e+02	552.3	11.43	11.43	0.094	
$\hat{\beta}_3$ (power)	4.749e+01	11.84	5.016	-10.14	11.08	-0.34	-1.833	

**Table 3.**  $\hat{\beta}$  coefficients for the complete dataset

High significance in almost all cases (p < 0.01) for Productivity and Power Mutation and GC are also significant in many cases (p < 0.05) Only minimax showed low significance (p > 0.1)

### How well do the Models fit the Results?

Par	rfib	Sumeuler		NBody		DI	DFT		Queens		Matmult		Minimax	
$Adj. R^2$	RMSE	$Adj.R^2$	RMSE	Adj. R <sup>2</sup>	RMSE	Adj. R <sup>2</sup>	RMSE	Adj. R <sup>2</sup>	RMSE	$Adj. R^2$	RMSE	$Adj. R^2$	RMSE	
97.6%	186.75	98.7%	243.41	98.8%	653.26	98.1%	251.96	98.5%	281.48	97.4%	189.65	99.2%	349.71	

**Table 1.** Adjusted  $R^2$  and root mean square error (RMSE) values for the complete dataset

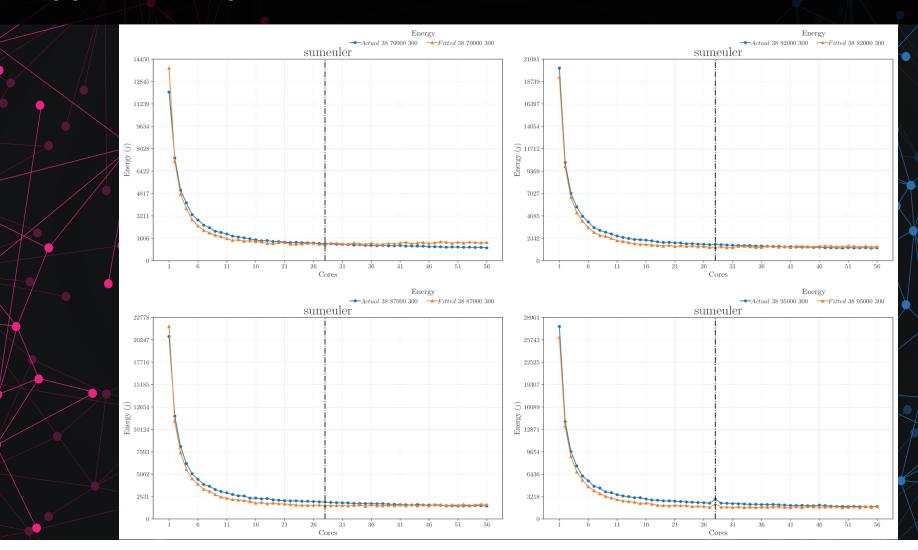
R<sup>2</sup> values show good fits between models and predictions in all cases (> 97%)



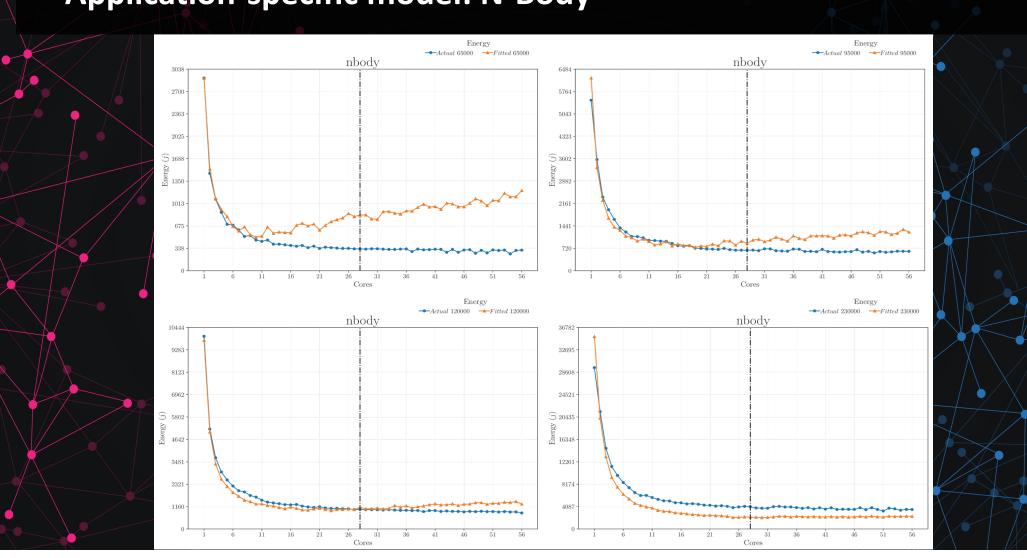


- 4 new problem sizes compared with predictions from models
  - Two smaller and two larger data sets considered

### **Application-specific model: Euler Summation**

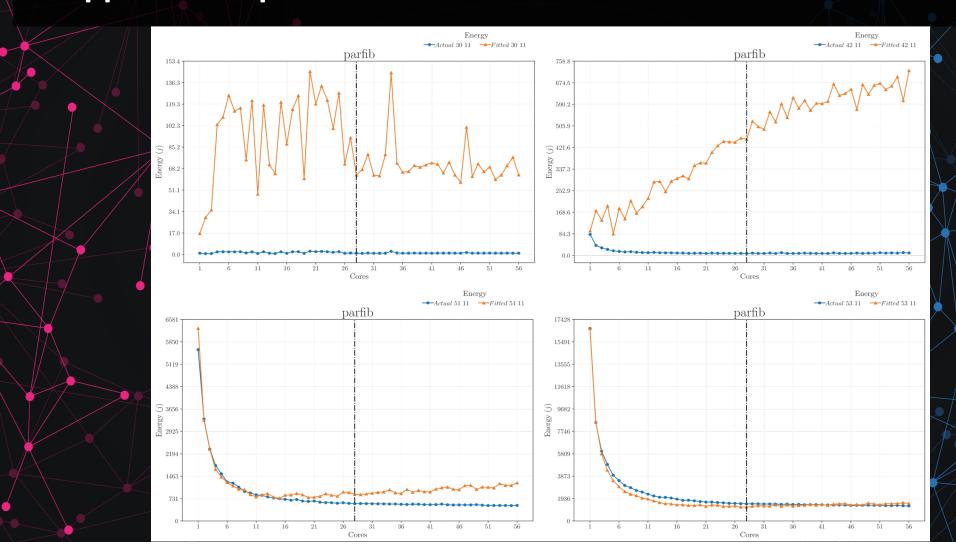


### **Application-specific model: N-Body**

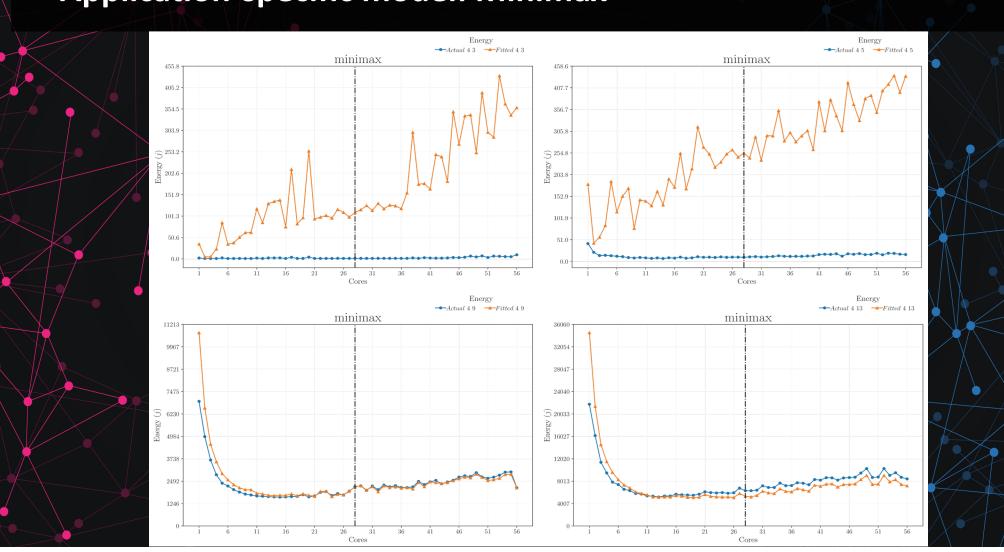


# **Application-specific model: DFT** → Actual 2300 2 → Fitted 2300 2 → Actual 2950 2 → Fitted 2950 2 (i) 742 -594 -Energy Energy 4459 -© 13497 1486 -2699

### **Application-specific model: Parallel Fibonacci**



## **Application-specific model: Minimax**





- Application-specific Models are a Good fit at larger sizes
  - Even to previously unseen data sets
- There is often a Poor Fit at smaller sizes
  - More likely to be affected by small variations?
  - Less consistent energy usage?

## New Research Question

Q3: Can we construct a "model of models"?



# A Generalised Energy Model

```
Residuals:
 Min 10 Median 30 Max
-12291.8 -400.0 -167.2 291.7 7129.2
Coefficients:
             Estimate Std. Error t value Pr(>|t|)
(Intercept) -188.8143 28.3854 -6.652 2.99e-11 ***
               4.9193 0.5146 9.559 < 2e-16 ***
power
mutation -1580.8171 67.9912 -23.250 < 2e-16 ***
    -1544.8994 68.0120 -22.715 < 2e-16 ***
gc
productivity 0.1210 0.0086 14.071 < 2e-16 ***
total time 1615.2182 67.9903 23.757 < 2e-16 ***
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
Residual standard error: 819 on 15681 degrees of freedom
Multiple R-squared: 0.9398, Adjusted R-squared: 0.9397
F-statistic: 4.892e+04 on 5 and 15681 DF, p-value: < 2.2e-16
```

### **General model: DFT** → Actual 2300 2 → Fitted 2300 2 → Actual 2950 2 → Fitted 2950 2 dft 1103 -S 1506 473 -904 Energy Energy → Actual 6150 2 → Fitted 6150 2 21501 4277 16126 Energy (2) 13438 (f) 3564 -2851 -1426 -5375 2688

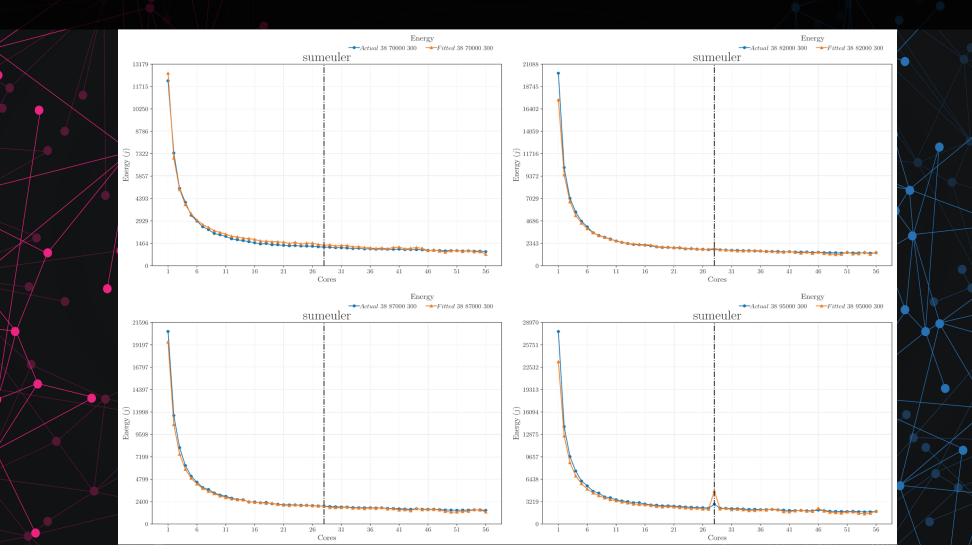
### General model: Matrix Multiplication → Actual 500 3 10 → Fitted 500 3 10 → Actual 1000 3 10 → Fitted 1000 3 10 matmult matmult 352.9 132.6 308.8 113.7 264.7 44.1 Energy Energy matmult matmult 696.1 2953 596.7 (j) 497.2 397.8 1688 1266

### **General model: Parallel Fibonacci** → Actual 30 11 → Fitted 30 11 → Actual 42 11 → Fitted 42 11 parfib 61.0 30.5 Energy Energy → Actual 51 11 → Fitted 51 11 parfib parfib 5213 15499 3910 11624 Energy (2) 8984 (3) (4) (5) 8984 (4) (5) 8984 (5 Energy 2606 -1303 -3875

### **General model: N-Body** → Actual 65000 → Fitted 65000 nbody nbody 2365 © 1689 S 3500 Energy 1351 1014 1400 -Energy Energy nbody nbody 30718 6964 23038 (f) 5804 (f) 6804 (f) 4643 S 19199 Energy 15359 3482 11519 2321 7679

### **General model: N-Queens** Energy Energy →Actual 11 →Fitted 11 →Actual 12 →Fitted 12 queens queens 111.283.1 97.3 71.2 83.4 (j) 59.3 47.5 © 69.5 Energy 55.6 23.7 27.8 11.9 13.9 Energy Energy -Actual 13 -Fitted 13 →Actual 17 →Fitted 17 queens 168727 66.7 126545 Energy (j) 44.5 Euergy (j.) 105454 - 84363 - 84363 63273 42182 21091

### **General model: Euler Summation**

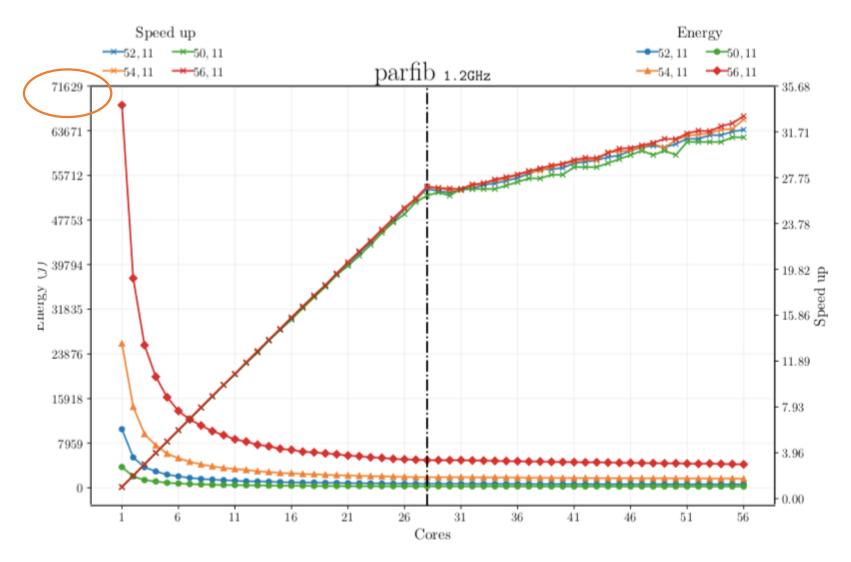


#### **General model: Minimax** Energy Energy → Actual 4 3 → Fitted 4 3 →Actual 4 5 →Fitted 4 5 minimax minimax 69.44 60.7648.69 52.0817.36 8.11 8.68 0.00 26 Cores Energy Energy →Actual 4 9 →Fitted 4 9 → Actual 4 13 → Fitted 4 13 minimax minimax 7113 -22571 62245335 -16928 Euergy (j) 4446 - 3557 -Euergy (j.) 14107 -8464 2668 5643 2821

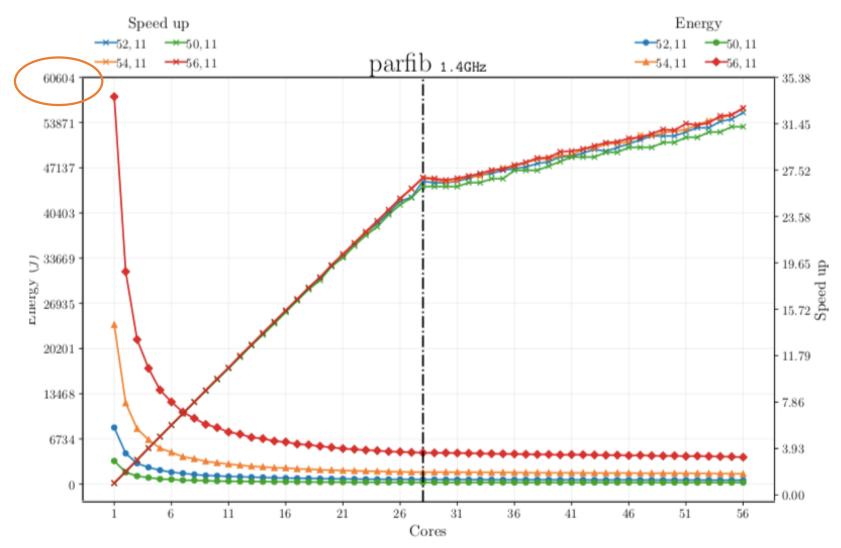
### New Research Question

### Q4: How does clock frequency affect energy usage?

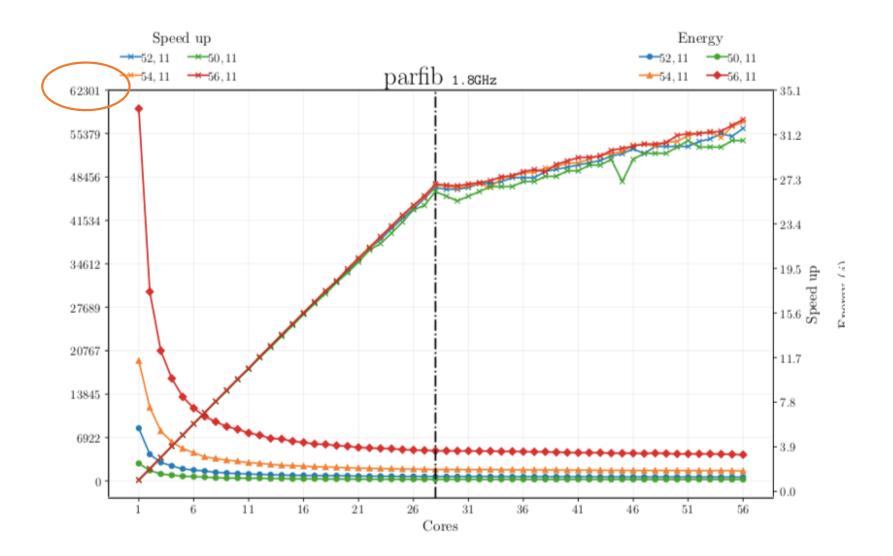




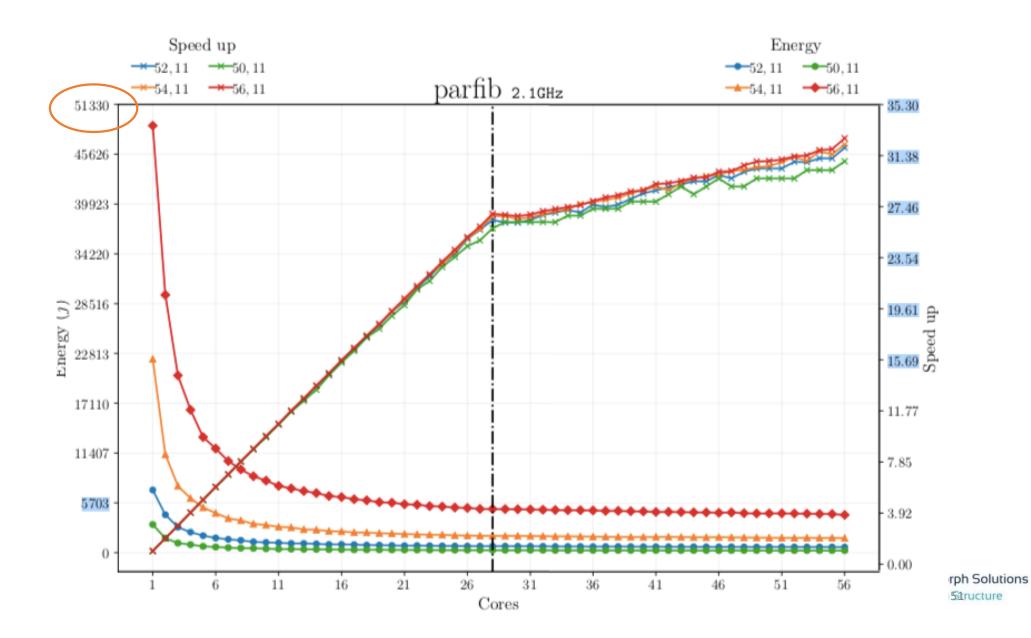


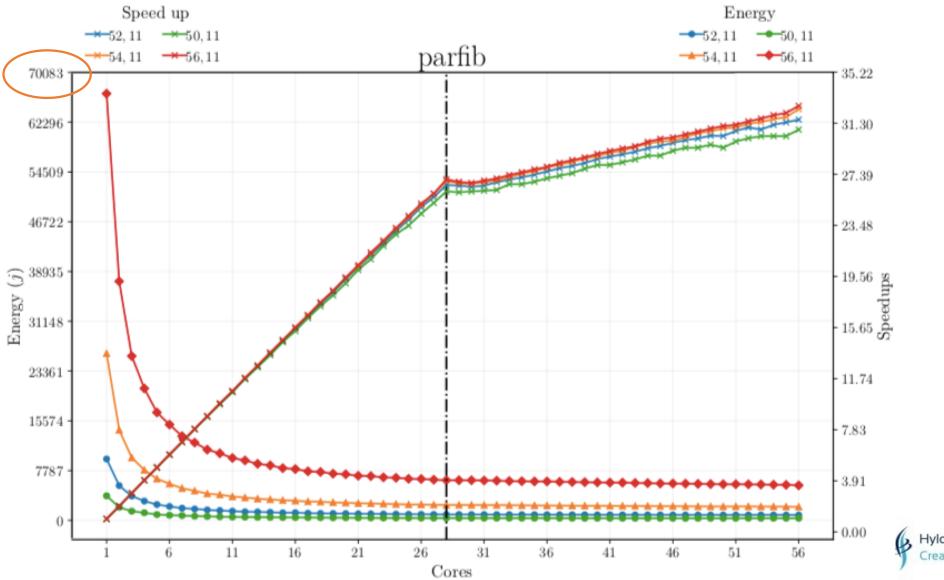




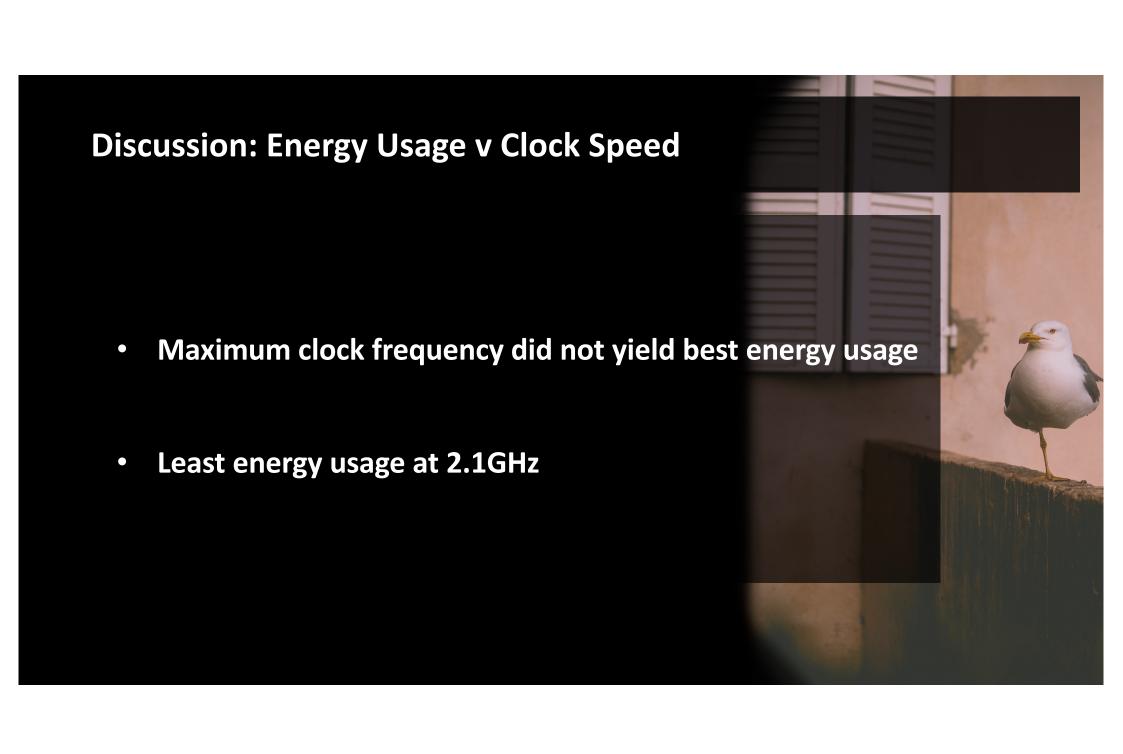












### **Conclusions**

- Runtime energy usage is predictable for Parallel Haskell
- General close correlation between performance and energy usage
- BUT highest clock frequency does not equal lower energy
- General model gives better overall predictions than Application-Specific ones

# Thank you!

- kevin@kevinhammond.net
- ya8@st-andrews.ac.uk

