Towards Analysis-Driven Multiphysics Software Architecture

Damian Rouson
Reacting Flows Dept., Combustion Research Facility
Sandia National Laboratories

Project: Morfeus

Sponsor: Office of Naval Research

Sandia National Laboratories

Motivation

HPC Survey Results

SMG

What Are the Biggest Bottlenecks Today in Creating Custom Parallel Applications?*

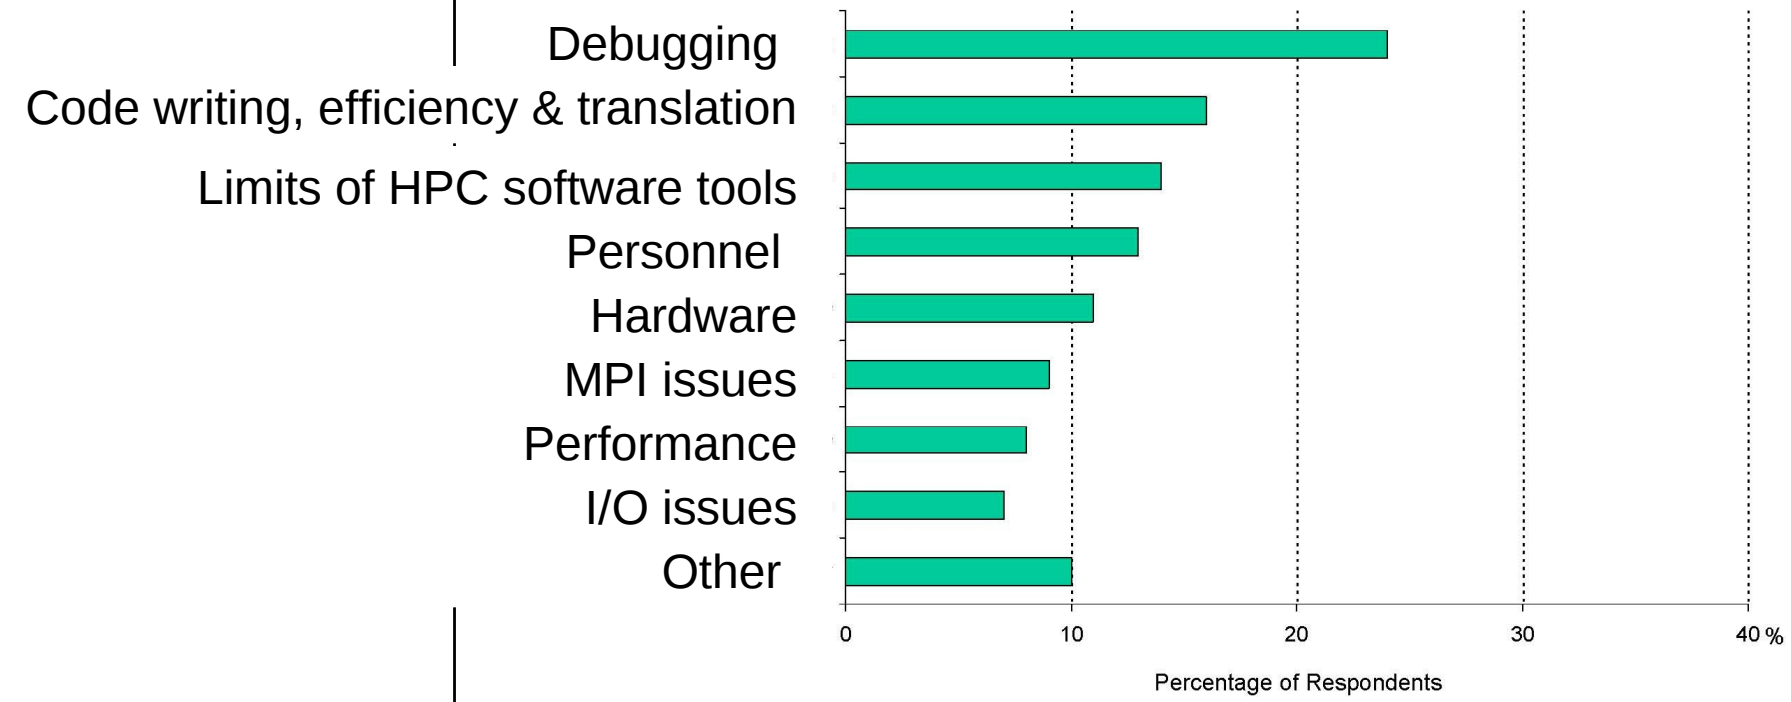

* Note: Multiple responses allowed.

Sandia National Laboratories

Objectives

1. To analyze how development time scales with program size & how this depends on the choice of abstractions.
2. To develop strategies for reducing development time.
3. To demonstrate *scalable development* of multiphysics models.

Outline

1. Analyses of the impact of
   a. Coding efficiency on total solution time.
   b. Programming paradigm on debugging.
   c. Abstraction choice on interface content.
2. Multiphysics model demonstrations
3. Conclusions & Future Work

Outline

1. Analyses of the impact of
   a. Coding efficiency on total solution time.
   b. Programming paradigm on debugging.
   c. Abstraction choice on interface content.
2. Multiphysics model demonstrations
3. Conclusions & Future Work

Conventional Development

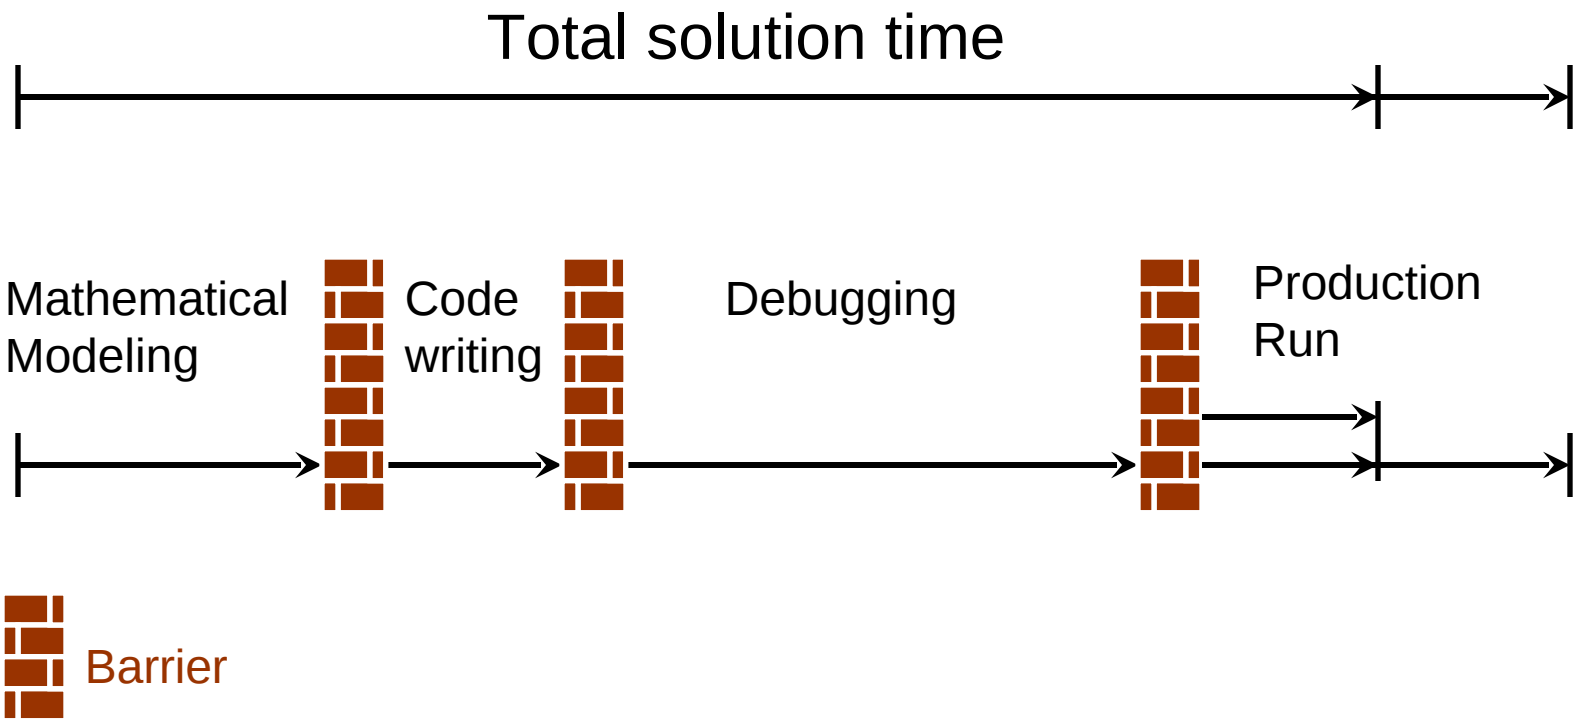

Amdahl's Law

Representative case study for a published run[1,2]:

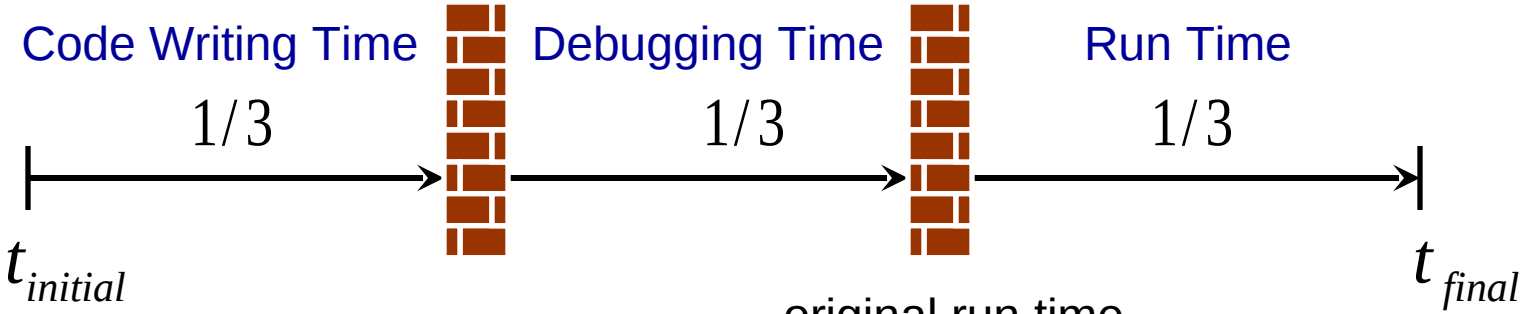

Run-time speedup: $S_{run} \equiv \dfrac{\text{original run time}}{\text{optimized run time}}$

Total speedup: $S_{tot} = \dfrac{1}{\frac{2}{3} + \frac{1}{3}\frac{1}{S_{run}}} \Rightarrow \lim_{S_{run}\to\infty} S_{tot} = 1.5$

The speedup achievable by focusing solely on decreasing run time is very limited.

[1]Rouson et al. (2008) *Physics of Fluids*.

[2]Rouson et al. (2008) *ACM Transactions on Mathematical Software*.

Sandia National Laboratories

Case Study: Isotropic Turbulence

Procedure	Inclusive Run-Time Share (%)
`main`	100.0
`operator(.x.)`	79.5
`RK3_Integrate()`	47.8
`Nonlinear_Fluid()`	44.0
`Statistics_`	43.8
`transform_to_fourier`	38.7
`transform_to_physical`	23.6

Calls

- 5% procedures occupy nearly 80% of run time.
- Structure 95% of procedures to reduce development time.

Sandia National Laboratories

Total Solution Time Speedup

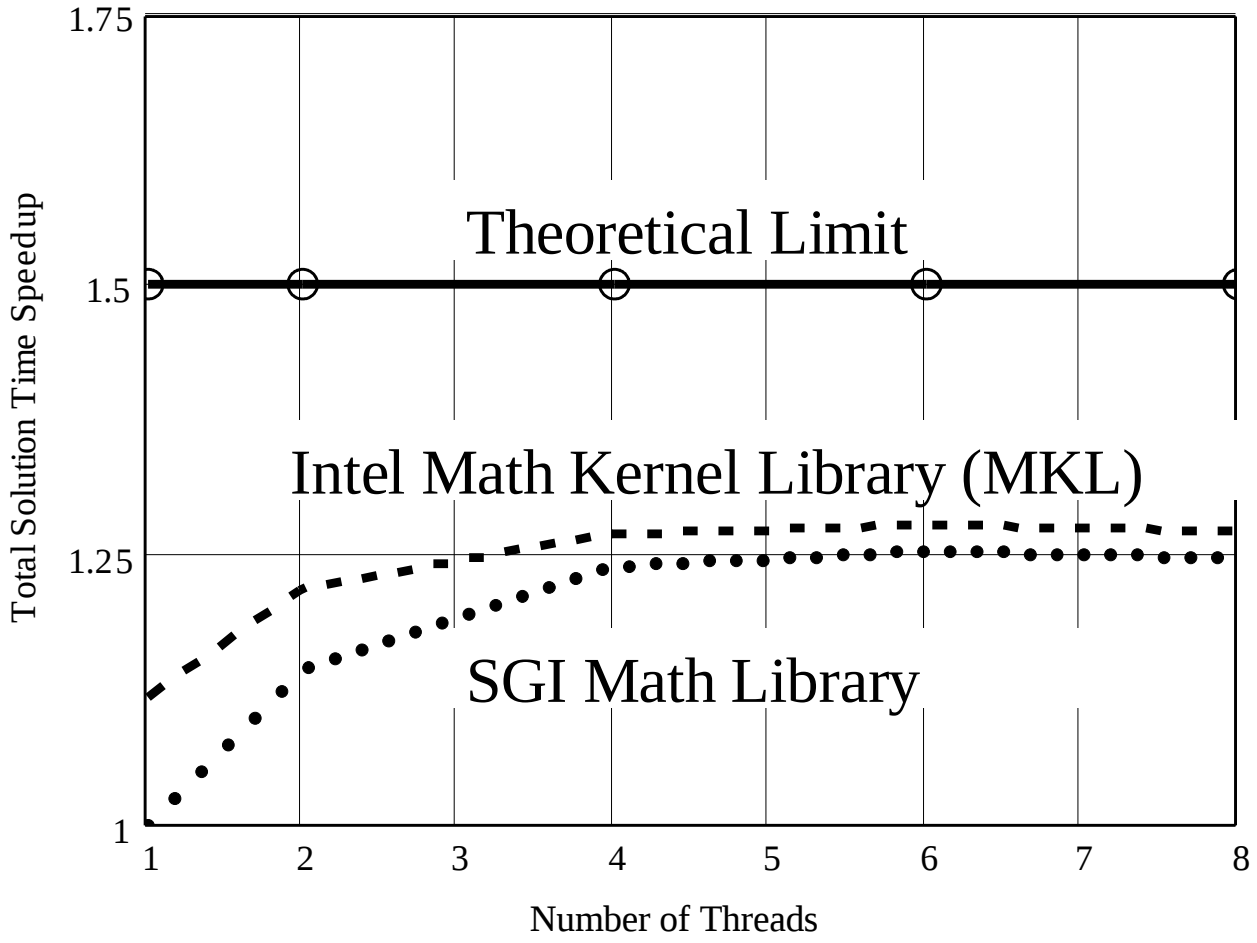

Pareto Principle

When participants (lines) share resources (run time), there always exists a number $k \in [50,100)$ such that (1-k)% of the participants occupy k% of the resources:

Limiting cases:

- k=50%, equal distribution
- k→100%, monopoly

Rule of thumb: 20% of the lines occupy 80% of the run time

Scalability requirements determine the percentage of the code that can be focused strictly on programmability:

$$S_{\max} = \lim_{S_{k\%} \to \infty} \frac{1}{0.2 + 0.8 / S_{k\%}} = 5$$

Sandia National Laboratories

“Separate the physics from the data.”

Jaideep Ray

Sandia National Laboratories, ca. 2005

Sandia National Laboratories

“Software abstractions should resemble blackboard abstractions.”

Kevin Long

Texas Tech. Univ., ca. 2007

Sandia National Laboratories

Abstract Data Type Calculus

Blackboard abstraction

$$T = T(x, y, z, t)$$

$$T(x = 0, y, z, t) \equiv T_0$$

$$T_t \equiv \frac{\partial T}{\partial t}$$

$$T^{n+1} = T^n + \Delta t T_t^n$$

$$\frac{\partial T}{\partial t} = \frac{1}{\alpha} \nabla^2 T$$

$$\nabla^2 T \equiv T_{xx} + T_{yy} + T_{zz}$$

Software abstraction (Fortran 2003):

```
type(field) :: T,dT_dt,laplacian_T
```

```
call T%boundary(x,0,T0)
```

```
T%t()
```

```
T = T + dt*T%t()
```

```
dT_dt = (1./alpha)*laplacian(T)
```

```
laplacian_T = T%xx()+T%yy()+T%zz()
```

Sandia National Laboratories

Morfeus

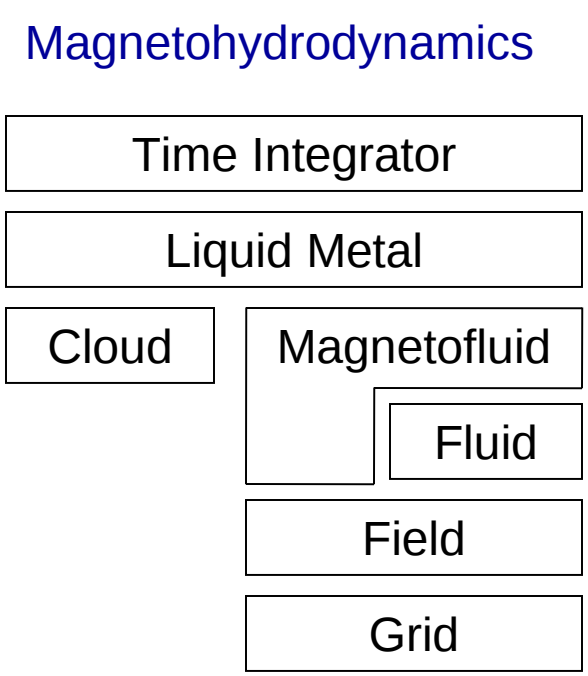

Lattice Boltzmann bio-fluid dynamics:

Xu & Lee (2008) “Application of the lattice Boltzmann method to flow in aneurysm with ring-shaped stent obstacles,” *Int. J. Numerical Methods in Fluids.*

Sandia National Laboratories

Particle-Laden MHD

Strong magnetic fields damp velocity variations in the field direction, leading to 2D/3C state:

$$\frac{\partial}{\partial t}\vec{u}(x,t) = \ldots + \frac{1}{\eta}\nabla^2\left(\vec{B}_A^{ext}\cdot\nabla\right)\vec{u}(x,t) + \ldots$$

Cross-stream dispersion segregates inertial particles:

$$dr_i / dt = \mathrm{v}_i$$

$$d\mathrm{v}_i / dt = \left[u_i(r,t) - \mathrm{v}_i\right] / St$$

$$St \equiv \tau_p / \tau_f$$

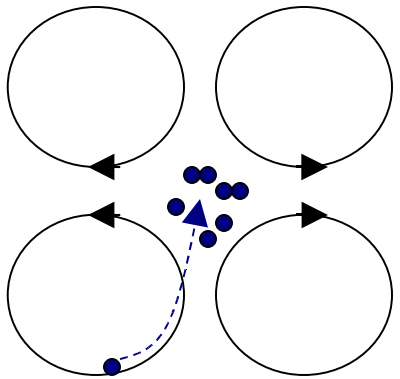

Sandia National Laboratories

Outline

1. Analyses
2. Multiphysics model demonstrations
   a. Particle transport in magnetohydrodynamics.
   b. Quantum turbulence in superfluid ^{4}He.
   c. Atmospheric boundary layer.
   d. Lattice-Boltzmann bio-fluid dynamics.
3. Conclusions & Future Work

Particle-Laden MHD

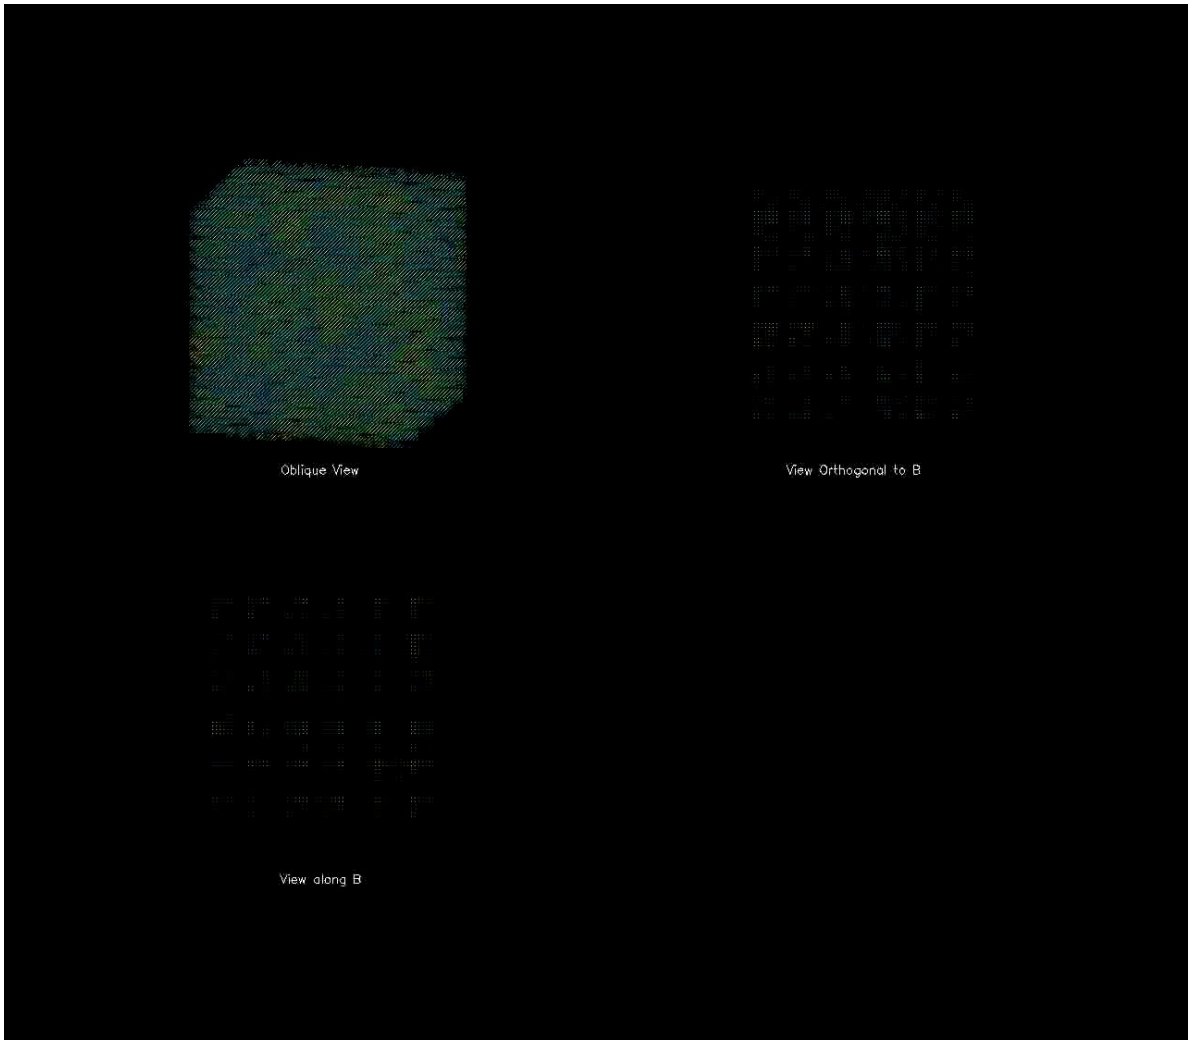

Rouson et al. (2008), “Dispersed-phase structural anisotropy in magnetohydrodynamic turbulence at low magnetic Reynolds numbers,” *Physics of Fluids*.

Sandia National Laboratories

Particle-Laden MHD

Particle positions colored by speed (red=fastest, blue=slowest).

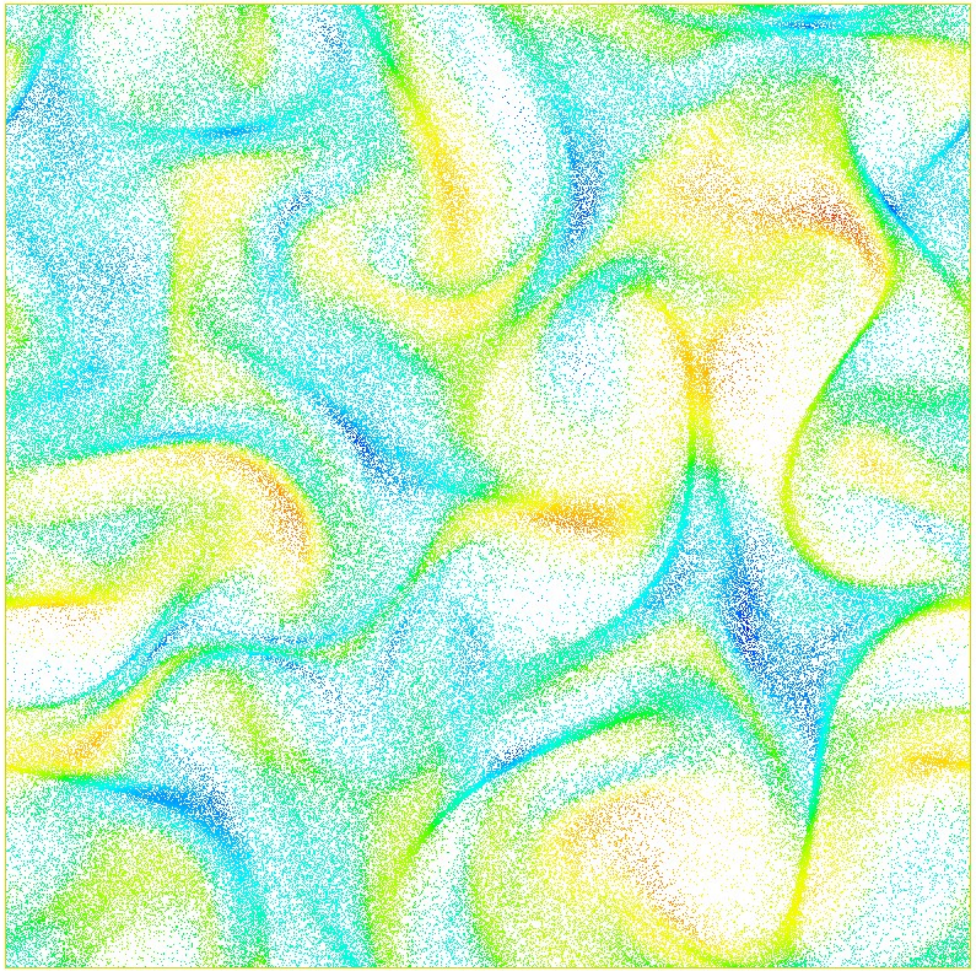

Physics of Fluids Feb. 2008 cover image.

Sandia National Laboratories

Quantum Turbulence

Below 2.17 K, liquid ^{4}He flows as a two-fluid mixture with mutual friction between the two components:

1. Normal viscous fluid

$$\frac{\partial}{\partial t}\vec{u} + \vec{u}\cdot\nabla\vec{u} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\vec{u} + \vec{f}$$

$$\nabla\cdot\vec{u} = 0$$

2. Inviscid superfluid with quantized vortex circulation

$$\kappa = h/m_{He}$$

$$\mathbf{v}_i = \frac{\kappa}{4\pi}\int (\mathbf{S}_0 - \mathbf{r}) \otimes d\mathbf{S}_0 / \left\|\mathbf{S} - \mathbf{r}\right\|^3$$

$$\frac{d\mathbf{S}}{dt} = \mathbf{v}_i + \alpha\mathbf{S}' \otimes (\mathbf{v}_n - \mathbf{v}_i) - \alpha'\mathbf{S}' \otimes [\mathbf{S}' \otimes (\mathbf{v}_n - \mathbf{v}_i)]$$

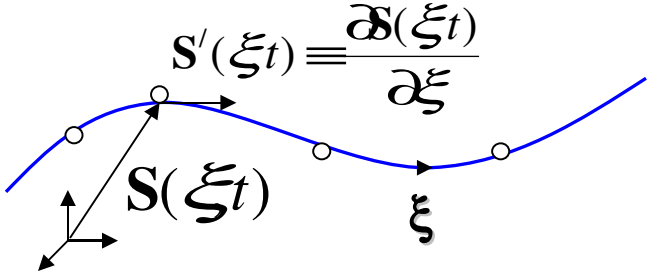

Sandia National Laboratories

Quantum Turbulence

Quantum vortices driven by forced, isotropic normal-fluid turbulence at 2.1 K:

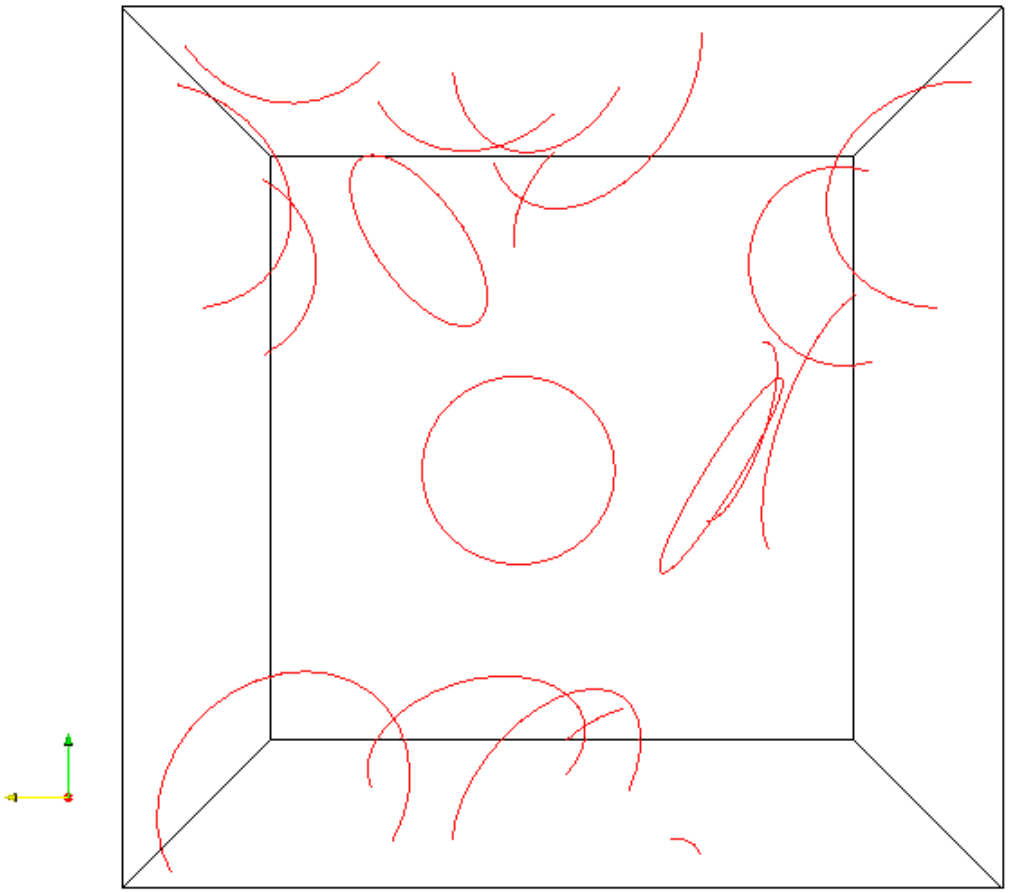

Sandia National Laboratories

Vortex Locking

Quantum vortex alignment with classical vortices in frozen normal-fluid turbulence:

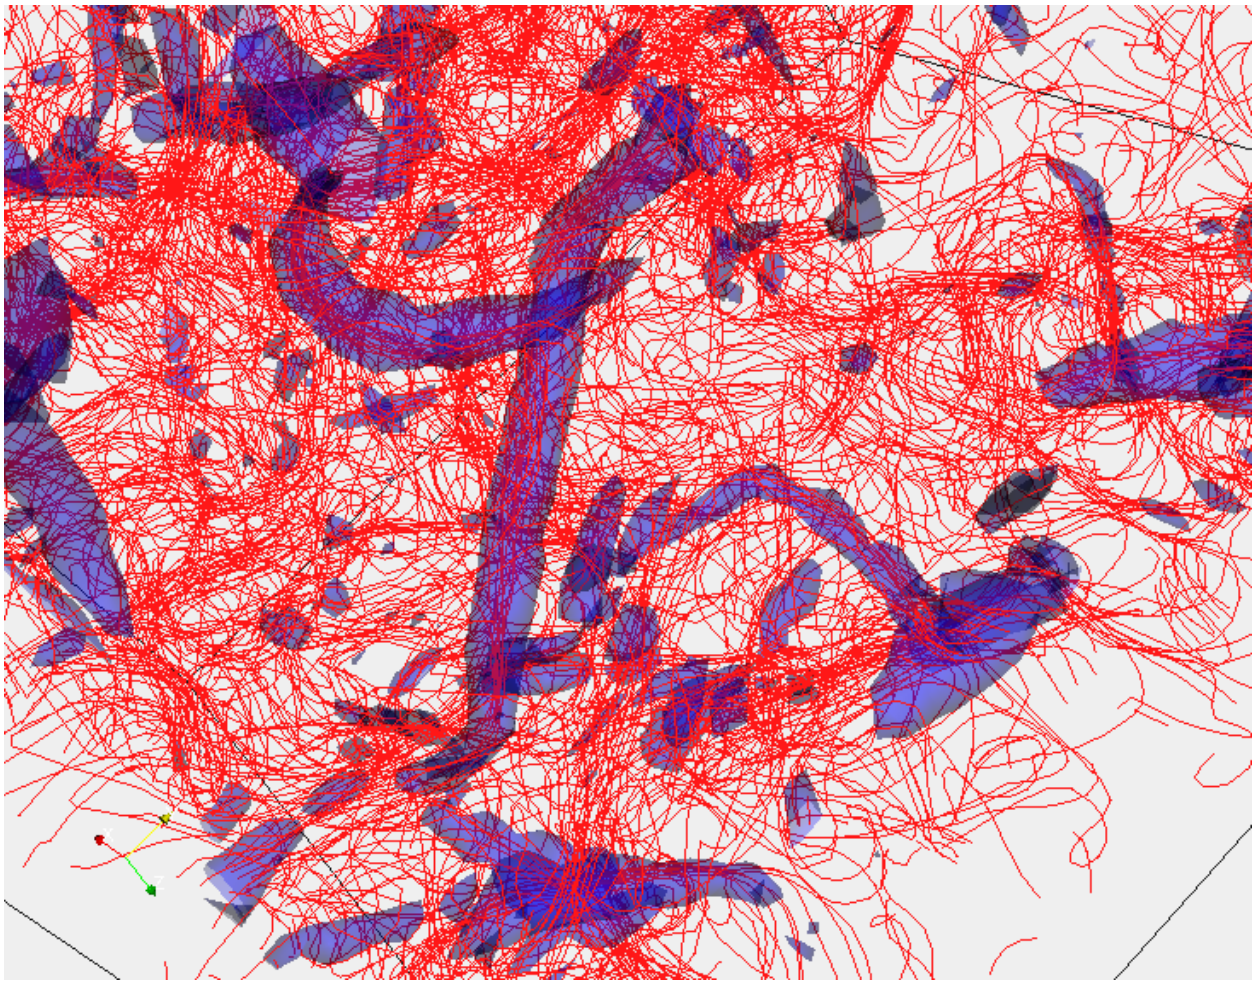

Morris, Koplik & Rouson (2008) “Vortex locking in direct numerical simulations of quantum turbulence,” *Phys. Rev. Lett.*

Sandia National Laboratories

Outline

1. Analyses of the impact of
   a. Coding efficiency on total solution time.
   b. Programming paradigm on debugging.
   c. Abstraction choice on interface content.
2. Multiphysics model demonstrations
3. Conclusions & Future Work

“Procedural programming is like an N-body problem.”

Lester Dye, Stanford University, ca. 1994

Sandia National Laboratories

“What are the metrics?”

Oyekunle Olukotun, ca. 1996

Stanford University

Sandia National Laboratories

“The overwhelming amount of time spent in maintenance and debugging is on finding bugs... The actual fix is relatively short!”
Shalloway & Trott, *Design Patterns Explained*, 2002.

“The hardest part of debugging is finding where the bug is.”
Oliveira & Stewart, *Writing Scientific Software: A Guide to Style*, 2006.

Debugging Structured Programs

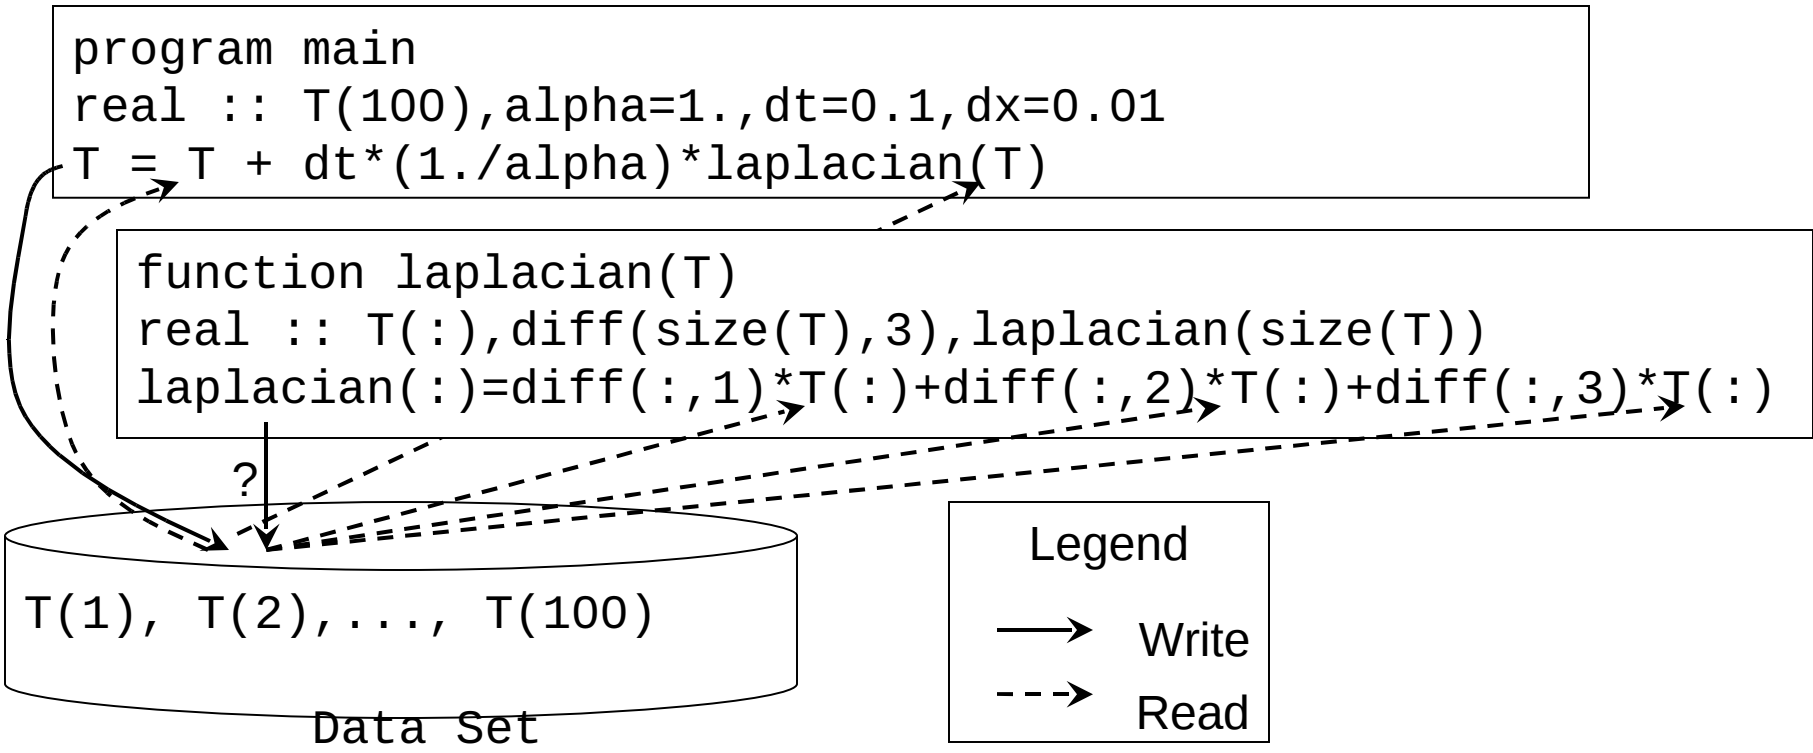

Bisection Search Time

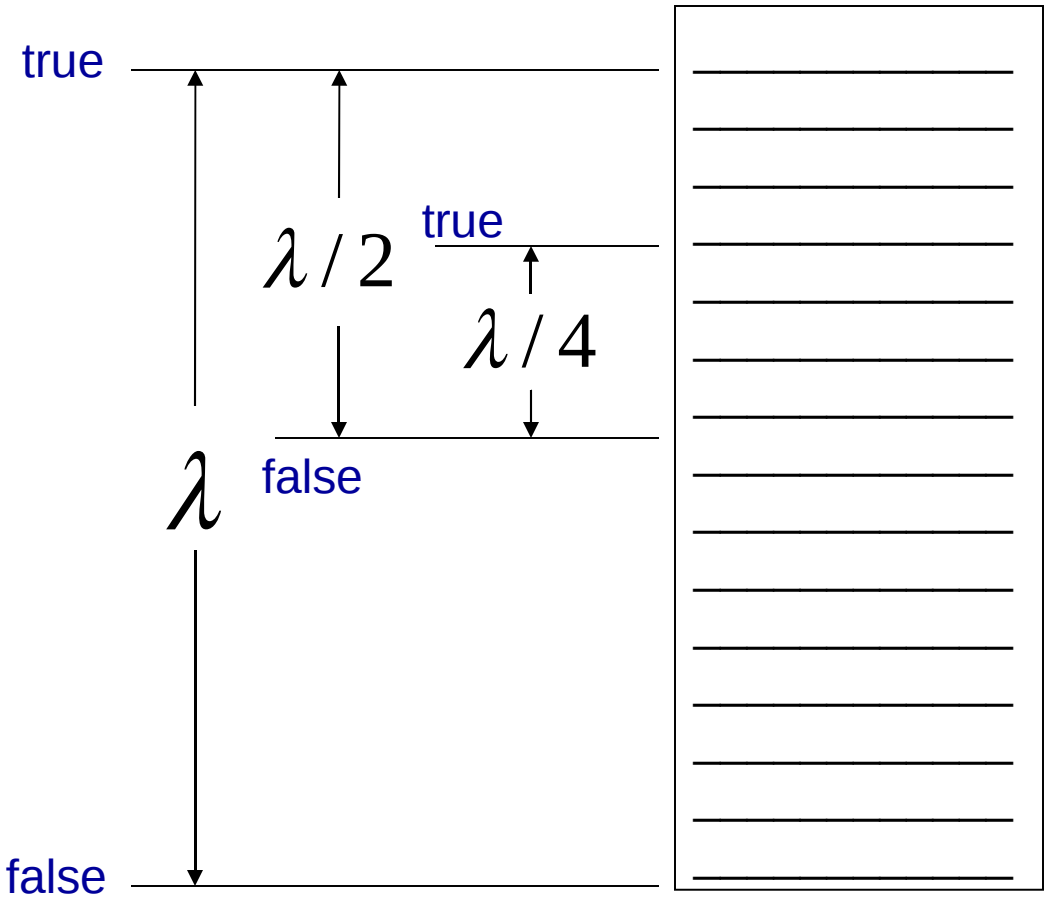

Localization error:

$$\lambda_n = \frac{\lambda}{2^n}$$

Convergence criterion:

$$\frac{\lambda}{2^q} = 1$$

$$q = \log_2 \lambda$$

Search time metric:

$$\lambda_{\text{searched}} = \log_2 \lambda$$

Sandia National Laboratories

Fault Rate

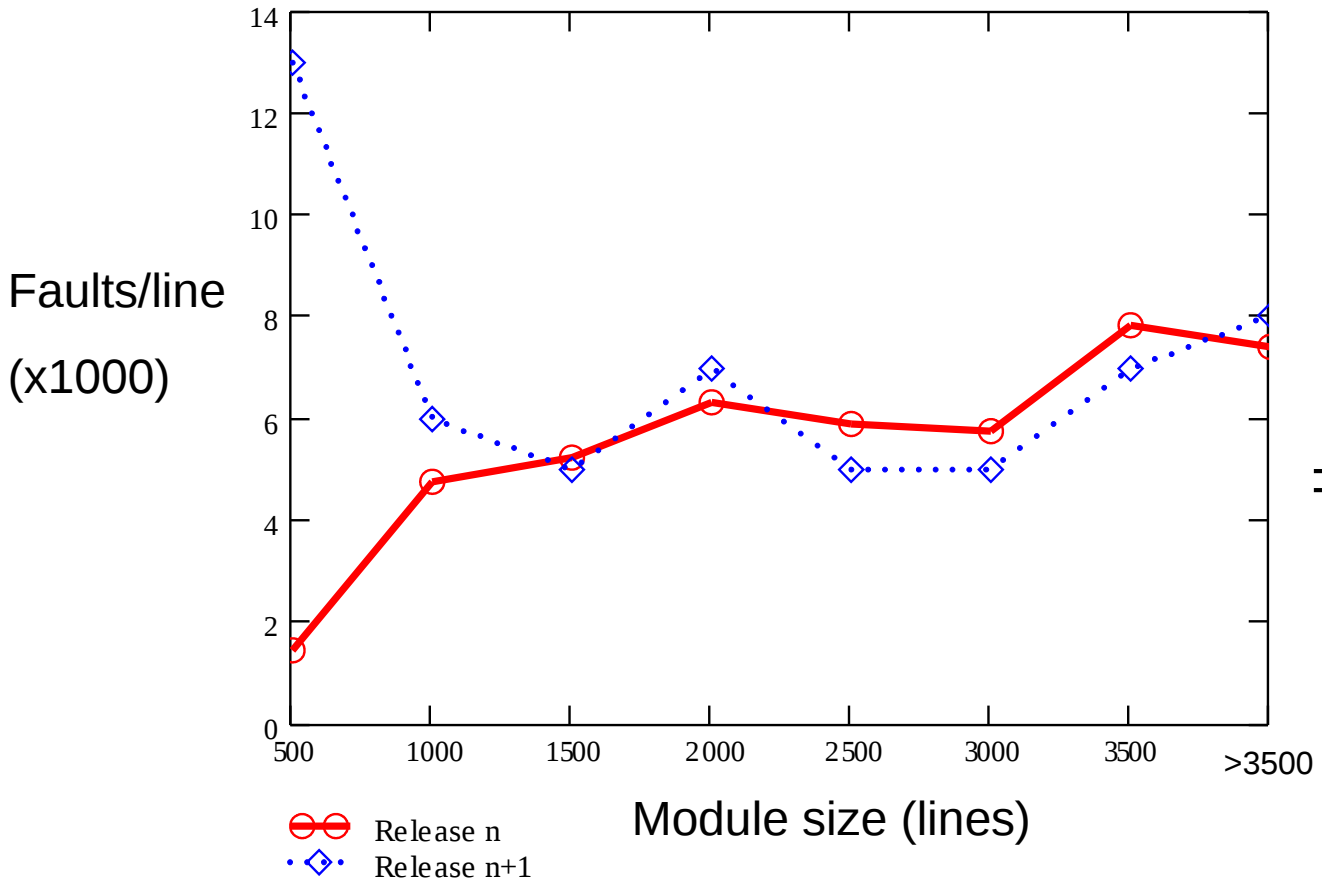

$$\Rightarrow r \approx 6/1000$$

Source: Fenton & Ohlssen (2000) “Quantitative analysis of faults and failures in a complex software system,” *IEEE Trans. Soft. Eng.*

Sandia National Laboratories

Scientific Code Faults

Observed faults in *commercially released* code*:

- 8 statically detectable faults/1000 lines of C code
- 12 statically detectable faults/1000 lines of Fortran 77 code
- more recent data finds 2-3 times as many faults in C++

$$\Rightarrow r \approx 0.006 - 0.036$$

$$t_{search} = (\#\,bugs)(lines\ searched\ per\ bug)(\bar{t}_{line\ review})$$

$$= (r\lambda)[(\lambda/2 - 1)/2]\bar{t}_{line\ review})$$

*Source: Hatton, L. (1997) "The `T' Experiments – Errors in Scientific Software," *Comp. Sci. Eng.*

Sandia National Laboratories

Object-Oriented Program

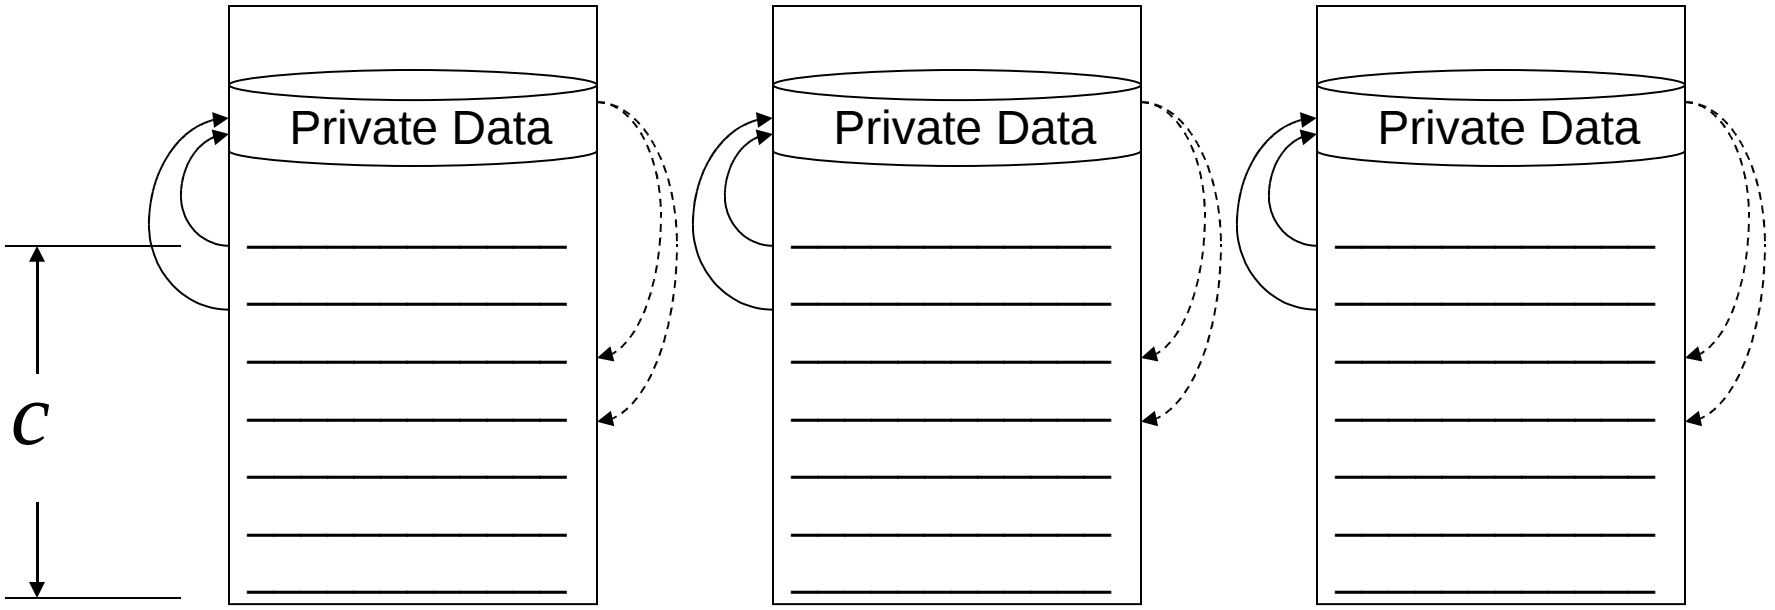

$$\lambda_{searched} = rc \log_2 c$$

$$c << \lambda$$

ADT Calculus

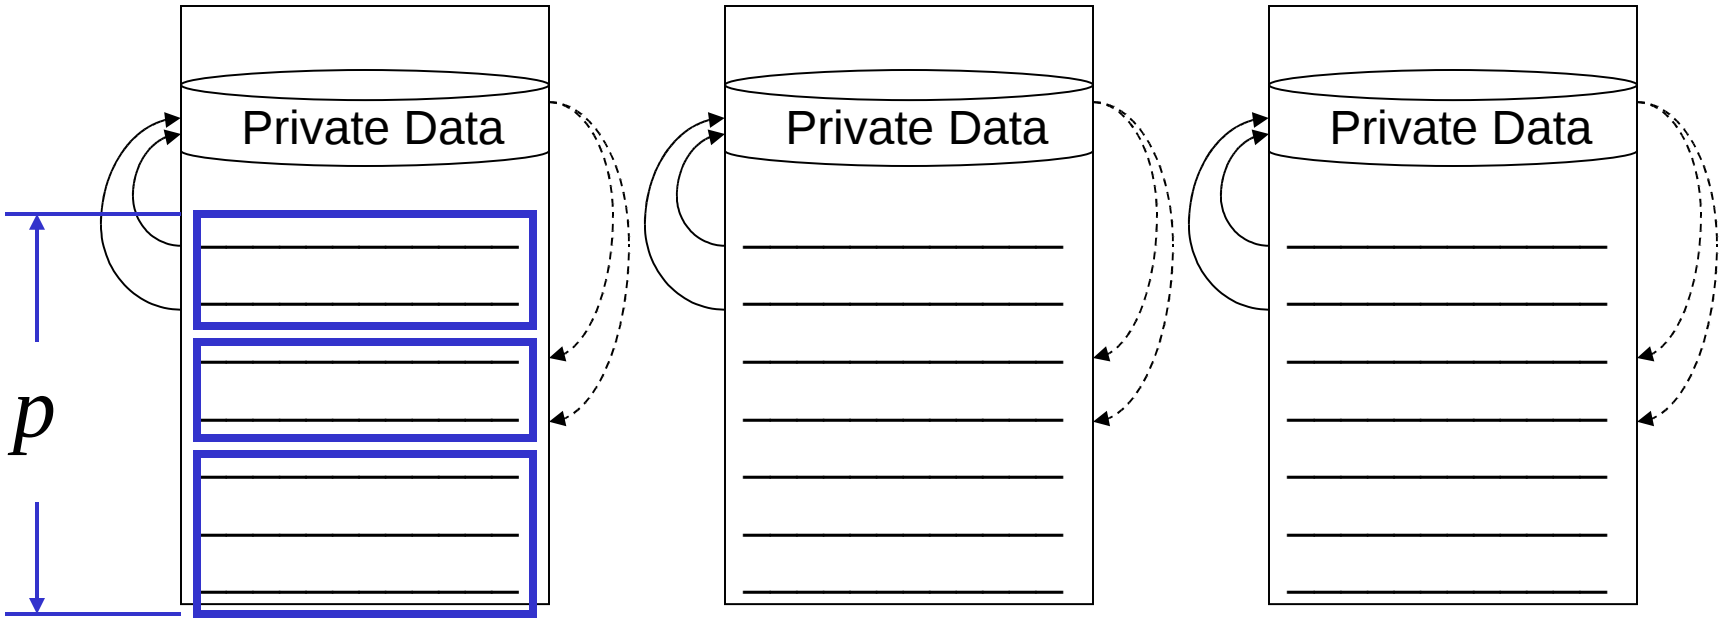

Procedural line density:

$$\rho \equiv \frac{c}{p} = \frac{\textit{lines per class}}{\textit{procedures per class}} = \lambda_{searched} \log_2 p$$

For ADT calculus, $\rho \approx const$, $p \approx const$ $\Rightarrow \lambda_{searched} \approx const$

Sandia National Laboratories

Outline

1. Analyses of the impact of
   a. Coding efficiency on total solution time.
   b. Programming paradigm on debugging.
   c. Abstraction choice on interface content.
2. Multiphysics model demonstrations
3. Conclusions & Future Work

Interface Content

Abstract class interface (Unified Modeling Language):

integrable_model

`+ operator(+)(integrable_model,integrable_model) : integrable_model`
`+ operator(*)(real,integrable_model) : integrable_model`
`+ t(integrable_model) : integrable_model`

A single interface describes all of the public information for all classes that extend this class.

Information Entropy

Shannon (1948) “A mathematical theory of communication,” *Bell System Tech. J.*

The class interfaces embody inter-developer communications. Consider the set of all (*N*) possible messages that can be transmitted between two developers:

> “If the number of messages in the set is finite, then this number or any monotonic function of this number can be regarded as a measure of the information produced when one message is chosen from the set, all choices being equally likely.”

Shannon chose the logarithm because it satisfies several constraints that match our intuitive understanding of information:

$$H = -\sum_{i=1}^{N} p_i \log_2 p_i = -\sum_{i=1}^{N} \frac{1}{N} \log_2 \frac{1}{N} = \log_2 N$$

Sandia National Laboratories

Minimum Information Growth

```
subroutine integrate(integrand)
  class(integrable_model) :: integrand
  integrand = integrand + dt*integrand%t()
end subroutine
```

If only one class extends integrable_model, the executable line only has one possible interpretation, so *H=0*. Each subsequent subclass increases the information content by

$$\Delta H = \log_2(N+1) - \log_2 N$$

which is the minimum information growth.

Sandia National Laboratories

Conclusions

- Applying Amdahl's law to the *total* solution time suggests that optimizing runtime only severely limits speedup.
- The Pareto Principle determines the percentage of the code that can be focused strictly on programmability rather than runtime efficiency.
- ADT calculus renders bug search times very nearly scale-invariant and reduces interface information content.

“More computing sins are committed in the name of efficiency (without necessarily achieving it) than for any other single reason - including blind stupidity.” W. A. Wulf

Source: S. Kargl, comp.lang.fortran, 11/6/08.

Sandia National Laboratories

Abstract Data Type (ADT)

```
module field_class                    !C++ namespace
  implicit none
  private
  type, public :: field               !C++ class
    private                           !C++ data members
    real, dimension(:,:,:), allocatable :: nodalValues
  contains
    procedure :: boundary             !C++ member functions
    procedure :: plus                 !C++ overloaded operator
    generic, public :: operator(+)=>plus
  end type
contains
  subroutine boundary(this,direction,location,value)
    class(field) :: this              !C++ dynamic dispatching
    ...
  end subroutine
  function plus(lhs,rhs) result(total)
    class(field), intent(in)  :: lhs,rhs
    class(field), allocatable :: total
    ...
```

Sandia National Laboratories

Future Directions

- Demonstrate runtime scalability.
- Add empirical support for reductions in
  - Fault localization time.
  - Information entropy*.

*Kirk & Jenkins (2004) “Information theory-based software metrics and obfuscation.” *J. Systems & Software.*

Sandia National Laboratories

Fault Localization

"Computational" complexity theory: Derive a polynomial time estimate for fault localization in a chronological list of the unique program lines executed:

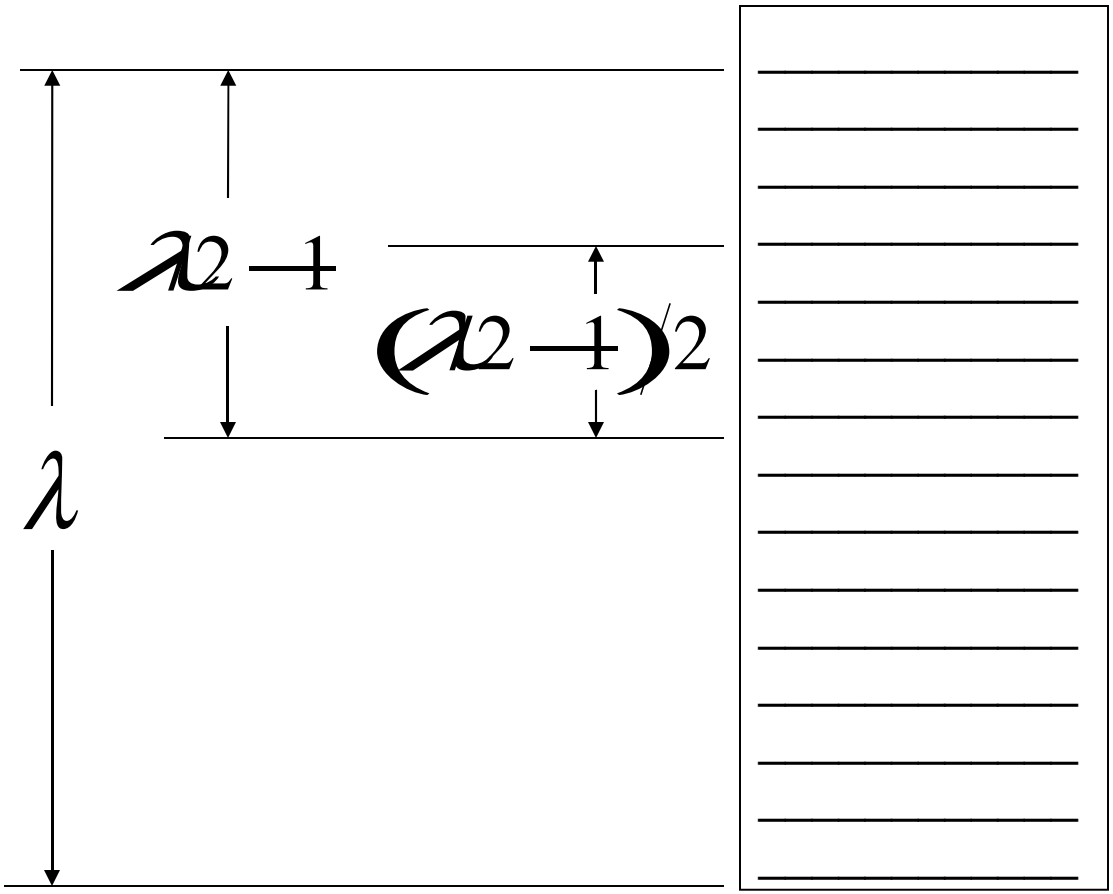

$\alpha < 0$ (bug)

$T(2) < 0$ (symptom)

Sandia National Laboratories

Structure Tensors

One-point turbulence structure tensor statistics:

$$\left\langle u^2 \right\rangle \delta_{ij} \equiv R_{ij} + D_{ij} + F_{ij}$$

$$\mathbf{u} = \sum_{\mathbf{k}} \hat{\mathbf{u}}_{\mathbf{k}} e^{i\mathbf{k}\cdot\mathbf{x}}$$

$$R_{ij} \equiv \left\langle u_i u_j \right\rangle, \quad D_{ij} \equiv \int \frac{k_i k_j}{k^2} E_{nn} d^3\mathbf{k}, \quad F_{ij} \sim \int \frac{\overline{\hat{\omega}_i \hat{\omega}_j^*}}{k^2} d^3\mathbf{k},$$

"Componentality" (Reynolds stress)

"Dimensionality"

"Circulicity"

Sandia National Laboratories

Dispersed Phase

$$D_{ij}^{p} \sim \int \frac{k_i k_j}{k^2} \left\langle \hat{n}\hat{n}^{*} \right\rangle d^3\mathbf{k},$$ ← Describes preferred orientation of preferential concentration

$$d_{ij}^{p} \equiv D_{ij}^{p} \Big/ D_{mm}^{p}$$ ← Minimize impact of non-stationarity

$$\tilde{d}_{ij}^{p} \equiv d_{ij}^{p} - \delta_{ij} \Big/ 3$$ ← Remove information about *degree* of preferential concentration.

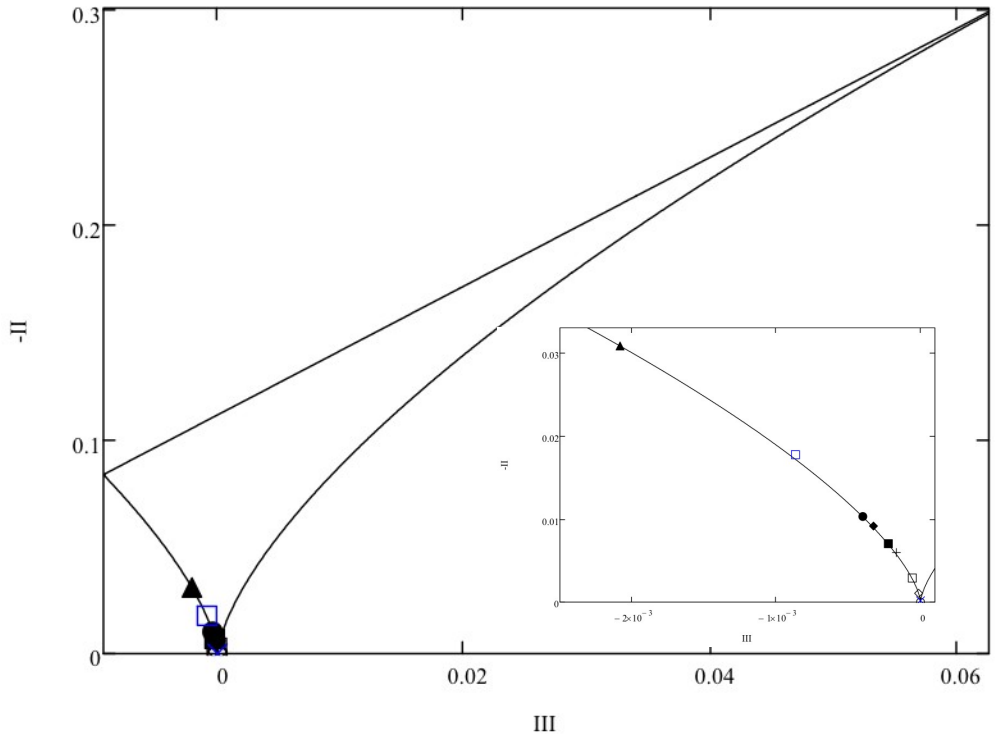

Sandia National Laboratories