



Productivity and HPC

**Erik Hagersten,
CTO, Rogue Wave Software AB**

Developing parallel, data-intensive applications is hard.

We make it easier.

-
- **Chief architect high-end servers Sun Microsystems 1994 – 1999**
 - **Professor Uppsala University 1999 –**
 - **Startup: Acumem 2006 – 2010 (Now Rogue Wave Software AB)**
 - **CTO Rogue Wave Software AB 2010 –**
 - **On the board of Swedish National Infrastructure for Computing (SNIC) 2003 –**

Rogue Wave Software AB

- ➔ **Unique world-leading technology that makes performance optimization fast and simple.**
- **In ITEA2 EU project: H4H**
 - Examples of collaborators: Bull, BSC, Jülich, ...
- **In FP-7 EU project: HOPSA**
 - Examples of collaborators: Jülich, BSC, TUD... MSU, T-Platform (Russia)

*The national security research we conduct requires computational technology that allows **programmers to optimize their codes** and thereby get the best performance from their systems. By working with industry leaders, such as Acumem, LLNL is **able to do** some of the most advanced software **development for some of the world's most powerful supercomputers faster while reducing costs.***

Mike Kumbera, Lawrence Livermore National Labs

The HPC Landscape is Changing in Many Dimensions

- **Architectures are getting more complex**
 - Multicores with deep memory hierarchies and many bottlenecks
 - Heterogeneous tendencies: ISA, NUMA, memory capacity, ...
- **Demand for parallelism is higher**
 - Achieving good performance from a single CPU chip requires ~10 threads.
 - Number of threads per job is growing
- **New user groups**
 - New emerging application areas: Linguistics, biology, ...
 - Typically, computer science is not a strong discipline
- **Moving into more commercial settings**
 - New requirements and new rules for success
 - Time to market, productivity and reliability are key


These Changes are not Compatible

- **Architectures are getting more complex**
 - Requires extreme performance experts
 - Requires wading through massive amount of performance data for week
- **Demand for parallelism is higher**
 - Requires extreme algorithm experts
 - Writing and debugging parallel code is time-consuming (some say 10x...)
- **New user groups**
 - Typically, computer science is not a strong discipline
 - Neither knowledge or urge to spend >50% time programming
- **Moving into more commercial settings**
 - May have some of the expertise needed
 - Need to work

➔ Need better languages, environments and tools...

... for the experts as well as ordinary users

Need Good Tools and Technology at Each Development Stage



Rogue Wave Today



The largest independent provider of cross-platform software development tools and embedded components for the next generation of HPC applications

Visual Numerics®

Leader in embeddable math and statistics algorithms and visualization software for data-intensive applications.



Leading provider of intelligent software technology which analyzes and optimizes computing performance in single and multi-core environments.



Industry-leading interactive analysis and debugging tools for the world's most sophisticated software applications.

Rogue Wave Vision

WE ENABLE HPC

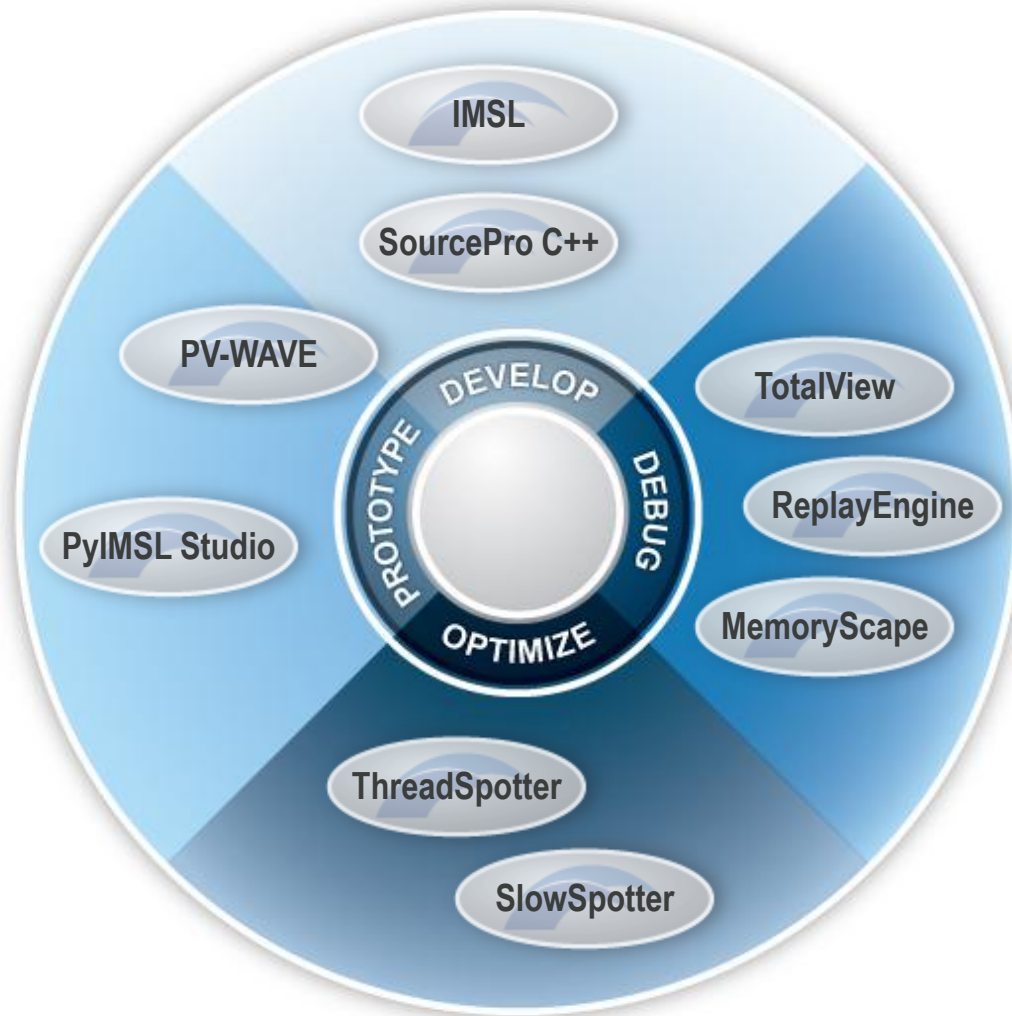
- **Best of breed HPC development products**
- **Cross-platform tools and embedded components**
- **Leading solutions for complex data analysis**
- **Increased productivity and decreased complexity**



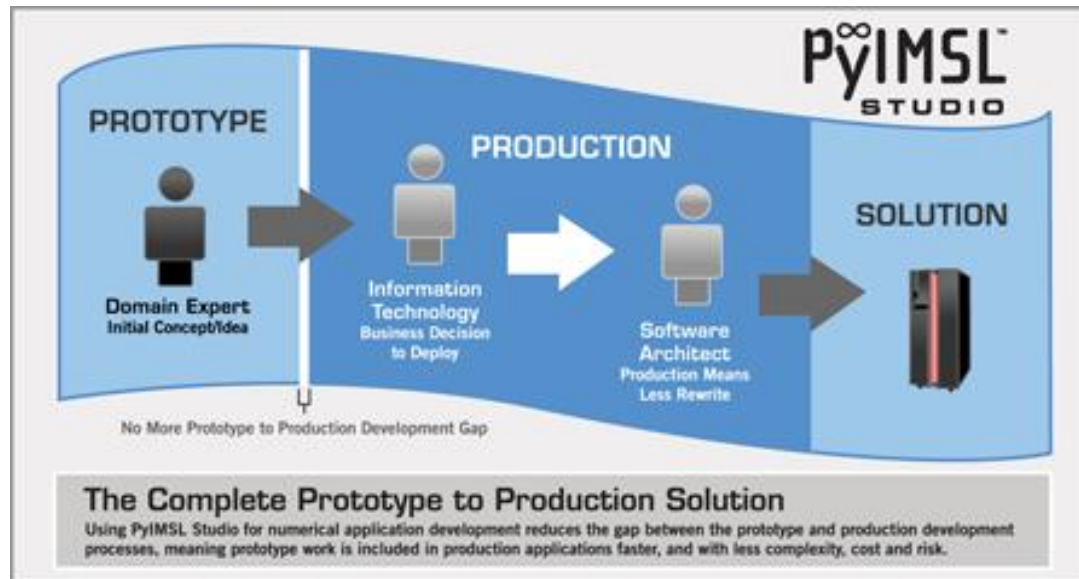
**Developing parallel, data-intensive applications is hard.
We make it easier.**



Rogue Wave Product Offerings

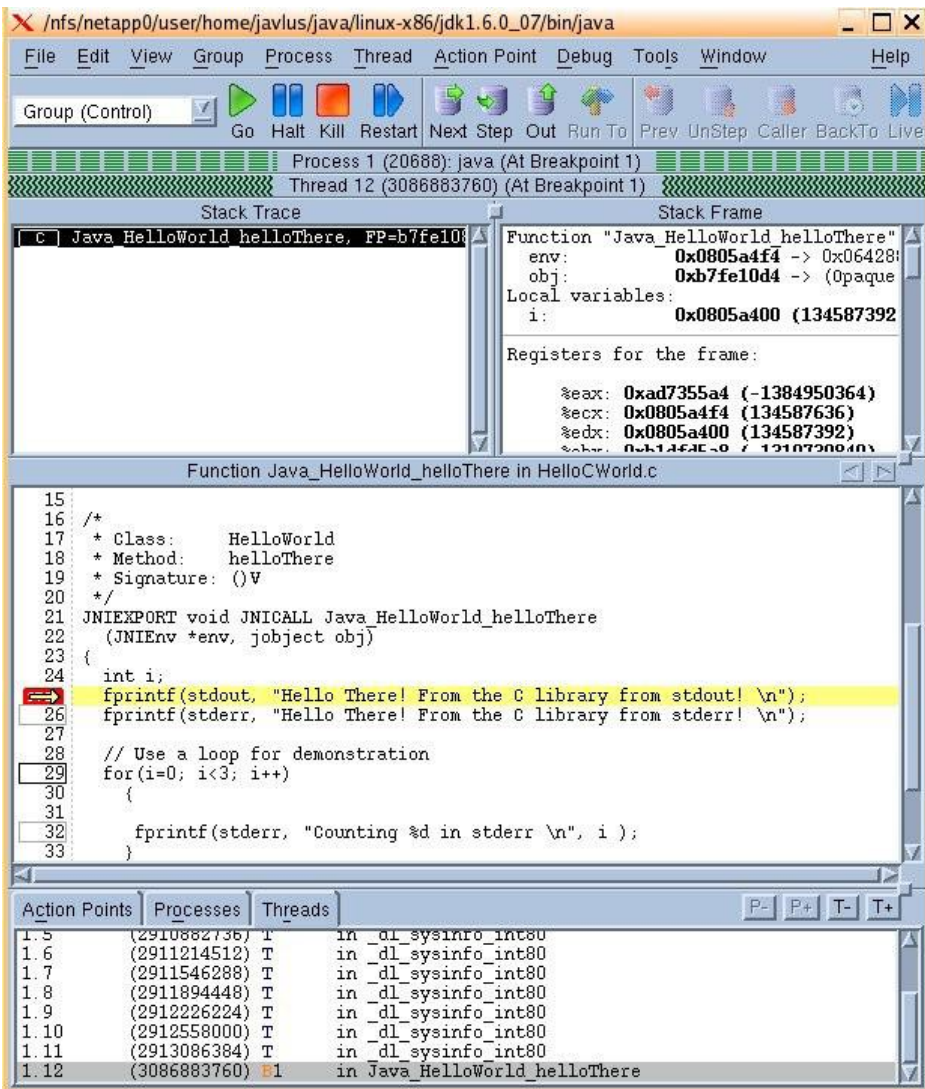


Prototype & Develop: PyIMSL and IMSL



- **PyIMSL – prototyping environment**
 - Python environment including wrappers for the 800 thread-safe IMSL functions
 - Domain-expert to prototype and test in minutes
- **IMSL – highly optimized and portable math and stat. functions**
 - PyIMSL prototype easily ported to highly optimized implementation
 - Includes some parallel and GPU function implementations

Debug: TotalView



Advanced Parallel Debugger:

- Supports many arch (incl. CUDA)
- Individual control over threads
- Attach to a part of a server
- Scripting debugging
- Execution points ...

MemoryScope:

- Automatically finds hard-to-detect memory bugs

Fast Deterministic Replay (x86)

- Makes data-raced repeatable
- Can execute backwards



Optimizing: ThreadSpotter

Source:
C, C++, Fortran, ADA...

```

/* Unoptimized Array Multiplication: x = y * z  N = 1024 */
for (i = 0; i < N; i = i + 1)
  for (j = 0; j < N; j = j + 1)
    {r = 0;
     for (k = 0; k < N; k = k + 1)
       r = r + y[i][k] * z[k][j];
    }
}

/* Unoptimized Array Multiplication: x = y * z  N = 1024 */
for (i = 0; i < N; i = i + 1)
  for (j = 0; j < N; j = j + 1)
    {r = 0;
     for (k = 0; k < N; k = k + 1)
       r = r + y[i][k] * z[k][j];
    }
}
    
```

What?

How?

The screenshot shows the Acumem VPE interface. On the left, a table lists performance issues:

Loop / Issue	Filter On	Summary	% of fetches	Utilization	HW-Prefetch	Randomness
Poor utilization			29.4%	12.4%	100.0%	Low
Loop fusion			29.4%	12.4%	97.6%	Low
Inefficient loop nesting			29.2%	12.0%	0.0%	Low
Loop fusion			4.4%	11.8%	97.3%	Low
Poor utilization			4.4%	23.7%	100.0%	Low

The 'Issue #8: Cache line utilization' section shows statistics for instructions of this issue:

Category	Value
% of misses	0.0%
% of fetches	4.4%
Fetch rate	51.7%
Cache line utilization	23.7%
HW prefetch probability	100.0%
Access randomness	Low
Worst instruction	mem30, countor_457

The code editor on the right shows the corresponding C code with annotations for the analyzed instructions.

Where?

Any Compiler

Sampler

Model

Binary

Host System

- Click one button to get advice — WHAT!, WHERE! and HOW!
- Context-driven manual (no clip)
- Enables mortals to optimize
- Makes experts more productive
- Efficient use of HW resources
- Saves energy



Summary

- **Bringing productivity into high-performance computing**
- **Enables mortals to do more**
- **Makes experts more productive**

... across all the steps in the development process

**Developing parallel, data-intensive applications is hard.
We make it easier.**



Thank You

