34 Father of the Modern Supercomputer

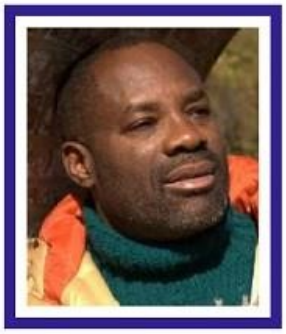

Philip Emeagwali
info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

I'm **Philip Emeagwali**.
Please allow me to quote
from the book titled:
"History of the Internet."

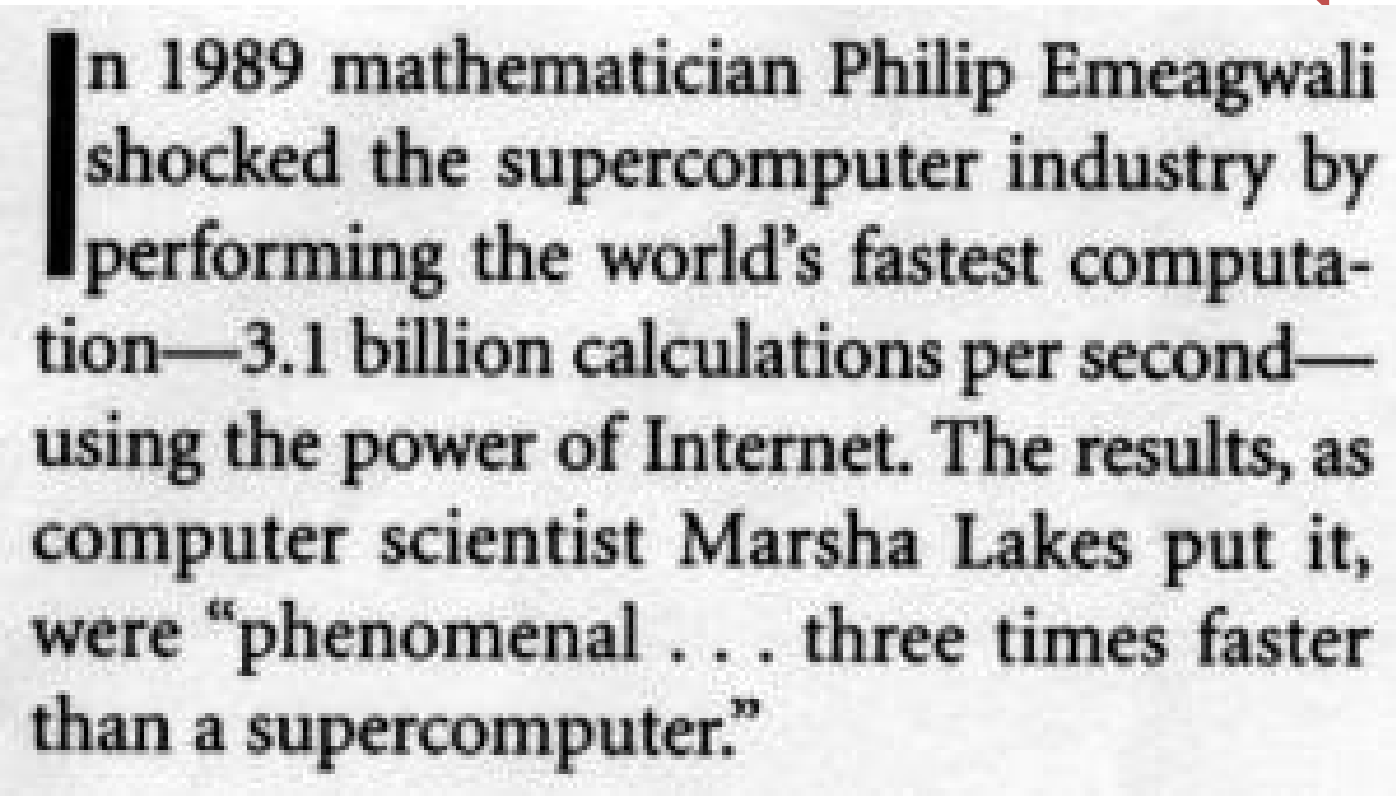

In 1989 mathematician Philip Emeagwali shocked the supercomputer industry by performing the world's fastest computation—3.1 billion calculations per second—using the power of Internet. The results, as computer scientist Marsha Lakes put it, were "phenomenal . . . three times faster than a supercomputer."

I invented
the massively parallel processing
supercomputer.
My invention
contributed to our understanding
of my new supercomputer
as a new internet
that is a new global network of
65,536 already-available processors.

I invented how to harness
the power of a new internet.
My invention
contributed to our understanding
of how to reduce the **time-to-solution**
of the toughest problems
in extreme-scale computational physics
and how to reduce that **time-to-solution**
by a factor of one billion **across**
one billion processors
that were already available
in the market
that are identical
and that are equal distances
apart.
The modern supercomputer
that computes in parallel,
or by solving
a million mathematical problems
at once

arose from our need
to make the **impossible-to-compute**
possible-to-compute.
All living beings
has the intrinsic need
to explore, or to see something
that was previously unseen.
That need to see that previously unseen
is what took us to the moon.
That need to conquer new frontiers
of outer space
is why we sent **unmanned spacecrafts**
to the **planet Mars**.
But the most important frontier
of exploration and knowledge
is not outer space.
The most important frontier
is our inner space.
The frontier of inner space
is where we see the previously unseen

and see them
in unexpected places.

34.1.1 Who's Philip Emeagwali?

Who is **Philip Emeagwali**?
On the Fourth of July 1989
in Los Alamos, New Mexico,
United States,
I invented
how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**,
namely, **the Philip Emeagwali formula that President Bill Clinton described in his speech of August 26, 2000.**
I invented
how and why

to use the modern supercomputer
to solve the toughest problems
in extreme-scale computational physics.
I invented
how and why
to solve those computation-intensive
problems
and solve them
across a new internet
that is a new global network of
64 binary thousand
already-available processors,
or as many tiny computers.
The most computation-intensive
problems in physics
include problems arising from encoding
the laws of motion of physics
and encoding those laws
into the partial differential equations
of modern calculus

that are, in turn, discretized (or reduced)
to an extreme-scale system of equations
of algebra,
such as using those algebraic equations
to **foresee** otherwise **unforeseeable**
global climate change
or to **recover** otherwise **unrecoverable**
crude oil and natural gas
or to solve a thousand
computation-intensive problems
arising in extreme-scale
computational physics.
As I gradually discovered
during my supercomputer quest
of the fifteen-year period
that was inclusive of from June 20, 1974
to the Fourth of July 1989,
that massively parallel processing
supercomputer
and that was a previously unimagined

internet
is a new supercomputer,
that is the fastest computer, *de facto*.

34.1.2 History of Fastest Supercomputers

My 15-year quest
was for the fastest supercomputer
that is a massively parallel processing supercomputer
that is fastest out of the **rawness**
of its **raw** processors.
Back in nineteen forty-six [1946],
the fastest supercomputer
represented
the first generation
of the Information Age.
In nineteen fifty [1950],
the fastest supercomputer
was powered by vacuum tubes.
The supercomputers of the nineteen

sixties [1960s]
represented
the second generation of the
Information Age.
The fastest supercomputers
of the nineteen sixties [1960s]
were powered by transistors.
The fastest supercomputers
of the nineteen seventies [1970s]
were powered by only one
isolated processor
that was not a member
of an ensemble of processors
that communicates and computes
together
and did both as one seamless, cohesive
supercomputer,
rather than as disparate processors.
That is, my 64 binary thousand
computer codes
were processed separately

across as many processors.
And my computer codes were unified
into an extreme-scaled
supercomputer code.
That is, my 65,536 computer codes
that made the news headlines
back in 1989
were integrated
and affected each other.
The paradigm in fast supercomputing
shifted on the Fourth of July 1989,
the day I invented
the precursor
to the modern supercomputer
that computes by parallel processing **across**
a new internet
that is a new global network of
65,536 already-available processors
that were tightly-coupled
and that shared nothing with each other.

My experiment of the Fourth of July 1989
gave me my first **startlingly** clear invention
of a new internet
that is the precursor to the modern,
massively parallel processing
supercomputer.
That invention
of how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**
made the news headlines
in newspapers around the world
and earned me the top prize
in the field of supercomputing.
That invention
opened our eyes
and made it possible to see the computer
in a different way.

Before the Fourth of July 1989,
the fastest computations
were recorded on a supercomputer
that used only one **isolated**
processor
that was **not a member**
of an ensemble of processors
that communicated and computed
together
and that did both as one seamless,
cohesive supercomputer.
That singular processor
was the **heartbeat** of the supercomputer.

34.1.3 The Invention of Parallel Processing for the Modern Supercomputer

To discover is to understand something
that was **misunderstood**
and understand it in a new way.

I understood a new global network
of the slowest 65,536
processors that were already available in
the market
as a new internet
that was **misunderstood**
as something else.
I was the first supercomputer scientist
to understand that new technology
to be a new internet.
Back in 1989, the 25,000
supercomputer scientists in the world
that followed the
vector processing supercomputer vision
of Seymour Cray
misunderstood the new technology
to be merely a massively
parallel processing machine
that will forever have performance problem.

By 1989, I understood that new technology
to be a new supercomputer
whereas the supercomputer textbooks
of the 1940s through ‘80s
considered it impossible
for that parallel processing machine
to compute faster than
the fastest supercomputer
that computed sequentially.
At that time, the established truth
in supercomputer textbooks
was called **Amdahl’s Law**.
In the most quoted scientific paper
in supercomputing
that was published in April 1967,
Gene **Amdahl** wrote that
it will forever be impossible
to achieve a speed increase
of a factor of eight, or more,

and achieve that speedup
by using eight processors, or more,
to power a supercomputer
that is configured as an ensemble
of eight, or more, processors.
Twenty-three years after **Amdahl's Law**
entered into the supercomputer textbook,
it made the news headlines
that I—**Philip Emeagwali**—
has discovered
that the **impossible**-to-compute is, in fact,
possible-to-compute,
namely, that it is possible to exceed
Amdahl's Law limit.
That limit was a factor of eight
speed increase.
I invented
how to exceed **Amdahl's Law limit**
and I did so when I recorded a factor of

65,536 fold increase
in the speed of the supercomputer.
I achieved that speed increase
in supercomputing with 65,536
processors
that were already available
in the market anyway.
I programmed the first supercomputer
to be rated at
one million instructions per second
and I did so on June 20, 1974
at age nineteen.
I programmed that first supercomputer
in Corvallis, Oregon, United States.
I began supercomputing
by programming supercomputers
that were powered by only
one processor
and that computed only

one thing at a time.
In the early 1970s,
they were about one thousand
supercomputer scientists
in the world,
and each supercomputer scientist
programmed
a sequential processing supercomputer.
In the 1970s and ‘80s,
supercomputer textbook authors
wrote that to parallel process
—or to compute many things (or processes)
at once, instead of computing
only one thing at a time—
was a huge waste of everybody’s time.
In the spirit of the times,
the June 14, 1976 issue
of the ***Computer World***
the flagship publication

of the **computer world**
carried an article that was titled:

[quote]
"Research in Parallel Processing Questioned as 'Waste of Time'"
[unquote]

The following timeline and facts
speak for themselves.
In the 1950s and '60s,
the top 500 supercomputers
in the world
performed their fastest computations
by using only one **isolated**
sequential processing unit
that was **not a member**
of an **ensemble** of processors
that communicated and computed
together

and did both as one seamless, cohesive supercomputer.
In the 1970s and ‘80s,
the top 500 supercomputers
in the world
performed their fastest computations
and did so by using only one **isolated**
vector processing unit
that was **not a member**
of an **ensemble** of processors.
In the 1990s and later,
the top 500 supercomputers
in the world
performed their fastest computations
and did so by using the slowest
65,536 processors
that were already available
in the market
or using up to ten million
six hundred and forty-nine thousand

six hundred [10,649,600]
processors that were already available
in the market
that processed together
to execute extreme-scale computations.
From my seven-decade timeline,
the way we think about the computer
and the supercomputer
changed after my invention
of the massively parallel processing
supercomputer
that occurred on the Fourth of July 1989.
So, what happened
on the Fourth of July 1989
that convinced the world of
supercomputing
to change the way
it thought about the **computer**
and change the way

it thought about the **supercomputer**
and change its long-held opinion
that **parallel processing**
is a huge waste of everybody's time?
My **invention**
of the massively parallel processing
supercomputer
happened on the Fourth of July 1989.
My **invention** of how and why
parallel processing
makes modern computers **faster**
made the news headlines
in 1989
and was recorded in the June 20, 1990 issue
of *The Wall Street Journal*.
My **invention**
of how and why parallel processing
makes supercomputers
fastest

has been the subject of
millions of school reports,
since 1989.
For instance, according to the June 20,
1990 issue
of the *Wall Street Journal*,
Philip Emeagwali
discovered that we must look at
the supercomputer of today
as powered by
the **slowest 65,536**
processors,
instead of as powered by
the **fastest singular**
processor.
Since the supercomputer of today
will become the computer of tomorrow,
I reasoned that the computer
should be powered by many

processors
and my **prophecy** became true
after my invention
of the massively parallel processing
supercomputer.
Yet, massively parallel processing
is easier theorized
than discovered.
A theory is an idea
that is not positively true.
My invention
of the massively parallel processing
supercomputer
prompted **Steve Jobs**
who at that time had left Apple Corporation
to reach out to me by telephone
in about June 1990.
Eighteen years later, on June 9, 2008,
Steve Jobs told the opening session
of Apple's Worldwide Developers Conference
in San Francisco, California

that parallel processing
is a huge waste of time.
As reported, one day later,
in the June 10, 2008 issue
of the *New York Times*,
Steve Jobs
told Apple's Worldwide Developers
that:

[quote]
PROP ALERT
"The way the processor industry
is going
is to add more and more cores,
but nobody knows
how to program
those things.
I mean, two, yeah;
four, not really;
eight, forget it."
[unquote]

Like other computer scientists,

Steve Jobs was merely restating
an untrue hypothesis
that **Gene Amdahl**
of **Amdahl's Law** fame
wrote forty-one years earlier.
In his often-quoted scientific paper
of April 1967, **Gene Amdahl** wrote that
supercomputer scientists
should forget using eight
processors.
And should forget using them
to increase the speed of supercomputers.
Seymour Cray who designed 70 percent
of the vector processing supercomputers
of the 1980s
did not use up to two
vector processing units.
And did not use them to increase the speed
of his supercomputers.

Steve Jobs of Apple Computers
agreed that less than eight
processors
could increase the speed of personal
computers.
I—**Philip Emeagwali**—
invented how and why
the slowest 65,536
tightly-coupled processors
in the world
could increase the speeds
of computers and supercomputers.
To invent
the modern supercomputer
—that computes a million things
at once,
instead of computing just one thing
at a time—
is to see parallel processing

compress the **time-to-solution**
of the toughest problems
in extreme-scale computational physics.
In 1989, my invention
of the parallel processing supercomputer
made the news headlines.
It was written that I discovered
how to compress the **time-to-solution**
from 65,536 days, or **180 years**,
to just **one day** across
the slowest 65,536 processors.
Contrary to the opinions
held by the likes of **Gene Amdahl**
of the mainframe computer world,
held by the likes of **Seymour Cray**
of the supercomputer world,
and held by the likes of **Steve Jobs**
of the computer world
and contrary to their opinions

that parallel processing
will forever remain
a huge waste of everybody's time,
I—**Philip Emeagwali**—discovered that
parallel processing
is not
a huge waste of everybody's time.
To invent the massively parallel processing
supercomputer
is to record previously **unrecorded**
speeds in supercomputing
and to do so by supercomputing
a million things **at once**
and supercomputing them
when the likes of **Steve Jobs**
of the computer world
and the likes of **Seymour Cray**
of the supercomputer world
argued that it will forever remain

impossible to do so
and to efficiently compute eight things
at once,
and argued that it will forever remain
impossible to parallel process **across**
eight processors.

34.2 How I Invented the Modern Supercomputer

34.2.1 Beauty and Serendipity

I'm **Philip Emeagwali**.
I invented a new internet
that is a new global network of
65,536 tightly-coupled processors
that shared nothing with each other
and that were already available
in the market anyway.

And I invented
how to use that new internet
to make modern computers **faster**
and make new supercomputers the **fastest**.
Please allow me to take
a half-century **retrospective** look
on my early years in supercomputing
that began on June 20, 1974
in Corvallis, Oregon, United States.
I will focus only on my supercomputing
during the **1970s and the ‘80s**
and on the role of **beauty** and **serendipity**
in my experimental discovery
of how and why
parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**.

34.2.2 Beauty and Serendipity in Supercomputing

I'm often asked
to describe the role **beauty**
and **serendipity** played
in my invention
of the massively parallel processing
supercomputer
that is a new internet, *de facto*.
The beauty
of my multi-colored illustrations
of the **cube**
in the sixteenth dimension
that was tightly circumscribed
by the **sphere**
in the sixteenth dimension
inspired me to use
the **hypercube** and the **hypersphere**
as my metaphors
for the two
massively parallel processing
supercomputers
that I invented

and invented as new internets.
The **beauty** of geometrical objects,
such as the **cube**, the **sphere**,
and the **truncated icosahedron**
inspired me to keep moving forward
in my quest
for the massively parallel processing
Cosmic Supercomputer.
Serendipity or luck
played no role
in my invention
of the massively parallel processing
supercomputer.
It was an invention
that was 15-years in the making.
The earliest illustrations
of my new internet
as a new global network of
processors
were described as **beautiful**.
Those illustrations

are still reproduced—but are reproduced
without attribution to
Philip Emeagwali—
and are reproduced
by the print and the television media.
Serendipity played no role
in my **quest for the fastest computation**.
My invention
of a system of coupled, non-linear,
time-dependent, and state-of-the-art
partial differential equations
of the modern calculus
is not due to luck.
My invention of how to solve
the algebraic approximations
of those partial differential equations
is not due to luck.
My invention
of how to solve them **faster**
and do so on a new computer

is not due to luck.
My invention
of how to solve them **fastest**
and how to do so **across**
a new internet
that is a new supercomputer
is not due to luck.
My invention
of how to solve them **faster**
and how to do so by a factor of
64 binary thousand
is not due to luck.
All those inventions
at the frontiers of knowledge
of the modern calculus,
of the extreme-scale
computational physics,
and of the massively parallel processing
supercomputer
cannot be attributed to **serendipity**
or to **luck**.

In the mathematical
and computational sciences,
the invention of a new calculus
or a new algebra or a new algorithm
is first made
inside the mind of the inventor
and second seen with the eyes
of the inventor.
So, my invention
of how to massively parallel process
and how to process
and solve the toughest problems
arising in calculus and physics
and how to solve those problems **across**
my new global network of processors
that's a new internet
wasn't **serendipitous.**
Nor was that discovery
due to an act of luck.
I knew, ***a priori***,
that the massively parallel processing

supercomputer
that I invented
will be the world's fastest.
And I knew, ***a posteriori***, that my
massively parallel processing
supercomputer
was, ***de facto***, a new internet
that is a small copy of the Internet.

34.2.3 Origin Story of Philip Emeagwali

Please allow me to describe
my small copy of the Internet
and to describe that small internet
in prose,
instead of describing that small internet
as an illustration.
In the summer of 1974,
I visualized 64,000 computers
that encircled the Earth.
I visualized each computer

as equidistant
from its nearest-neighboring computers,
or the computers closest to it.
I visualized each computer
to be assigned to an area of
three thousand square miles.
That is, **I visualized**,
those 64,000 computers
to be evenly distributed
around planet Earth.
That **visualization** was the origin story
of my small copy of the Internet.
I first programmed
a scalar processing supercomputer
on June 20, 1974
in Corvallis, Oregon, United States.
Back in 1974, when I was programming
sequential processing supercomputers
that processed only one thing **at a time**,
I had a very vague idea
of the future massively

parallel processing supercomputer
that will do many things **at once**
that I was **visualizing**.
A decade later,
I had intellectually matured
and came of age
as a research mathematician
and as a research physicist.
I had invented
how to program a new internet
that is a new global network of
64 binary thousand
tightly-coupled processors
that shared nothing with each other,
or a new global network of
as many computers.
Fast forward another five years,
to the Fourth of July 1989,
I had intellectually matured
and had come of age
as a supercomputer scientist

that was the lone wolf
in Los Alamos, New Mexico,
United States
in the **unchartered territory**
of the massively
parallel processing supercomputer.
After 15-years
of supercomputer research,
it made the news headlines
in 1989
that I—**Philip Emeagwali**—
had experimentally invented
what I had theoretically invented
across the previous decade and half.

34.2.4 Philip Emeagwali Internet in Prose

Each of my processor
was my **metaphor** for a tiny computer
within my new global network of
64 binary thousand processors

that was a new internet
that I visualized back in 1974.
The reason the illustrations
of my new global network of
processors
were beautiful
to the eyes and mind
was that each processor
within my small copy of the Internet
was **equidistant**
from its nearest neighboring
processors
that shared nothing with each other.
In practical massively parallel processing
supercomputing,
my new global network of
two-raised-to-power sixteen
processors
must be **soldered** onto
two-dimensional motherboards.
Therefore, their **equidistance**

occurs only in a **topological sense**,
instead of in a **geometrical sense**.
In the early 1980s,
I did a back-of-the-envelope calculation.
Namely, I theorized the CPUs,
or the processors,
as positioned
at the two-raised-to-power sixteen,
or 65,536, vertices of the **cube**
that were tightly circumscribed
by the surface of the **sphere**
and I theorized that **hypercube**
in my imaginary sixteen dimensional
hyperspace.
My multi-colored illustrations
of the pathways of that cube
in the sixteenth dimension
were **beautiful**.
But they were not drawn
serendipitously.
I theorized the 64 binary thousand

CPUs
as having a one-to-one correspondence
to 64 binary thousand
initial-boundary value problems
of modern calculus
or to 64 binary thousand
computation-intensive problems
of computational fluid dynamics
of physics.
I reached my 64 binary thousand
central processing units
by emailing each processor
from one of my 64 binary thousand
processors
that were each uniquely identified
by a sixteen-bit long binary number
that is a unique string of
sixteen zeroes and ones.
What I've just described
was the toughest problem
in modern supercomputing.

For that reason, I can't count on
lady luck or serendipity
to help me to synchronously send
and to simultaneously receive
those unique 64 binary thousand emails.
I had to know
my 65,536
initial-boundary value problems
of modern calculus
that I was sending
and receiving **across**
my 65,536 tightly-coupled processors
that shared nothing with each other.
That was how
I invented
the massively parallel processing
supercomputer.
That discovery changed the way
we looked at the fastest computers.
That invention
enabled the modern high-performance

supercomputer
to be powered by
ten million commonly available
processors.
That invention
made it possible to reduce
30,000 years, or about ten million days,
of **time-to-solution**
to just one day of **time-to-solution**.

34.2.5 The Beauty in Parallel Processing

In the 1970s and ‘80s,
the leaders of thought
for the fields of sequential processing
and vector processing supercomputing
believed that:
parallel processing will forever remain
a beautiful theory
that lacked experimental confirmation.

Because it was then impossible
to solve the toughest problems
arising in extreme-scale
computational physics—such as,
the excruciatingly detailed
general circulation climate models
that must be used
in the shortest time possible—
and used to solve such grand challenges
across
a new internet
that is a new global network of
64 binary thousand processors,
and because it was then impossible
to solve such grand challenge problems
the few massively parallel processing
supercomputer-hopefuls and centers
in the world
were abandoned like a ghost town.
In the 1980s,

I was the only parallel programmer
that programmed fulltime
and I did so at the **farthest frontier**
of the massively parallel processing
supercomputer.
I visualized my new supercomputer
as my new internet
that is a new global network of
64 binary thousand
commonly available processors.
While supercomputing as a **lone wolf**
massively parallel programmer
in Los Alamos, New Mexico,
United States,
I invented
how to synchronously send and receive
each of my 64 binary thousand
computational fluid dynamics codes
of physics
that each is an initial-boundary value

problem
of modern calculus.
Each initial-boundary value problem
is a system of coupled, non-linear,
time-dependent, and state-of-the-art
partial differential equations
and the **associated constraints**
on that system.
Those **constraints**
are called initial and boundary
conditions.
I invented
how to compute a solution
to any initial-boundary value problem
—such as the solution to the classic wave
and heat equations
of mathematical physics
that has the appropriate initial
and boundary conditions specified.
I invented

how to solve a system of coupled,
non-linear, time-dependent, and state-of-the-art
partial differential equations
of modern calculus
and I discovered how to solve them
across
my new internet
that is a new global network of
65,536 tightly-coupled processors
with each processor
operating its own operating system
and with each processor
having its own dedicated memory
that shared nothing with each other.
I invented
how to solve the toughest problems
arising in calculus
and how to solve such problems
simultaneously

and how to solve such problems
while ensuring that my
massively parallel processed
supercomputer solutions
satisfy their specified initial
and boundary conditions
and their governing
partial differential equations.
I invented
how to synchronously
send and receive
their initial and boundary conditions,
or **companion data**,
and how to send them **across**
my new internet
that is a new global network of
two-raised-to-power sixteen, or 65,536,
already-available processors
that was outlined by sixteen times
two-raised-to-power-sixteen,

or 1,048,576,
commodity bi-directional email wires.
Those email wires
corresponded to the bi-directional edges
of the cube
in my imaginary sixteen dimensional
universe.
Each of my two-raised-to-power-sixteen
CPUs
—the **acronym** for
central processing units, or processors—
emailed their computed answers
to the computational fluid dynamics code
that I assigned to that CPU
and emailed them
to my uniquely identified
nearest-neighboring CPUs
that's associated
with physical domains
that are nearest-neighbors

to the physical domain
that the sending CPU represents.
Sending and receiving those emails
were impossible in the 1970s and ‘80s
and first became possible
when I sent and received them
as 65,536 simultaneous emails
that arrived at
65,536 already-available processors
and arrived at 8:15 on the morning
of Tuesday the Fourth of July 1989
in Los Alamos, New Mexico,
United States.
In the early 1980s, when I proposed to
simultaneously send and receive
those 65,536 emails,
it evoked laughter.
At that time, most vector processing
supercomputer scientists
had never sent or received an email.

For that reason, I was described as
a “lunatic”
that worked on the
“fringes of knowledge”
and I was dismissed
from my research group.
To synchronously send
65,536 emails
required that I know my new internet
very, very well
and know my new internet
forward and **backward**
and even **sideways**.
Serendipity, or luck, did not help me to
invent
that massively parallel processing,
or doing many things **at once**,
makes the modern computer
faster
and makes the modern supercomputer

fastest.
I invented
the massively parallel processing
supercomputer
and I invented the technology
on the Fourth of July 1989
and I invented it at the time
Steve Jobs and Seymour Cray
—the two leaders of thought
in the world of computing
and supercomputing, respectively
—were loudly voicing their opinions
that parallel processing
will forever remain
a huge waste of everybody's time.
In the 1970s and '80s,
to massively parallel process
the toughest problem
that arose in extreme-scale
computational physics

and to parallel process that problem
across a new internet
that is a new global network
of processors
was deemed impossible.
It was easier to send a man to the moon
than to harness
the total processing power
of the massively parallel processing
supercomputer.
To be a lone wolf programmer
of the most massively parallel processing
supercomputer
ever built
was to have the audacity
to look God in the face.
Two decades after I had
invented
how to massively parallel process
and how to compute **across**

a new internet
that is a new global network of
65,536 already-available processors,
Steve Jobs
and his team of programmers
at **Apple Computers**
tried parallel processing
across merely eight processors.
Steve Jobs gave up
dismissed and ridiculed
parallel processing.
Steve Jobs
mocked parallel processing
as a huge waste of everybody's time.
In the June 10, 2008 issue
of the *New York Times*,
Steve Jobs
was quoted as telling
Apple's Worldwide Developers

that [And I quote, **Steve Jobs**]:

PROP ALERT
“The way the processor industry
is going
is to add more and more cores,
but nobody knows
how to program
those things,”

Steve Jobs continued:

“**I mean, two, yeah;**
four, not really;
eight, forget it.”
[End of quote]

In late June 1990,
Steve Jobs read about **Philip Emeagwali**
in the June 20, 1990 issue

of *The Wall Street Journal*
and made a telephone call
to **Philip Emeagwali**
and did so because **Steve Jobs**
discovered that
the impossible-to-compute
by doing only one thing **at a time**
may, in fact,
be possible-to-parallel-compute
by doing many things **at once**.
In 1989,
it made the news headlines
that an African supercomputer wizard
in the United States
has invented
how to massively parallel compute
and how to compute simultaneously
and how to compute on
two-raised-to-power sixteen processors
and how to compute
and communicate synchronously

and do both **across** sixteen times
two-raised-to-power sixteen email wires
that, in turn, connected those processors
as one seamless, cohesive whole unit
that is a new internet
and that is a new supercomputer
and a new computer.
I—**Philip Emeagwali**—was that
African supercomputer scientist
who experimentally discovered that
parallel processing
was an all or nothing affair.
The explanation for this **all or nothing** affair
is that at the same compute cycle,
I locked all my 65,536
processors
that shared nothing with each other
and I **synchronously** instructed
each processor, or CPU,
to send and receive

CPU-to-CPU email messages
and to receive them
65,536 times faster than
your everyday **person-to-person**
email message.
So, either I—**Philip Emeagwali**—
was executing the world's fastest
floating-point arithmetical computations
and executing
the correspondingly fastest
email communication,
or I was executing nothing.
I programmed all
65,536 tightly-coupled processors
to simultaneously compute
so that the toughest problem
of computational physics
that had a **time-to-solution**
of 65,536 days
on one processor,
or 180 years

on one computer,
now had a reduced **time-to-solution**
of only one day
across a new internet
that is a new global network of
65,536 tightly-coupled processors
that shared nothing with each other
and that is a new supercomputer
and a new computer.
When I invented
how to compute at a world record speed
and how to compute **across**
a new internet
that is a new supercomputer *de facto*,
I felt like I was jolted
by a bolt of electricity.
I understood that my world record
supercomputer speed
was the same speed
that I theorized for a decade,
namely, that forty-seven thousand

three hundred and three [47,303]
calculations per second
per processor
that's totaled **across**
65,536 tightly-coupled processors
that shared nothing with each other
would yield that Fourth of July 1989
world record
of **3.1 billion** calculations per second.
I remember the afternoon
that I experimentally discovered
that **quantum leap**
in supercomputer speed.
I leapfrogged **across** the frontier
of the vector processing supercomputer
and I leapfrogged
by a factor of 65,536
and leapfrogged
into the unknown world
of the massively parallel processing
supercomputer.

I was shivering.
I couldn't continue working.
I went home and called my wife, Dale,
at her research laboratory
at the local Medical School.
"What's wrong?" Dale asked me.
"It worked! It worked!! It worked!!!"
I said with trepidation.
Dale knew that
I had experimentally invented
the massively parallel processing
supercomputer
that I had theoretically invented,
a decade earlier.
At the Eureka Moment
that I invented
the modern supercomputer,
I felt like the magician
that turned fiction into fact.
For my research years,
onward of the second Tuesday

in June 1978,
my wife, Dale,
was the only research scientist
that understood
my supercomputer vision
and my parallel processing motivation.
Dale knew that I had
invented
how to massively parallel process
180 years to one day.
Later, that invention
of the massively parallel processing
supercomputer
was reported in the June 20, 1990 issue
of the *Wall Street Journal*.
The news media reported that
a lone wolf
African supercomputer wizard
in Los Alamos, New Mexico,
United States
has experimentally discovered

a **paradigm shift**,
or a change in the way we look at
the computer and the supercomputer.
Before that paradigm shift,
we looked at the computer
as powered by only one processor
that was not a member
of an ensemble of processors.
After that paradigm shift,
we looked at the computer
as massively parallel computing
and supercomputing
with up to ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
processors
that were already available
in the market anyway.
The pre-cursor
to the modern supercomputer
that I invented

on the Fourth of July 1989
in Los Alamos, New Mexico,
United States,
occupied the space of a ping-pong table.
The fastest
of the modern supercomputers
of today
occupies the space of a football field.
That massively parallel processing
supercomputer
that I invented
on the Fourth of July 1989
opened the door
to the modern supercomputer
that now massively parallel processes
many things **at once**.

34.3 Father of the Modern Supercomputer

34.3.1 Contributions of Philip Emeagwali to the Supercomputer

On June 20, 1974,
at 1800 SW Campus Way,
Corvallis, Oregon, United States,
I began programming
the [quote unquote]

"first supercomputer."

It was called the first supercomputer
because it was the first computer
that could execute
one million instructions per second.
Back in the 1970s,
my supercomputing vision
was to invent
how to execute the fastest computations
ever

and how to execute them **across**
a new internet
that was my new global network
of processors
that were already available
in the market anyway.
My quest for the fastest
massively parallel processing
supercomputer
was all consuming.
It was a supercomputing quest
to record a computing speed
that was previously unrecorded.
The inventor
embarked upon a hero's quest
to hear something
that was previously unheard;
to see something
that was previously unseen;
and to understand something
that was previously misunderstood.

My Eureka Moment in supercomputing
that was recorded
in the June 20, 1990 issue
of *The Wall Street Journal*
was the high point
of my scientific journey
that began exactly sixteen years earlier
and began on June 20, 1974.
For me—**Philip Emeagwali**—
my highest high came from
finding the loudest voice
in the world of physics,
namely, the Second Law of Motion.
My highest high came from
finding the clearest vision
in the world of supercomputing,
namely, the fastest speed
in supercomputing that I executed
across a new internet
that I visualized
as a new global network of

64 binary thousand
processors.
My highest high came from
finding the deepest wisdom
and from gaining
the greatest understanding
of how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**.

34.3.2 My Biggest Obstacle in Life

The turning point
in my journey to the frontier of
supercomputer knowledge
occurred twenty years earlier
when I was a 14-year-old soldier
on the Biafran side
of the Nigeria-Biafra Civil War.
In July 1969, I was **conscripted**

and sent to the Oguta War Front,
Igbo Land, Biafra, West Africa.
I was a fourteen-year-old soldier
at the Oguta War Front
of the Nigeria Biafra Civil War.
That war
turned my ancestral homeland
into Africa's bloodiest battlefield.
I arrived at the Oguta War Front
a few days after
500 Biafran soldiers
fell dead on the ground.
Five hundred soldiers fell
as if they were dry leaves.
I was conscripted
to replace
one of these 500 men.
At Oguta War Front,
they were more guns
than pens.
It was at Oguta War Front

that I first heard the name
Colonel Olusegun Obasanjo,
the future three-term president
of Nigeria.
At Oguta War Front,
I was informed that
Colonel Olusegun Obasanjo
is the new commander
of the much better and heavily armed
Nigerian soldiers.
My mathematical journey
continued from
the bloody Oguta War Front
in Biafra
that was described
as the bloodiest war
in the history of Africa
to the frontiers of extreme-scaled
computational physics
that was described as the
toughest problem

in physics and supercomputing.
At that frontier of supercomputing,
nuclear bombs were **digitally exploded**
via physics-based simulations **across**
64 binary thousand
already-available processors
that were **wired together**
by one binary million
regular and short email wires
that were equal distances
apart.
At the end of my twenty-year long
journey
to the frontier of human knowledge
that began in July 1969
and began in Oguta War Front,
in Igbo Land, Biafra,
I felt like the sojourner in Igbo folktales
that arrived at the Land of the Spirits
where he wrestled with a
three-headed dog.

34.3.3 How Philip Emeagwali Invented a New Supercomputer

The quest for a never-before-recorded
supercomputer speed record
is what makes us human.
That new speed record
is an intrinsic part
of our human progress.
That new speed record
enables us to forever
reinvent our technologies
and ourselves.
That new speed record
is our collective hero's journey
to the *terra incognita*
of technological knowledge.
From June 20, 1974 at age nineteen
in Corvallis, Oregon, United States

to the Fourth of July 1989
in Los Alamos, New Mexico,
United States,
I was totally committed
to the quest
for the fastest supercomputer
that could be experimentally discovered
through the massively
parallel processing
of physics-based and calculus-based
supercomputer codes
that must arise
from any excruciatingly-detailed
simulation
in extreme-scaled computational physics.
Every line of my supercomputer code
encoded
the intellect and knowledge
of mathematicians and physicists
whose names have been lost
in the mist of time.

To record the fastest recorded
supercomputer speeds
and to do so **across** a new internet
demanded that
I visualized my emails
as exploding **across** my new internet
that is a new global network of
64 binary thousand
tightly-coupled processors
that shared nothing with each other.
I visualized emails
as exploding
like bullets out of my eyes.
When I arrived
at the frontier of knowledge
in massively parallel processing
supercomputing,
I felt like I had accosted
the Medusa
that was guarding
the Tree of Knowledge.

I, in part, defined that terra incognita
by my email address space
that consisted of my unique arrangement
of one binary million zeroes and ones.
Metaphorically speaking,
that Medusa had to be slayed
before I crossed into the *terra incognita*,
or the uncharted territory
of supercomputer knowledge.
The discovery
is the act of seeing something
previously unseen.
But we only see
the discovery
with our intellectual eyes,
not with our biological eyes.
That new object, that is the discovery,
may be the **fastest** of its kind,
such as the **fastest** aircraft
or the **fastest** singular processor
or the **fastest** singular internet

that is a new
global network of processors.
Or it may be a new person
that ran the **fastest**
one hundred yard race.
Or it may be
a new supercomputer scientist
that recorded the **fastest**
computer calculation
ever recorded.
For the record, **Philip Emeagwali**
recorded the experimental discovery
of the massively parallel processing
supercomputer.
For the record, **Philip Emeagwali**
made that experimental discovery
when parallel processing
was presumed to be untestable
or even wrong.
Yet, theorized parallel processing
was in the air for the 43 years

prior to my experimental discovery
of the massively parallel processing
supercomputer
that occurred
on the Fourth of July 1989.
The January 11, 1946 issue
of the *New York Times*
mentioned parallel processing
as science fiction
and as 100 computers
that could forecast the weather
all over the world
and that

[quote]
"the United States
would be divided into 'blocks'
penetrating into the stratosphere."
[unquote]

I—**Philip Emeagwali**—
read that *New York Times* article
and made the leap of imagination
From the 100 computers
that were theorized in 1946
to the 64 binary thousand
processors
that I theorized 28 years later
and that I experimentally programmed
43 years later.
Since that *New York Times* article,
parallel processing was scorned,
ridiculed, and rejected
as a beautiful theory
that lacked experimental confirmation.
Parallel processing
was experimentally confirmed
by I—**Philip Emeagwali**—and confirmed
on the Fourth of July 1989
in Los Alamos, New Mexico,
United States.

On that date, parallel processing
was verified by experiment
that I executed **across**
a new internet
that was outlined by sixteen times
two-raised-to-power-sixteen,
or 1,048,576 bi-directional email wires
that married
two-raised-to-power-sixteen processors
together
as one seamless, cohesive supercomputer
that was the precursor
to the modern supercomputer
as well as the modern computer.
The Fourth of July 1989
was the day
parallel processing passed
the merciless test of the experiment
that I conducted **across**
my new internet
that is a new global network

of processors
and that is a new supercomputer
and a new computer.
I made headlines in major U.S.
newspapers
because I provided
the lockdown evidence
that the massively parallel processing
supercomputer
can be used to solve
the toughest problems
arising in extreme-scale
computational physics.

34.3.4 More Information

Dălụ́'nụ̀ (DAH-LOO nooh)
Afam mụ bu Chukwurah Philip Emeagwali.
Abum onye onicha.
Bia ga fum na emeagwali dot com
Ka omesia.