54 Father of the Fastest Computers—Part 5 of 10

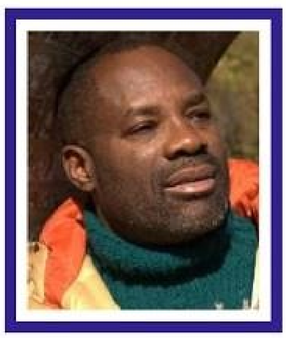

Philip Emeagwali
info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture 170926
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video:
https://www.youtube.com/watch?v=kG7Hn40TANQ&t=2407s

https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

54.1 A Messenger From God

54.1.1 A Messenger From God

The scientific discovery
is the **nothing**ness
from which new knowledge sprang.
That new knowledge
makes the discoverer
a messenger from God.
For me, **Philip Emeagwali**,
my quest for the fastest supercomputer
that can do one million things
at once
was a very difficult
and a **never-ending** journey.
My quest began on a
sequential processing supercomputer
that was at
1800 SW Campus Way,
Corvallis, Oregon, United States.
My quest for the fastest supercomputer

began on the morning of
Thursday June 20, 1974.
Three weeks after I began **that quest**
for the fastest supercomputer,
the photo of a nineteen-year-old
named "**Philip Emeagwali"**
appeared on the front page
of an Oregonian-American newspaper.
Two years before I began that quest
for the fastest supercomputer,
the name "**Philip Emeagwali**"
was mentioned in the science column
of the *Daily Times* of Nigeria.
For a half century, the *Daily Times*
was Nigeria's only national newspaper.
One year after the name "**Philip Emeagwali**"
was mentioned in the *Daily Times* of Nigeria,
I won a scholarship to the United States.
On the sixteenth anniversary

of when I began my quest
for the fastest supercomputer,
and on June 20, 1990,
The Wall Street Journal
reported that **Philip Emeagwali**
has experimentally discovered
how and why parallel processing
makes modern computers faster
and makes the new supercomputer
the fastest,
or discovered a new supercomputer.

54.1.2 School Reports on Philip Emeagwali

Since 1989, school children
in the United States, United Kingdom,
and Canada
are asked to do a school report
on the contributions of **Philip Emeagwali**
to the development of the modern

supercomputer.\
Back in 1989, it made the news headlines\
that a lone wolf\
African supercomputer wizard\
in Los Alamos, New Mexico, United States,\
made an experimental discovery\
of how to solve\
65,536 mathematical problems **at once**.\
That discovery is called\
“massively parallel processing.”\
I experimentally discovered that\
the fastest arithmetical speeds\
in supercomputing\
can always be recorded\
with massively parallel processing\
technology\
and used to solve the largest-scale\
problems in algebra\
those algebraic problems, in turn,\
arose from\
the most abstract problems\
in calculus.

54.1.3 Forging a Path From Oregon to New Mexico

For the sixteen years
onward of June 20, 1974,
I conducted my scientific
and technological research
as a **lone wolf** supercomputer programmer.
I programmed supercomputers
at the Computer Center
at 1800 SW Campus Way,
Corvallis, Oregon, United States.
Each morning, at 7 a.m.,
I parked my red, two-speed bicycle
in front of the nearby mathematics building
that was named Kidder Hall
that was at 2000 SW Campus Way,
Corvallis, Oregon.
But I came to the public attention
—sixteen years later—
and for my experimental discovery

that occurred in **Los Alamos**, New Mexico, 2646
United States.
Los Alamos is a small town
that is the birthplace of the atomic bomb.
But to research supercomputer scientists,
Los Alamos is the capital of
supercomputing.
Let's put supercomputer simulation
of nuclear explosions at **Los Alamos**,
in real world context.
The first atomic bomb was detonated
by the United States
on **August 9, 1945**
over Hiroshima, Japan.
The second atomic bomb was detonated
three days later
over Nagasaki, Japan.
Following world-wide outcry
to ban nuclear use and testing,
the fastest supercomputer
is used to quietly simulate the detonation

of atomic bombs. 2647
Paradoxically,
the simulated nuclear explosion
is more real
and has greater degrees of freedom
and can be repeated again and again
than the actual nuclear testing
that is extremely limited.
So supercomputers are secret tools
used to simulate nuclear explosions
at a densely populated mega city
and to evade
the spirit of the Comprehensive Nuclear Test Ban Treaty.
The **Los Alamos** National Laboratory,
in **Los Alamos**, New Mexico, United States
was where I experimentally discovered
how parallel processing
becomes the computational engine
that makes modern computers **faster**
and makes the new supercomputer

the **fastest,**
namely, **the Philip Emeagwali formula**
that then United States President
Bill Clinton described
in his White House speech of
August 26, 2000..
I experimentally discovered that
parallel processing
that makes the **fastest** computations
could arise from the bowels
of a new internet.
Faster computers
could arise from perhaps,
quantum computers that will solve
travelling salesman problem.
But, in my opinion,
the quantum computer
will have limited use
for its seemingly unlimited
supercomputing speeds.
That is, the quantum supercomputer

could never be harnessed
to solve any of the twenty
grand challenge problems
in supercomputing.

54.2 Making the Impossible-to-Solve Possible-to-Solve

54.2.1 Exiled to Unchartered Territory

The reason I invented parallel processing alone
was that in the 1970s, it was impossible
for a black African supercomputer scientist
to join an all-white supercomputing team,
such as the team
in **Los Alamos** National Laboratory
that was simulating nuclear explosions.
As a black African-born supercomputer
scientist, I felt like I was in **exile**
wherever I was.

2650

I'm in **exile** in the United States.
I was in **exile** in Nigeria (Africa).
I was in **exile**
in the then uncharted territory
of the massively parallel supercomputer.
Algebra is the way
the computer experience calculus.
Calculus is the way
mathematics experience physics.
And physics is the way
animals experience the universe.
At the U.S. national laboratories,
operated by the departments of energy
and NASA
each member of the large multidisciplinary
supercomputing team
brought his or her unique knowledge
of physics or mathematics or computing
to the team.
A member could bring expertise
from the frontier of knowledge

of biology or chemistry or physics, 2651
such as her mastery
of the laws of physics.
Or a member could bring expertise
from the frontier of knowledge
of number theory or large-scale algebra
or advanced calculus,
such as her mastery
of the partial differential equations
of calculus
and her mastery of how to solve
the largest system of equations
of algebra,
or her mastery
of how to assembly code
the most computation-intensive
algebraic kernels.
Or that member could bring expertise
from the frontier of knowledge
of supercomputing
such as his mastery of

massively parallel processing.
I mastered the parallel processing
technology **at the time**
it was ahead **of the times.**
I mastered parallel processing theoretically
before I experimentally discovered it.
To discover parallel processing
required both theories and experiments.
Because massively parallel processing
was a grand challenge problem
that **transcends** mathematics,
inventing how to solve
the most computation-intensive problems
in computational physics
required a polymath
that is a jack-of-all-sciences
and not merely a mathematician
that has expertise
only in algebra or calculus.
As the lone wolf
research internet scientist

of the decades of the 1970s and ‘80s, 2653
I had to be a polymath
or a jack-of-many-sciences
that sent and received
two-raised-to-power sixteen emails
to and from sixteen-bit long
email addresses
that were around a new internet
that I visualized
as a new global network of
sixteen times
two-raised-to-power sixteen
short, regular, and bi-directional
email wires.
That polymath
had to **be at the frontier** of knowledge
of physics,
and possess the command
and the mastery
of computational fluid dynamics.
That polymath

had to **be at the frontiers** of knowledge 2654
of mathematics,
and possess the command
and the mastery
of how to formulate
and then solve
the partial differential equations
of calculus
and how to code and parallel process
the supercomputer solutions
of the largest-scaled systems of equations
of algebra.
That polymath
had to **be at the frontier** of knowledge
of supercomputing,
and possess the command
and the mastery
of how to parallel process
across a new internet
that is a new global network of
two-raised-to-power sixteen processors

that were equal distances **apart**. 2655
It took me sixteen years of
multidisciplinary training
to **be at the frontiers**
of that uncharted territory of knowledge.
To arrive at the frontiers of scientific
and technological knowledge
required that
I attended 500 weekly research seminars
and attend them
during the decade of the 1970s and ‘80s.
Each research seminar
was given by a foremost research scientist
that was at the frontier
of science and technology.
It took me sixteen years
to become that triple threat
in physics, mathematics,
and supercomputing.
Most importantly,
I had to be forward-looking

and be way-out in unfamiliar territory 2656
of the technology of parallel processing.
I was the only research
internet scientist
of the early 1980s
that gave research lectures
on how to harness
massively parallel processing
and use the technology
to solve the toughest problems
in computational physics.
In the early 1980s,
I was ridiculed
and I was warned that
I would fall through the cracks
in the massively parallel processing
supercomputer.
When I announced
my experimental discovery
of massively parallel processing,
everybody warned that

I had made a mistake
but everybody was mistaken.

54.2.2 Seeing the Unseen

What kept me moving forward
was a **back-of-the-envelope**,
theoretical calculation
—on how to compress time-to-solution
in extreme-scale computational physics.
that I did in the 1970s.
I mathematically discovered
that processors that could only calculate
forty-seven thousand
three hundred and three [47,303]
floating-point arithmetical operations
per second
can be **integrated across**
a small internet.
I mathematically invented
that new internet

as a new global network of 2658
65,536
cooperating processors
that are equal distances **apart**.
I mathematically invented
that new internet
to be faster than any supercomputer
that computes less than
3.1 billion calculations per second.
It was my experimental re-confirmation
of my mathematical discoveries
that made the cover stories
in the world of mathematics,
including the May 1990 issue
of *SIAM News*.
The *SIAM News*
is the leading mathematics publication.
That invention
of how to solve a million, or a billion,
mathematical problems
and solve them **at once**

and across as many processors 2659
became my contribution
to mathematical knowledge.
In the 1980s,
I was the lone fulltime programmer
of the only
massively parallel supercomputer
that's an ensemble of
65,536
processors.
It made the news headlines
that I—**Philip Emeagwali**—invented
how to harness those computing units
and invented how to harness them
to compute simultaneously within
and **across**
65,536 tightly-coupled processors
and compute together
as one cohesive whole unit,
or massively communicate and compute
in parallel.

That is, in the 1970s and '80s, 2660
my technological quest was for
massively parallel processing across
a new internet
that is a new supercomputer
de facto.
That is, I was searching for an actual
parallel processing supercomputer
while theorists were theorizing their
way through
a non-existent
parallel processing supercomputer.
A theory is not a discovery.
A theory is an idea
that is not positively true.
In the 1960s, '70s, and '80s,
parallel processing
was ridiculed as a beautiful theory
that lacked experimental confirmation.
And you cannot experimentally discover
the fastest speeds in supercomputing

by merely and only theorizing about it. 2661
In 1989, I experimentally confirmed
that massively parallel processing
can enable the large-scale
computational mathematical physicist
to compress her **time-to-solution**
of the toughest problems
in arithmetic, algebra, calculus,
and physics,
and to compress her **times-to-solution**
from 180 years to just one day.
The experimental discoveries
in massively parallel processing
that followed in the subsequent
three decades
was a series of cleanups and refinements
and re-discoveries.
Three decades later,
my experimental discovery
of massively parallel processing
enabled China

to copy massively parallel processing technology 2662
and use it to massively compress
their **times-to-solution**
for solving the most
computation-intensive
mathematical problems
and compress them from
thirty thousand [30,000] years
to just one day.
That was how China
—that did not invent massively
parallel processing—
has massively parallel processed its way
to the world's fastest supercomputers.

54.3 A New Way of Computing

54.3.1 Turning Point of 1973

For me, **Philip Emeagwali**, 2663
I entered into the unknown world
of massively parallel supercomputing
and entered at a time parallel processing
was ridiculed
as a huge waste of everybody's time.
My turning point was on
September 10, 1973
when I received a four-year scholarship
to study mathematics
in Oregon, United States.
To put things in perspective,
back in 1973,
it was rare for a teenager in Igbo Land
[Nigeria, Africa]
to get a scholarship
and get it directly
from the United States.
In 1973,
fewer than a hundred scholarships,
or one in a million,

were given by the Nigerian government. 2664
That was the reason
I was the only teenager
at the airport, in Ikeja (Nigeria)
in the late afternoon of
Saturday March 23, 1974.
I was the fourth Nigerian
to live in Monmouth, Oregon,
in the Pacific Northwest Region
of the United States.
So my four-year scholarship
to Monmouth, Oregon
was what set me on the right path
and set me on my personal quest
for parallel processing,
the technology that now makes
modern computers **faster**
and makes the new supercomputer
the **fastest**.

54.3.2 Right Place, Right Time

2665

Recording previously unrecorded
supercomputer speeds
is the only objective measure
of the computer inventor's contributions
to the development of the computer.
Only those computer inventors
that worked directly
with the potential **fastest**
supercomputers
can discover that parallel processing
will make the supercomputer
fastest.
I had to be at the right place
—which was **Corvallis**, Oregon
(United States),
instead of **Onitsha**, East Central State,
(Nigeria)—
and be at that unchartered territory
of human knowledge

of fastest computing
to contribute to the development
of the fastest supercomputer.
I had to be in the Computer Center
at 1800 SW Campus Way,
Corvallis (Oregon, United States)
at the right time,
or on June 20, 1974.
I was doing scientific research
as a teenager
both in Nigeria and the United States.
As an aside,
an independent scientific research
that as a 17-year-old
“**Philip Emeagwali**”
of Christ the King College,
Onitsha (Nigeria)
was mentioned in the science column
of the *Daily Times* newspaper
of Nigeria.
And a 19-year-old “**Philip Emeagwali**”

was on the cover of a local newspaper 2667
in Oregon, United States.
On September 10, 1973
and about a year after I was mentioned
in the *Daily Times* newspaper of Nigeria,
I won a four-year scholarship
to train as a mathematician
in the United States.
That four-year scholarship of 1973
was the beginning
of my nineteen-year scholarship
in the United States
that led to my training
as a polymath
that is at home
at the frontiers of knowledge
in physics, mathematics, and computing.
I won that four-year scholarship
of 1973
before I applied for the scholarship
and before I applied for admission

into any American university. 2668
That scholarship
was my breakthrough moment,
that made it possible
for me to start
programming supercomputers
and programming them at a time
there was no supercomputer in Africa.
That scholarship letter
that was dated September 10, 1973
put me at the right place
at the right time, namely,
being a teenager in the 1970s
that was programming supercomputers
and doing so at 1800 SW Campus Way,
Corvallis, Oregon, United States,
instead of being a teenager
in July 1969
at **Oguta War Front**,
of the Nigeria Biafra Civil War.
Five years before I started

programming the sequential processing 2669
supercomputer,
I was a 14-year-old soldier
on the Biafran side
of the **Oguta War Front**.
The opposing Nigerian side
at the **Oguta War Front**
was commanded by
Colonel Olusegun Obasanjo,
who later rose to become
the three-time president of Nigeria.
It’s not a coincidence that
many of the computer pioneers
of today
were teenagers in the early 1970s.
It’s not a coincidence that
Steve Jobs and Bill Gates
were born a few months apart
from my birthdate of August 23, 1954.
Steve Jobs
was also about nineteen years old then

and was then living in nearby Portland, 2670
Oregon, United States,
not in Lagos, Nigeria, West Africa.
And **Bill Gates**
was also about nineteen years old
and was then living in nearby Seattle,
Washington,
United States,
not in Nairobi, Kenya, East Africa.

54.3.3 Exploring Uncharted Territories

In the 1980s,
I was the **lone wolf** full-time
research programmer
of the most massively
parallel processing supercomputer
that was ever built.
I was exploring the uncharted territories
of supercomputing,

called massively parallel processing. 2671
In the 1980s,
twenty-five thousand
research computational scientists
were exploring the charted territories
of computing.
That unknown world
was called vector processing
supercomputing.
The difference
between those 25,000
research physicists,
research mathematicians,
research computer scientists,
and myself
was that they were **re-plowing**
or researching computational physics
or calculus or algebra
or vector processing supercomputers.
They were **re-plowing**
or researching fields of study

that had been **plowed**
They were **re-plowing**
or researching supercomputers
that one million computer scientists
that preceded them
had **plowed** since
the automatic, programmable computer
was invented back in 1946.
Back in the 1970s,
I did not believe in **re-plowing**
the field of sequential processing
supercomputing
that had been plowed
during the preceding three decades.
Nor did I believe in **re-plowing**
the field of vector processing
supercomputing
that had been **plowed**
since the 1960s.
In general, r**e-plowing** the fields of
sequential processing

and vector processing supercomputers
merely yields a factor of two increase
in computing speed.
That doubling of computer speed
that occurred every two years
is called Moore's Law.
I paradigm shifted
and **plowed** the field of
massively parallel processing
supercomputers.
Doing so made the **news headlines**
in 1989.
I became known
as the African supercomputer wizard
in the United States.
I became known
for experimentally discovering
a factor of 65,536-fold increase
in parallel processing
supercomputing speeds.
I became known

for inventing a new supercomputer 2674
that can compress
65,536 days, or 180 years,
of **time-to-solution**
on one processor
and compress it
to just one day of **time-to-solution**
across
my new global network of
65,536 tightly-coupled processors
that shared nothing
that defined, outlined, and powered
that parallel processing supercomputer.
Back in the 1970s,
the leaders of thought
in vector processing supercomputing
rejected my research
in parallel processing supercomputers.
They rejected my theory
that I could solve
65,536 problems **at once**,

and compute faster than those
that solved only one problem **at a time**.
Back in the 1980s,
the **leaders of thought**
in vector processing supercomputing
gave me bad advice
that did me more harm than good.
The **leaders of thought**
in vector processing supercomputing
insisted that I—**Philip Emeagwali**—
sign a written contract
on the supercomputer speedup
that I will experimentally discover
within vector processing.
The **leaders of thought**
in vector processing supercomputing
insisted that I, *a priori*, explain how
I will experimentally discover
that supercomputer speedup
and, most importantly,
experimentally discover that speedup

by disavowing
the parallel processing
supercomputer technology
that was absolutely necessary
for experimentally discovering
my speedup of 65,536.
I wasted the decades
of the 1970s and '80s
fighting against the **resistance**
and the **skepticism**
about massively parallel processing
supercomputers.
That lack of support
forced me to work alone
and in exile
from the supercomputing community.
Those leaders of thought
in vector processing supercomputing
ridiculed parallel processing
as a beautiful theory
that lacked experimental confirmation.

2676

Due to those ridicules,
my experimental confirmation
of parallel processing
was rejected in the 1980s
and rejected
by the vector processing
supercomputing community.
In that decade of the 1980s,
my discovery of parallel processing
was **rotting**
in the **drawers** of prominent
supercomputer scientists.
My supercomputing quest
that began on a sequential processing
supercomputer
and began on June 20, 1974
and began at 1800 SW Campus Way,
Corvallis, Oregon, United States
was to **experimentally confirm**
parallel processing
and to confirm the technology

by solving
65,536 problems
in extreme-scale computational physics
and solving them **at once**,
instead of solving only one
smaller-scaled
computational physics problem
at a time.
To compare **re-plowing**
vector processing supercomputers,
instead of **plowing**
parallel processing supercomputers
is like comparing
re-paving an existing
Lagos-Onitsha Road
to building a new road
between Lagos (Nigeria)
and Johannesburg (South Africa).
The supercomputer of today
will become the computer
of tomorrow.

The massively parallel processing
supercomputer
that I experimentally discovered
back in 1989,
that was ridiculed by Steve Jobs,
had become the computer of today.

54.4 How I Invented the Modern Supercomputer

54.4.1 Changing the Way We Look at Computers

You've changed the way
you look at your computer.
In the 1980s and earlier,
you looked at your computer
as powered by
only one isolated processor
that was not a member
of an ensemble of processors.

Those processors 2680
communicate and compute together
and as one seamless, cohesive
supercomputer.
Today, you look at your computer
as powered by
a hundred, and sometimes,
a thousand processors.
In **mathematical studies**
that I conducted in the 1970s
and in
companion **laboratory experiments**
that I conducted in the 1980s,
I invented
how to use a new internet
that is a global network of
65,536
processors.
I invented
how to use that new internet

to solve the toughest problems 2681
in extreme-scale computational physics.
That invention
set me apart from
Thomas Edison
and his invention of the light bulb.
That invention
set me apart from
Alexander Graham Bell
and his invention of the telephone.
That invention
set me apart from
Albert Einstein
and his discovery of the
theory of relativity.
And that invention
set me apart from
heroes like George Washington Carver
and heroines of science
like Marie Curie
that are scientific role models

and that are studied in American 2682
and European schools.
The reason I am the subject
of school reports
is that I contributed to the development
of the fastest supercomputers.
My contribution
to the development of the computer
is my experimental discovery
of massively parallel processing.
My contribution
to the development of the computer
was as **objective** as my speed increase
of a factor of 65,536.
My contribution
to the development of the computer
is as **measurable**
as the fastest speed
in the history of computing.
My contribution
to the development of the computer

is as **tangible** as a new computer 2683
and as **visible** as a new internet.
That visible internet
is a new global network of
65,536 processors
that are cooperating with each other
and that are tightly-coupled.
The reason my contributions
made the news headlines
in 1989
was that it changed the way
we think about the supercomputer.
After 1989, we think of the
supercomputer
not as solving only one problem
at a time
but as solving one million problems
at once.
A contribution to the development
of the modern supercomputer
that solves many problems

2684

at once
is **concrete** and **durable**
because faster computations
are **objective** and **measurable**.
That contribution
to the fastest supercomputer
will stand as long as
the computer exists.
That contribution
will stand as long as
the internet, or global network of
computers, exists.
That contribution will,
perhaps, **stand as long as**
the **rock of Gibraltar** exists.
In the 1970s, supercomputer scientists
followed an erroneous line of reasoning.
The leaders of thought in
supercomputing pointed out that
the human brain is not wired
to process many things **at once**.

For that reason, they reasoned that 2685
the supercomputer should not be wired
to process many things **at once**.
In the 1970s and ‘80s,
research computational scientists
only focused on ways
to use sequential and vector processing
supercomputers
and on how to use them
to solve the most computation-intensive
problems.
In the 1970s,
I believed that the modern
supercomputer
should be an ensemble
of the **slowest** processors
that communicates and computes
together
and as one seamless, cohesive
supercomputer.
In the 1980s,

Seymour Cray, 2686
who was the then leading luminary
in the world of supercomputers
argued that the supercomputer
must be powered by
a single vector processing unit
that is extremely fast.
Seymour Cray
believed in the old proverb that
[quote]
"He who tries to do many things
at once,
will accomplish nothing."
[unquote]
In 1989, it made the news headlines
that an African supercomputer wizard
in the United States
had discovered that
Seymour Cray was wrong.
That African discovered that
the modern supercomputer

can process many things
(or processes) **at once**.
I—**Philip Emeagwali**—was that
African supercomputer scientist.
My contributions
to the development of the supercomputer
was recorded
in the June 20, 1990 issue
of the *Wall Street Journal*.
Yet, I was not an overnight success
in supercomputing.
I began programming supercomputers,
exactly sixteen years earlier,
and at age 19.
I began programming supercomputers
on **June 20**, 1974.
I began programming supercomputers
at 1800 SW Campus Way,
Corvallis, Oregon, United States.

54.5 How I Turned Science Fiction to Technology

54.5.1 Fiction Became Non-Fiction

To invent
is to turn fiction to non-fiction.
The reason my experimental discovery
that parallel processing makes computers
faster
and makes supercomputers
fastest
made the news headlines
in 1989
was that it was the biggest **paradigm shift**
in the history of the computer.
At first, my experimental discovery
of massively parallel processing
was dismissed
as a scientific non-accomplishment
and dismissed
as a technological non-achievement.

At first, my experimental discovery 2689
of massively parallel processing
was, in fact,
cited to discredit my intellect.
Instead of discrediting my intellect
that invention
of how to solve the toughest problems
arising in both calculus and algebra
put me on the **perennial list**
of the ten most intelligent persons
in the world.
Back in the 1980s, when I announced
my experimental discovery
of parallel processing,
everybody said I made a mistake
and my discovery was promptly rejected.
My discovery of parallel processing
was ruled as a scientific offside,
akin to a soccer offside,
or the supercomputer programming mistake
of performing something illegal
in mathematical parallel processing
computations.

Before my experimental discovery
of parallel processing
in 1989
none of the then fastest supercomputers
were powered by parallel processing.
After my experimental discovery
of 1989,
all supercomputers and most computers
were powered by
parallel processing technology.
In the 1960s, ‘70s, and ‘80s,
the list of the computer pioneers
that **ridiculed** parallel processing
looked like it was drawn
as a roster of the Who’s Who
in the world of
high-performance computer technology.
Gene Amdahl of Amdahl’s Law fame
of April 1967
used Amdahl’s Law
to **ridicule** parallel processing.
Gene Amdahl is best known for
arguing that parallel processing

could not be harnessed
to solve the toughest problems
in extreme-scale computational physics.
Seymour Cray of vector processing fame
of the 1970s,
ridiculed parallel processing.
Seymour Cray caricatured
parallel processing
with his famous chicken versus oxen quote.
That quote of Seymour Cray
is the most often repeated quote
in supercomputing.
That quote of Seymour Cray
is always cited
by those that opposed parallel processing.
And Steve Jobs of Apple Computers
of the 1970s and '80s
ridiculed parallel processing
and **rejected** the technology
as a huge waste of everybody's time.
In the 1980s, twenty-five thousand [25,000]
computational scientists were only
using vector processing supercomputers

to solve their problems. 2692
Those twenty-five thousand [25,000]
computational scientists
and their leaders, **Seymour Cray**
and **Gene Amdahl**,
constantly **ridiculed**
parallel processing
and dismissed the technology
as a huge waste of everybody's time.
I—**Philip Emeagwali**—was the lone wolf
fulltime programmer of the most massively
parallel processing supercomputer
ever built.
By 1989,
the **digital divide** in supercomputing
can be summed as follows:
The lone wolf fulltime programmer
of the **precursor**
to the modern supercomputer
was **black**
and was born in Akure, colonial Nigeria
in the **heart of colonial Africa**.
I was the invisible black man

in the 1970s and '80s
world of supercomputing
who became very visible
after 1989.
I was the black person
that was as invisible
as a black goat at night.
I was also the **black** supercomputer scientist
that was unafraid of the darkness.
I was the black lone wolf
in the **uncharted territory**
of massively parallel supercomputing.
I was the lone wolf
on his quest
to cross the frontier of knowledge
and to discover
how and why parallel processing
makes supercomputers
fastest.
The reason American students
write school reports on "**Philip Emeagwali**"
is that I experimentally discovered
a new way of looking at computers

and understanding 2694
how and why parallel processing
makes computers
faster.
I experimentally discovered
a new **paradigm**
that changes how we compute
and how we do so at the fastest speeds.
Discovering how and why parallel processing
makes the computer
faster
will always prompt the leaders
from the world of computers
and supercomputers
to ask for your telephone number.
In the 1980s,
every vector processing
supercomputer scientist
that I spoke to
insisted that I had made a mistake
in my programming
of the world's fastest supercomputer.
I programmed that supercomputer

not as a computer *per se*. 2695
I programmed that supercomputer
as my new internet *de facto*,
or as my new global network of
65,536 tightly-coupled processors
that shared nothing
and that were married together
by another new global network of
1,048,576 email wires.
For that alleged mistake,
no vector processing
supercomputer scientist
took the time to read
my entire 1,057-paged
supercomputing research report.
That negative attitude changed
after the June 20, 1990 issue
of the *Wall Street Journal*
mentioned the names
"**Philip Emeagwali**" and "**Steve Jobs**"
on subsequent paragraphs.
So, I was not surprised
that **Steve Jobs** somehow got my

telephone number. 2696
On June 20, 1990,
both **Steve Jobs**
and the *Wall Street Journal*
did not take it for granted
that I invented
how to harness
the total computing power
of a new internet
that I described
as a new global network of
65,536
identical processors
that were cooperating,
that were tightly-coupled,
and that were equal distances **apart**.
That invention
opened the door
to the technology of parallel processing
that is at the core of the
modern computer
and today's supercomputer.
That invention

opened the door 2697
to parallel processing
one **million**
or even a **billion** problems (or processes)
at once.
Looking back, the June 14, 1976 issue
of the *Computer World*
carried an article titled:
[quote]
"Research in Parallel Processing
Questioned
as 'Waste of Time."
[unquote]
Today, we take it for granted
that the world's fastest supercomputer
manufactured in China
is fastest
because it harnessed
the total computing power
of ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
processors.

But to Steve Jobs 2698
my invention of 1989
of how to compress
65,536 days, or 180 years,
of **time-to-solution**
and how to compress that **time-to-solution**
across a new internet
that I described as my new global network of
65,536 processors
and how to compress that **time-to-solution**
to just one day
of **time-to-solution**
was like watching science fiction
become non-fiction.
In the 1980s,
I dreamt that in a century,
a school report said something to this effect:
"Once upon a time,
25,000 supercomputer scientists
embarked on a quest
for a **magic** supercomputer
that could be faster than
their fastest supercomputers

powered by only one processor." 2699
"Every two years,"
the school report continues
"their supercomputer only doubled in speed.
Then one day, a new magician arrived
from a sixteen dimensional universe.
That new magician
invented
the two-raised-to-power sixteen computers
as just one internet
that is *de facto*
one supercomputer, or processor,
that uniformly encircled a globe
in the magician's sixteen-dimensional
ancestral universe.
That new magician
invented
how to increase the speed of supercomputers
and increase them, in theory,
by a factor of two-raised-to-power sixteen
and increase them
by a factor of two-raised-to-power
sixty-four.

That new magician
opened the door
to the modern supercomputer
that can compress 180 years
of time-to-solution
to just one day
of time-to-solution.
That magician
opened the door
to the precursor
to their planetary supercomputer
that will enshroud the Earth
with its electronic cloth and consciousness.
That magician
turned fiction to fact."

54.6 Genius Responds to Polymaths, Not Mathematicians

54.6.1 The Polymath vs. The Mathematician

Solving the grand challenge problem 2701
of supercomputing
required a supercomputer
that is faster than the fastest computer
as well as scientific knowledge
that is beyond algebra, calculus, and physics.
Each of the twenty
grand challenge problems of supercomputing
was an interdisciplinary problem.
For that reason, it was impossible
for a scientist trained
in only one field
to even attempt to solve
the same grand challenge problem alone.
The research physicist
is on a quest for new knowledge
about how the universe works.
The research physicist
seeks to discover new laws of physics.
On the other hand,
the research mathematician
is not interested in physically discovering
the laws of physics *per se*.

The research mathematician
is interested in mathematically discovering
some laws of logic and numbers.
The research mathematician
is interested
on how to invent new equations.
As a research mathematician
my contributions
to mathematical knowledge
include my nine
partial differential equations
that were the cover story
of the May 1990 issue
of the *SIAM News*.
The *SIAM News*
is the bi-monthly news journal
that was the flagship publication
of the mathematics community.
The *SIAM News*
was published by the
Society for Industrial
and Applied Mathematics.
The research supercomputer scientist

is interested in **new physics**,
new mathematics and, most importantly,
on how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**.
The research supercomputer scientist
is on a quest to discover
new knowledge
that could push the frontier
of computational mathematics,
such as to solve a system of
coupled, non-linear, time-dependent, and
state-of-the-art
partial differential equations
that defines
an **initial-boundary value** problem
of calculus and
the most large-scale
computational fluid dynamics.
Such systems of equations
are called Grand Challenge equations
because the set of floating-point

arithmetical operations
for solving their algebraic analogues
are too computation-intensive
to be solved across
one processor,
or solves within only one computer.

54.6.2 Parallel Processing Hurricane Katrina

The research supercomputer scientist
is on a quest to experimentally discover
faster speeds in computation.
The **fastest** supercomputers
were used to push the frontier
of extreme-scale computational physics
and to make the impossible-to-compute
possible-to-compute.
The **fastest** supercomputers
were used to accurately forecast
Hurricane Katrina
and to forecast it
with excruciatingly-detailed

and coupled ocean-atmosphere 2705
general circulation models.
Back in 2005,
Hurricane Katrina of Southern United
States
forced insurance companies
to pay more than 40 billion dollars
in damages.
Hurricane Katrina
killed more than
one thousand eight hundred [1,800] people.
The fastest supercomputer
is needed to execute
the excruciatingly-detailed model
that was used to more accurately forecast
Hurricane Katrina.
The fastest supercomputer is needed
to reduce **forecasting uncertainties.**
The fastest supercomputer is needed
to save more lives and properties
and to make the world a better place,
and a more knowledgeable one.

54.6.3 How I Discovered Parallel Processing

Because massively parallel
supercomputing
transcended physics,
it could not have been discovered
from within physics
or by physical intuition alone.
Massively parallel supercomputing
was not discovered
from intuition and mathematical brawn,
alone.
Massively parallel supercomputing
had to be discovered
from the intuition,
brawn, and horsepower
that I generated **across**
my new internet
that I visualized
as my new global network of
65,536 tightly-coupled processors

2707

that shared nothing,
or a new global network
of as many computers.
I visualized
my sixteen-network-deep
networked motherboards
and the codes each executed
as **extensions**
of the singular motherboard that
extends the blackboard
and the differential and algebraic
equations
on it.
The new global network of
65,536 processors
is not, by itself, a contribution
to the development of the
supercomputer.
A processor
is like a **coffin**

that is merely a box
until you put somebody inside it.
However, the 65,536-fold increase
in supercomputer speeds
that compressed 65,536 days, 180 years,
of **time-to-solution**
to just one day of **time-to-solution**
and that **opened the door**
to the compression of
30,000 years of **time-to-solution**
to just one day of **time-to-solution**
is a **paradigm shift** of **tectonic scale**.
In theory, I invented
how to compress
a **time-to-solution**
that is as old as
42 times the age of the universe
and compress it
to just one second of **time-to-solution**
across a global network of
two-raised-to-power sixty-four

processors.
I visualized those processors
in sixty-four dimensional hyperspace.
My invention
of how to compress
thirty thousand [30,000] years
of **time-to-solution**
to just one day
of **time-to-solution**
changed the way we think about
the supercomputer of today
that will become the computer
of tomorrow.
That invention
of how, in theory, to compress
a **time-to-solution**
of 42 times the age of the universe
to a **time-to-solution**
of just one second
is the most important invention
in the history of the computer.

A research internet scientist
in quest
for the fastest supercomputer,
that parallel processed **across**
a new internet
had to be a polymath,
as well as a mathematician
and had to have the knowledge
and the command of materials
and the fluency of ideas
that are required to deliver lectures
on his contributions to scientific
knowledge.
The polymath's lectures
must be **fantastical**
and even make the mathematician
deeply uncomfortable.
The polymath must simultaneously
deliver his or her series of
multidisciplinary lectures
and speak from the **frontier**

of knowledge of large-scale
algebraic computations,
and speak from the **frontier** of
knowledge of advanced calculus,
and speak from the **frontier** of
knowledge of large-scale
computational physics,
and speak from the **frontier** of
knowledge of massively
parallel processing
and, in particular, processing
the grand challenge problem
of supercomputing
and solving that problem
across a new internet.
My new internet was a small copy
of the Internet
that is a global network of computers.
Like the Internet
that is a global network of computers,
or *de facto*

a global network of processors,
my new internet
is a new global network of
sixteen times
two-raised-to-power sixteen
regular, short, and equidistant
email wires
that married
two-raised-to-power sixteen
already-available,
tightly-coupled processors
together
and married them as one cohesive
whole unit
that is a not a computer *per se*
but is a supercomputer *de facto*.
It was a **paradox**
to believe that I solved
the toughest problem
in computation
and that I did so

at the crossroads of physics, calculus,
and computing
and believed that
I solved that grand challenge problem
of computing **alone**
and that I did so, without foremost,
arriving at the frontiers of knowledge
of physics, calculus, and computing.
It was a **paradox**
that was fueled by jealous academics.
It's only a biased person
that will watch
my videotaped lecture series
and continue to deny that
I was at the frontier of knowledge
of extreme-scale computational physics.
I put in 200,000 hours
in my quest
for how and why parallel processing
makes computers **faster**
and makes supercomputers

fastest,
namely, **the Philip Emeagwali formula**
that then United States President
Bill Clinton described
in his White House speech of
August 26, 2000.
It's a **paradox**
in their misunderstanding
that I solved an interdisciplinary
grand challenge problem
that spanned the disciplines of
mathematics, physics, and computing
and solved that interdisciplinary problem
without being at the frontiers of
knowledge of those three disciplines.
Put differently, it's impossible
to solve the toughest problem in calculus
and solve it
without, foremost,
learning how to solve
the easiest problems in calculus.

54.7 My Quest for the Philip Emeagwali Supercomputer

54.7.1 Quest for the Fastest Supercomputer

Faster supercomputers
yield more scientific discoveries
that push the frontier of
computer knowledge
as well as broaden our horizons.
The fastest supercomputer is used
to execute general circulation models
that is used to foresee
the worst-case future
of the planet we call home.
The 21st century supercomputer scientist
that embarked on his **hero's quest**
for the fastest parallel processing
supercomputer

should be armed with two swords.
That hero's quest
is for the frontier of knowledge
and for the unchartered territory
where ground-breaking scientific
discoveries can be discovered
and where technological inventions
can be invented.
That hero's quest
is for the primal place
where galaxies are astronomically large.
The frontier of astronomy
is where strange stars can be discovered.
The frontier of astronomy
is at astronomically far-away distances.
At the frontier of astronomy,
astronomically expensive,
multi-billion dollar telescopes
are required to discover
those astronomically far-away stars.
One such telescope

—called the **Square Kilometer Array**,
or SKA for short—
is a network of three thousand
15-meter dishes,
with a total collecting area
of one square kilometer.
The **Square Kilometer Array**
has a price tag of **1.5 billion** dollars.
The world's most expensive telescope
works together
with the equally expensive array
of **processor**s,
such as the new global network of
64 binary thousand
tightly-coupled **processor**s
that is married together
as one seamless, cohesive supercomputer
and married
by one binary million
regular, short, and equidistant
email wires

that I experimentally discovered 2718
how to program
to send and receive
astronomically large data
to and from
64 binary thousand
cooperating **processors**.
That hero’s quest
is for the primal time,
or the be**gi**nning of time,
when the big bang explosion occurred.
That hero’s quest
is for the primal things,
or the things that are extremely small.
That hero’s quest
is for the primal computations,
or the quantum computations,
that can occur at sub-atomic levels.
That hero’s quest
is for the primal communications
and computations

that made the news headlines 2719
when I experimentally discovered them
as the fastest computations
across one binary million
regular, short, and equidistant
email wires
that wired 64 binary thousand
processors
and wired them
as a global network
that I named a HyperBall
Supercomputer.
In our search, or rather re-search,
for distant galaxies colonized by
super-intelligent beings
—or for planets that exhibits evidence
of **extra**terrestrial intelligence—
that primal place
is at the edge of our visible, known
universe.
The primal place

where strange galaxies can be discovered 2720
is beyond the edge
of our visible universe,
which is 18 billion light years.
The distance of that primal place
is 18 billion times the miles
light travels in a year
while travelling at its speed in a vacuum.
That speed is always about
186,000 miles per second,
or exactly 299,**792**,458
meters per second.
That hero’s quest
is a response to a call for adventure
arising from humanity’s need
to **foresee** otherwise **unforeseeable**
climate change.
The **unforeseeable** global warming
is sometimes **seen**
by reducing 30,000 years
of **time-to-solution**

to only one day 2721
of **time-to-solution**,
or to executing
the excruciatingly-detailed
general circulation model
and solving
that initial-boundary value problem
of modern calculus
and computational physics
and solving the problem
as soon as possible.
The **first sword** on that hero’s quest
must be intellectual
and must be used
to understand parallel processing.
And, most importantly,
to go against convention
and to go against
the prevailing paradigm
to advocate parallel processing
and advocate the technology at a time

parallel processing was misunderstood 2722
and dismissed by **everybody**
as a huge waste of **everybody**'s time.
The **second sword** on that hero's quest
must be a **never-before-seen**
electronic instrument
that will be used to see an
ensemble of processors
that was previously unseen
as one unit
and unseen as a seamless, cohesive
supercomputer.
That **second sword** on that hero's quest
is used to compress 10.65 million days,
or 30,000 years
of **time-to-solution**
to only one day of **time-to-solution.**
That compression of **time-to-solution**
was discovered
across an ensemble
of 10.65 million processors.

The strong **intellectual sword** 2723
on that hero’s quest
is used to **slash**
the grand challenge problems
of supercomputing.
That **intellectual sword**
is used to understand
the laws of physics
and is used to invent
the partial differential equations
of calculus
that encode those laws of physics.
That **intellectual sword**
is used to formulate
the system of equations
arising in large-scale
algebraic computations
that approximates
those partial differential equations
that define
the initial-boundary value problem

that is at the mathematical core 2724
of the most vexing
grand challenge problem.
That **intellectual sword**
takes the inventor
of the modern supercomputer
to the frontier of knowledge
and into the **uncharted territory**
known as massively parallel processing.
The **physical sword,**
or the never-before-seen,
large-scale ensemble
of **processor**s,
takes the discoverer
inside the *terra incognita,*
or the unknown world,
of scientific knowledge
where floating-point arith**me**tical
operations
could be executed
at speeds previously **unrecorded**.

That unknown world 2725
is the **primal place**
where the previously unseen is seen.
That **primal place**
is where ground-breaking discoveries
are made.
That **primal place**
is a place of extremes.
The **physical instrument**
used at that **primal place**
is often as extremely expensive
as the **13.25**-billion-dollar
Large Hadron Collider
in Geneva, Switzerland.
The **physical instrument**
used at that **primal place**
is often as astronomically large
as the most powerful telescopes.
That **physical instrument**
is used to search for
infinitely small subatomic particles

—such as the **boson.**
That **physical instrument**
is a supercomputer-in-progress
that could be used as the prototype
for my proposed Cosmic Supercomputer
on the **North Pole**.
But, more importantly,
the **human instruments**
must be a large team of
renaissance scientists
and polymaths,
each a walking encyclopedia.
My experimental discovery
of parallel processing
is permanently embodied
inside your computer.
Parallel processing
enables us to obtain a **surer** and **deeper**
understanding
of our universe, and in particular,
enables us to **foresee** otherwise

2727

unforeseeable climatic changes
that enshroud the Earth.
Faster supercomputers
enable us to climb higher
and up the ladder of knowledge
and to make the **impossible**-to-compute
possible-to-compute.
Faster supercomputers
make it possible to **foresee**
otherwise **unforeseeable**
climate change.
Faster supercomputers
make it possible
to develop more powerful technologies.
Over the last five millennia,
we progressed from the wheel
to the automobile
and progressed from the planting of
cassava
to hoping for the discovery
of the cure of AIDS.

2728

Hopefully, our descendants will progress
from the abacus
as the computing aid
of three thousand [3,000] years ago
to our hoped-for super-intelligent
supercomputers
of three thousand years from today.
I foresee our descendants
of a thousand millennia
to be super-intelligent lizards
that are masquerading as post-humans
in the planet Mars.

54.7.2 The Supercomputer as Walking Stick

Year Million
is a thousand millennia away.
I was asked:
“How will our **Year Million**
post-human descendants

be supercomputing?" 2729
My answer is this:
"I don't know the answer?"
But I believe that the supercomputer
will be the walking stick
in humanity's million-year hero's journey.
That spiritual journey
to envision our post-human descendants
will be akin to metaphorically visiting
the Land of the Cyborgs,
where each cyborg
is half-human
and half super-intelligent computer.
That spiritual journey
to envision cyborgs
will be akin to metaphorically visiting
the Land of the Spirits
of my grandma's folklore
of my ancestral Igbo Land
of southeastern Nigeria.
I believe that by the end of our millennium,
our descendants
could reinvent themselves

as **asexual cyborgs** 2730
—that is, half humans, half computers—
each cyborg with a great sense of humor.
I believe that our cyborg descendants
of Year Million
will be **anthro**po**morphic**,
or have human attributes.
I believe that our cyborg descendants
of Year Million
will be human-like
because we humans will create them
in our own human image.
I believe that our cyborg descendants
will not have computers around them.
I believe that
the computer will be within them.
I believe that
our descendants will not need computers
because they will **BE** computers.