46 Why I Invented a New Supercomputer—Part 4 of 15

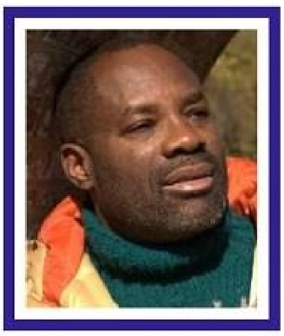

info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture 170930
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

46.1 My Quest for the Fastest Supercomputer

46.1.1 They Called Me "Calculus"

The fastest supercomputer
costs the budget of a small nation
and it is purchased because
the fastest supercomputer
gives meaning to life,
and because
the fastest supercomputer
makes the world a better place,
and because
the fastest supercomputer
makes humanity more knowledgeable
and because
the fastest supercomputer of today
will become the computer of tomorrow.
A faster supercomputer proves that
humanity is progressing
in the right direction.

A faster supercomputer
increases our level of civilization
and enables our children
to do better than us.
The supercomputer
that computes in parallel
is the most important instrument
in **ex**cru**cia**tin**gly**-detailed
computational physics
and in extreme-scale
computational mathematics.
In the history of mathematics,
the answers to the biggest questions
don’t come easy.
The solutions to the biggest questions
demand big ideas and hard work.
To the leaders of thought
in the world of mathematics,
the biggest question in calculus
is this:

How can the twenty-first century
computational mathematician
harness a new internet
that is a global network of millions,
or even billions,
of processors
or harness a global network of
as many tiny computers
and harness that new internet
to solve a system of coupled, non-linear,
time-dependent, and state-of-the-art
partial differential equations
of calculus?
My 1989 solution
of the toughest problem in calculus
was the front page story
of the top mathematics news journal,
namely,
the May 1990 issue
of the *SIAM News*
that was published by the Society

for Industrial and Applied Mathematics.
To the world leader,
debating how to mitigate
the consequences of climate change,
the toughest problem in mathematics
is this:
How can the supercomputer scientist
solve an initial-boundary value problem,
such as executing
a general circulation model
and executing it to foresee
otherwise unforeseeable
global warming?
Solving the toughest problem in calculus
demands the most powerful
processor, or supercomputer,
ever built or imagined.
Yet, the toughest problem in calculus
cannot be solved
with only one **isolated**
processor

that was not a member
of an **ensemble** of processors,
or only one isolated computer
that was not a member
of an **ensemble** of computers.
The computer is not powerful enough
to be used
to solve the toughest problems
in calculus.
Only the fastest supercomputer
can be used
to solve the toughest problems
in calculus.

46.1.2 Conversations in Calculus

It took me twenty years
to learn and discover
how to solve the toughest problem
in calculus.

During those twenty years,
onward of June 1970,
I studied physics, mathematics,
and computing.
I studied to be a **triple threat**
and to possess
the interdisciplinary knowledge
and to be a **polymath**
that possessed knowledge
that is beyond mathematics.
Only a **polymath**
can deeply understand
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.

That invention 2196
of the massively parallel processing
supercomputer
that was ridiculed as a
huge waste of everybody's time
required that I be a polymath
that is simultaneously
at the frontiers of knowledge
in algebra, calculus, physics, computing,
and, in particular,
that I master how to communicate
my computer codes
and how to send and receive those codes
via emails to sixteen-bit long addresses
and how to email them **across**
a new internet
that is a new global network of
65,536
commodity-off-the-shelf processors
that are tightly-coupled,
that are identical

and that are equal distances **apart**
with each processor
operating its own operating system
and with each processor
having its own dedicated memory
that shared nothing with each other.

46.1.3 Crossroad of Calculus and Computing

That parallel processing problem
of supercomputing
was named by the United States government
as the Grand Challenge Problem
of computing **across**
millions upon millions of processors.
The problem was a grand challenge
for a good reason.
It’s true that
parallel programming the pre-cursor
to the modern supercomputer

—and parallel programming
that massively parallel supercomputer
to compute in parallel,
or to process many things (or processes)
at once
is no walk in the park.
The technology of parallel processing
was so terrifying
to the vector processing
supercomputer scientist
that the June 14, 1976 issue
of the *Computer World*
carried an article titled:

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

I had the **confidence**
to make the impossible-to-compute

2199

possible-to-compute
and make it possible
at the crossroad of
physics, calculus, and computing.
I had that **confidence**
because I was trained for **twenty years**
and trained in physics, calculus, algebra,
and computing.
I had that **confidence**
because I had received positive feedbacks
from computational mathematicians.
I had that **confidence**
because I had given out-of-town
research seminars in supercomputing
in 1982, ‘83, and ’84
and seen the reactions
of the scientific community.
I had that **confidence**
because I had seen the looks on their faces
when I explained to
supercomputer scientists

how far I had progressed
in my then ongoing research
in massively parallel processing.
Another reason for my **confidence**
was that, in the 1980s,
I attended 500 weekly research seminars.
Each seminar speaker was invited
because he discovered a method
in computational physics
or contributed an equation
to calculus
or invented a technique
for solving a large-scale
system of equations
of algebra
and solving them
on a supercomputer.
I had that **confidence**
because I knew, firsthand,
that I—**Philip Emeagwali**—
as was more knowledgeable

than those 500 research scientists. 2201
I had that **confidence**
because I became a polymath
after **twenty years** of training
and competing against 25,000
computational mathematicians
that each had only ten years of training
and that each only had
programming experience
with vector processing supercomputers.
Metaphorically speaking,
I had the confidence
of a twenty-year-old male
professional wrestler
that was challenged
to a wrestling match
against a ten-year-old female
professional wrestler.
I was the twenty-year-old boy
amongst ten year old girls.
For those 25,000

computational mathematicians 2202
parallel processing was a **pipe dream**.
For I, the lone wolf
at the **farthest frontier**
of massively parallel processing,
it wasn't a **pipe dream**.
During my scientific quest
to experimentally discover
how and why parallel processing
makes the impossible-to-compute
possible-to-compute,
I was alone
and at the crossroad of
calculus and computing
and at the crossroad of
physics and supercomputing
and at the crossroad of
algebra and parallel computing.

46.2 How I Discovered the Fastest Supercomputer

46.2.1 Solving the Toughest Problem in Calculus

After the first decade
and from the frontiers of knowledge
of the **partial differential equations**
of calculus,
I saw the horizon
of the **uncharted territory**
of massively parallel supercomputing.
That technological vision
gave me the motivation
to spent my second decade
in narrowing my focus
to how to invent
the massively parallel processing
supercomputer.
In 1989, it made the news headlines

that I had invented 2204
how to solve the toughest problems
in computational physics
and invented
how to compute their answers **across**
a new internet
that I visualized
as a new global network of
64 binary thousand
tightly-coupled processors
that are equal distances **apart**.
I visualized
those two-raised-to-power sixteen
tightly-coupled processors
to be on the fifteen-dimensional surface
of a globe, or a **hypersphere**,
in a sixteen-dimensional universe.
I visualized
sixteen times
two-raised-to-power sixteen,
or one million

forty-eight thousand
five hundred and seventy-six [1,048,576]
bi-directional email wires
that are **uniform** and **regular.**
I visualized
those two-raised-to-power twenty
email wires
that were on the surface of a globe,
or a **hypersphere**,
in sixteen dimensional **hyperspace**.
I visualized
those one binary million wires
as the means to circulating emails
that nourished my 64 binary thousand
processors.
I visualized
those two-raised-to-power sixteen
tightly-coupled processors
as nourished with arithmetical data
that I needed to solve my large-scale
initial-boundary value problems

of computational physics and calculus. 2206
I visualized my boundary data
and the intermediate answers
to my initial-boundary value problem
as flowing via emails
to and from sixteen-bit long email addresses
and to and from each and every
of my two-raised-to-power sixteen
tightly-coupled processors.
I visualized my arithmetical data
as transmitted via emails
to and from along sixteen directions
that were metaphorically speaking
mutually orthogonal,
that is, sixteen directions
that are mutually perpendicular
in an imaginary sixteen dimensional
universe.
I visualized my emails around
my calculations
on the surface of the **hypersphere**,

in **hyperspace**.
I invented
how to unravel my email pathways
that **wind to** and **from**
metaphorically equidistant
processors
that I visualized
as on the surface of a **hypersphere**
in **hyperspace**.
I visualized how to map
those 64 binary thousand points
onto the surface of a globe
in our everyday space,
such as planet Earth.
I visualized my global network of
processors
as having applications
towards global circulation models.
I invented
the massively parallel processing
supercomputer

2207

and I invent the technology 2208
by narrowing my focus
to the **hypersurface**
of that **hyperball**
in **hyperspace**
and by throwing away
the less promising, nitty-gritty
engineering subject matters
and throwing unnecessary details
over the side.
For the decade of the 1980s,
I sat alone
starring at an abandoned
massively parallel processing supercomputer
powered by
up to 65,536 tightly-coupled,
commodity-off-the-shelf processors
and I sat alone because
everybody else **ridiculed**
that **never-before-seen** supercomputer

as a huge waste of everybody's time.

46.3 A Mathematician's Quest for the Fastest Supercomputer

46.3.1 Diary of a Lone Wolf Programmer

There was no instruction manual
on how to harness the power
of the then **never-before-seen**
massively parallel supercomputer
that was abandoned
for me to parallel program alone.
Nor was there a help desk
that could explain
how I could synchronously send
and receive
64 binary thousand emails.
Back on June 20, 1974,
when I began to program
sequential processing supercomputers

there were an average of
24 programmers
logged into the sequential processing
supercomputer
that I was programming.
Those were the good old days
when it was possible
to program
the world's fastest supercomputer
and do so alone.
On June 20, 1974,
in Corvallis, Oregon, United States,
I even had a supercomputer instructor,
had a supercomputer instruction manual
and a supercomputer help desk.
But a decade and half,
onward of June 20, 1974,
in Los Alamos, New Mexico, United States,
I was the lone wolf programmer
that was at the **farthest frontier**
that was the most massively parallel

supercomputer
ever built.
In the 1980s,
I had no massively parallel
supercomputer instructor.
The reason was that
parallel processing
was then considered
a waste of time.
In the 1980s,
I had no massively parallel
supercomputer instruction manual.
The reason was that
parallel processing
was at that time
not yet experimentally confirmed.
In the 1980s,
I had no massively parallel supercomputer
help desk.
The reason was that
in that decade

no supercomputer scientist
could answer questions
on how to parallel process and harness
the total computing power
of a 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.
That massively parallel processing
supercomputer
was the **pre-cursor**
to the modern supercomputer.
In the 1970s and ‘80s,
I was **ridiculed** and **laughed** at
for attempting
what seemed **impossible**-to-compute
namely, to program
massively parallel processing

supercomputers
and parallel program them
to solve the toughest problems
arising in computational physics
and parallel program such supercomputers—
powered by up to 65,536 processors—**alone**.
For that experimental supercomputer
of the 1980s,
only one person can control its
65,536
equidistant processors.
Only one person can lock
all its processors
and lock them simultaneously.
I—**Philip Emeagwali**—
was the lone wolf
that was at the **farthest frontier**
of massively parallel processing
and the supercomputer scientist
that controlled all those

65,536 processors
and controlled them at all times.

46.3.2 Father of the Modern Supercomputer

I experimentally programmed
a room-sized model
or a working prototype of the modern supercomputer
that could be constructively reduced
to the planetary-sized Internet.
I said that
I—**Philip Emeagwali**—working alone
invented **a supercomputer**
that encircled a room-sized globe
that I visualized
as an idealized model
of the Internet
that encircled the Earth.
I also said that

although the supercomputer 2215
had 25,000 vector processing
supercomputer programmers,
and had many fathers and mothers,
and uncles and aunts,
that I am the only **father**
of the supercomputer
that invented a new, parallel processing
supercomputer
and invented it as the lone wolf
at the **farthest frontier**
of the most massively parallel processing
supercomputer.
Because the modern supercomputer
is now automated
making it easier to parallel program
it seems every other person
could be a supercomputer programmer.
Unlike in the 1980s,
there is no lone wolf
massively parallel processing

supercomputer programmer.
Consider the quantum computer.
In May 2016,
IBM Corporation
made its **five-qubit** quantum computer
available to the public.
That **five-qubit** quantum computer
attracted 40,000 users
and attracted 275,000
quantum computer experiments.
Thirty years earlier, or in the 1980s,
I—**Philip Emeagwali**—was the lone wolf
programmer in Los Alamos, New Mexico,
United States,
and the only person
that was attracted
to the most massively parallel processing
supercomputer
ever built.
In the 1980s,
none of the twenty-five thousand [25,000]

programmers 2217
of vector processing supercomputers
and their leaders, Seymour Cray
and Gene Amdahl,
could harness
even eight
of the processors
of a massively parallel supercomputer.
Nobody, except I
that devoted sixteen years
to the parallel processing craft,
would touch the technology
and do so with a ten-foot pole.
You may ask:
"Why do we have
40,000 early adopters
for the first quantum computer
but only had one early adopter
for the first massively parallel
supercomputer?"
The reason is that, in the 1980s,

it was 40,000 times more difficult
to parallel program across a new internet
that is a global network of
65,536
processors
than it is to program
one quantum computer
of today.

Recently, a 12-year-old
that was writing a school report
on the invention of the
parallel processing supercomputer
asked:
"Why is Philip Emeagwali called
the **father** of the modern supercomputer
that computes in parallel?"
I answered:
"I'm called the **father**
of the parallel processing supercomputer
because, for sixteen years,
I was the lone wolf,

in a pack of 25,000
supercomputer scientists
in the world
that were led by Seymour Cray,
that stood alone
and at the **farthest frontier**
of supercomputing,
called massively parallel processing.
I was the first supercomputer scientist
to be widely recognized
for his **experimental discovery**
of how to compute many things **at once**
instead of computing
only one thing **at a time**."

46.4 Father of Parallel Processing Supercomputers

46.4.1 Discovering Parallel Processing

The reason the seventh-grader
is doing a school report
on **Philip Emeagwali**—as the Father of Parallel Processing Supercomputers—
is that I experimentally discovered
how and why massively parallel processing
makes the supercomputer
faster.
Today, massively parallel processing
is the technology
that powers the modern computer
as well as the fastest supercomputer.
I experimentally discovered
how and why
massively parallel processing across
64 binary thousand
processors

of a supercomputer
can be faster than
sequential or vector processing
on one supercomputer
computing with only one isolated
processor
that was not a member
of an ensemble of processors.
It is faster by a factor of
64 binary thousand.
My invention
of the massively parallel processing
supercomputer
was extended by Chinese supercomputer
designers
and extended from my factor of
64 binary thousand increase in speed
to their factor of ten million
six hundred and forty-nine thousand
six hundred [10,649,600] increase in speed.
In the 1980s,
my experimental discovery

2222

of the 65,536-fold speed increase
of the massively parallel processing
supercomputer
was science fiction
to authors of textbooks on supercomputers.
During the 1980s,
the twenty-five thousand
[25,000] supercomputer geeks
that used conventional vector
supercomputers
to solve their most computation-intensive
scientific problems
were like a roll call of the Who's Whos
across all disciplines
of science, technology, engineering,
and mathematics.
It's often said that supercomputing
is an interdisciplinary field.
The largest supercomputer manufacturer
—in the United States or Japan or China—
is *de facto* an interdisciplinary
research and development team.
The fastest supercomputer in the world

is custom manufactured
by an engineering team
of about one thousand engineers
and scientists.
The largest supercomputer
development project
can have a billion-dollar budget.

46.4.2 A Black Supercomputer Wizard in Oregon

I envisioned a future for humanity
that's powered by massively
parallel processing.
That technology is a hundred and eighty
[180] degrees apart
from sequential processing.
Briefly,
the heart of the computer
resides
on the processor
while the heart of the supercomputer

resides and spreads 2224
across millions upon millions
of processors.
So comparing the computer
to the supercomputer
is like comparing one wave
to the sea.
I discovered that
creating a new technology
creates a need for a new narrative
for the history of technology.
The date June 20, 1974
was the day I first logged into
a conventional sequential processing
supercomputer.
The first supercomputer I programmed
was marketed
as the worlds’ fastest computer
when it was manufactured
back in December 1965
and two years before the word
“supercomputer” was coined.
The reason it was called [quote unquote]

the “first supercomputer” 2225
was that it was the **first computer**
that could theoretically execute
the record speed of
one million instructions
per second.
The first supercomputer
that I programmed was inside
the Computer Center
at 1800 SW Campus Way,
Corvallis, Oregon, United States.
Only twenty-four programmers
were time-sharing that supercomputer.
All twenty-four supercomputer programmers
were logged-in
from twenty-four teletypewriters
that were scattered around
the state of Oregon.
In 1974,
the number of computer programmers
in the entire world
was counted in the thousands.
Today, the number of

information technologists in India, alone,
is counted in the millions.

46.4.3 Forging a Path to the Modern Supercomputer

As a researcher in quest
for the Holy Grail
of the fastest computation,
my central objective was to **forge a path**
to the modern supercomputer
that's powered by massively
parallel processing technology.
In the 1970s and 80s,
the difference between I
and the largest research groups
was that they had to write
theoretical papers
on a, sometimes, theoretical supercomputer
in other to get paid.
They don't get paid for attempting
to make the **impossible**-to-compute
possible-to-compute.

Because I was **black** and African, 2227
I wasn’t getting paid
which meant that
I was the only supercomputer scientist
with the freedom
to attempt to make
the **impossible**-to-compute
possible-to-compute
That unpaid, lone wolf
black supercomputer geek
had the freedom to devote a decade
to experimentally how to program
a massively parallel supercomputer
while the teams of six-figured salaried
white researchers
were writing hypothetical papers about how,
in theory,
a massively parallel supercomputer
would be a huge waste of their time.
For sixteen years,
I programmed sixteen
massively parallel supercomputers,
while theorists wrote papers

and textbooks
about a massively parallel supercomputer
that they never programmed.
Their theoretical papers
pertaining to the impossible-to-compute
gave rise to the saying that
parallel processing
is a beautiful theory
that lacked experimental confirmation.
To write a theoretical paper
on parallel processing
and write it
without conducting a companion experiment
on parallel processing
is akin to a fifth grader
writing a school report
on the contribution of **Philip Emeagwali**
to the supercomputer
and writing that school report
without reading a book
on the life of **Philip Emeagwali**
and writing that school report
without knowing the definition

of the word “supercomputer.”
It’s akin to a fifth grader
writing a book report
on a book he never read.
In the 1970s and ‘80s,
I was the lone programmer
going in the **opposite direction**
and paradigm
from theorists on
massively parallel supercomputing.
In my **opposite direction**,
I defocused on what the textbooks said:
namely, that parallel processing
will never work.
A supercomputer expert from M.I.T.
in Cambridge, Massachusetts
warned in the June 14, 1976 issue
of the **Computer World magazine**
that the massively parallel supercomputer
is **large** and **clumsy** and a waste of time.
As the lone wolf
at the **farthest frontier** of supercomputing,
it was intuitive that I focused

2230

on programming the unknown
massively parallel supercomputer,
instead of the known
vector supercomputer.
In the 1970s and 80s,
I didn't fit
and couldn't work as part of a team
in conventional vector processing
supercomputing
who believed that
parallel processing
is a huge waste of everybody's time.
I was **dismissed** from
vector processing supercomputer teams
and forced to become
a one-person parallel processing
supercomputer team.
I had to forge
a never-before-taken technological path
to enable me to reach
the *terra incognita*
that is the unknown world
of massively parallel supercomputing

and where I discovered
the fastest computations
across my new internet
that's a global network of
65,536 processors.
I had to forge a unique
supercomputer path
into the history of computing.

46.5 Origin Story of Philip Emeagwali Supercomputer

46.5.1 The Sword of Truth

Prior to my experimental discovery,
parallel processing was dismissed
as a beautiful theory
that lacked experimental confirmation.
I experimentally discovered
that a new internet
that's outlined and defined
by a global network of processors

that communicates and computes
in parallel
could become a new supercomputer.
In the 1970s and ‘80s,
I witnessed my global network
of the slowest 65,536
processors
change.
That new internet
changed like the caterpillar
that slowly transformed
into its final form.
That **final form**
was my metaphor for a small internet.
That new internet
that I visualized in the sixteenth dimension
was my **shapeshifting** magic sword.
The discoverer saw the unseen
and became the magician
that wields the sword of truth.
Like the caterpillar,
I changed during the sixteen years
onward of June 20, 1974.

At age nineteen [19],
I was programming a supercomputer
that computed only one thing
at a time.
For the sixteen years
onward of age nineteen,
I was always in the shadows,
and always on the horizons.

46.5.2 Birthdates of the Computer

To invent means to originate
or to create
a product of the inventor's ingenuity.
We cannot assign a precise birth date
and a place
to any complex invention
such as the modern computer
that continuously evolved
over the past seven decades.
We all agree that the computer

was not invented by a
super-intelligent lizard
from the Moon.
The speed of the modern computer
increased gradually
but made a **quantum increase**
of a factor of 64 binary thousand
in 1989.
That was the largest recorded speed up
since the programmable computer
was invented **in 1946**.
That first computer
was in Aberdeen Proving Ground
at the outskirt of Baltimore, Maryland
that is my wife's hometown
and my second hometown.

46.5.3 Birthdate of My Contribution to the Computer

My contributions to the development
of the new internet
that is a new global network of
65,536

tightly-coupled processors
made the news headlines
as the **experimental discovery**
of how and why parallel processing
makes modern computers **faster**,
as the **experimental discovery**
of how and why parallel processing
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.
It made news headlines as
the **invention**
of the world's fastest supercomputer.
That **invention**
of a new supercomputer
was recorded in the June 20, 1990 issue
of *The Wall Street Journal*.
So I ask:

What day was the internet
that is *de facto* a new supercomputer
invented?
To the historian of the internet
my date of conception was:
June 20, 1974,
or the date I began my quest
on a supercomputer
that was at 1800 SW Campus Way,
Corvallis, Oregon, United States.
Three weeks after I began
that supercomputing quest
and at age nineteen,
I was on the cover
of an Oregonian newspaper
that circulated only in the cities of
Monmouth and Independence, Oregon.
So, nineteen seventy-four [1974]
was the year **Philip Emeagwali**
began his quest for a new internet.
To the student writing a school report
on the contributions of **Philip Emeagwali**
to the development of the supercomputer

the answer is:
1989
or the year my invention
of the massively parallel processing
supercomputer
made the news headlines.
Nineteen eighty-nine [1989]
was the year **Philip Emeagwali**
and his new internet
that is a new supercomputer *de facto*
came of age,
or blossomed.

46.5.4 Birthdate of My Contribution to the Internet

There's no single flow chart
that describes how to invent
the first internet
and how to become
the father of the Internet.
Moreover, an invention
has three birthdates.
The first birthdate of an invention

is it’s date of
conception.
The second birthdate of an invention
is it’s date of
constructive reduction to practice.
The third birthdate of an invention
is it’s date of
physical reduction to practice.
In the summer of 1974,
I conceived a new internet.
My **new internet** was shifted off stage
but I constantly shifted it
towards the center stage.
Over the following decade and half,
I invented that **Philip Emeagwali** internet
as a new global network of computers
and it made the news headlines
when I invented
that new internet
as a new global network of
65,536
already-available, tightly-coupled
processors.

2239

In the late 1970s,
I **constructively reduced** to practice
a **new internet**
that is a new global network of computers
that arose from a kernel
that had a **one-to-one** correspondence
between the vertices
of the **truncated icosahedron**
and the processors
of that new internet.
It also had a one-to-one correspondence
between the edges
of the **truncated icosahedron**
and the regular, short, and equidistant
email wires
of that new internet.
I discovered that
more already-available
processors
and regular, short, and equidistant
email wires
can be added by **triangulation**
of the **truncated icosahedron**.

As the sole inventor 2240
of the **Philip Emeagwali** supercomputer,
I had to know, *a priori*, the topology,
or the locations,
of my ensemble of processors
and know where each and every processor,
or electronic brain, is located
and know those locations,
both forward and backward,
and know them with the completeness
an American airline pilot
had to know the geography of Nigeria
to understand where Abuja is located.
I visualized short email wires
for processor-to-processor communications
comprised of wires
printed onto circuit boards.
I visualized long email wires
comprised of fiber optic cables
and electrical cables.
I visualized my ensemble of processors
as communicating and computing
together

as one seamless, cohesive supercomputer.
I visualized using plentiful, powerful,
and inexpensive commodity,
or large numbers of available,
processors
that were designed for everyday computers
as opposed to a few
custom vector processors
that were designed for the
fastest supercomputers.

46.5.5 Closure—Inspirational

My primary quest
was to challenge the established truth.
That established truth—of the 1980s
and earlier—was that the slowest
processors
cannot work together
to solve

the most computation-intensive
problems
in algebra
or in large-scale computational physics.
As the sole inventor
of the **Philip Emeagwali** supercomputer,
I had to know, *a priori*, the topology,
or the locations,
of my ensemble of processors
and know where each
and every processor,
or electronic brain, is located
and know those locations,
both forward and backward,
and know them with the completeness
an airline pilot
had to know the geography of Nigeria
to fly from Lagos to Abuja.
I visualized short email wires
for processor-to-processor

email communications 2243
that **I visualized** as comprised of wires
printed onto circuit boards.
I visualized long email wires
comprised of fiber optic cables
or electric cables.
I visualized my ensemble of processors
as communicating and computing
together
and doing both
as one seamless, cohesive
parallel processing supercomputer
that is the world's fastest computer.
I visualized using commodity,
or large numbers of available,
processors
that were designed for everyday
computers
as opposed to using a few
custom vector processors

2244

that were designed for the
fastest supercomputers
of the 1970s and '80s.
My contribution
to the development of the supercomputer
is that I made the **impossible** possible
when I encoded the laws of physics
into the partial differential equations
of calculus
that I discretized
into systems of equations
of algebra
that I coded and solved
across
a primordial internet
that I invented
as a global network of
65,536
tightly-coupled processors.

In supercomputing,
to discover or invent
is to show that
the **impossible-to-solve** is,
in fact, **possible-to-solve**.
The genius is the ordinary person
that found the extraordinary
in the ordinary.
I conduct scientific research
to know what is **discoverable**
and **knowable**
but which nobody knows.
To witness a scientific discovery
that has rich, fertile, and far-reaching
consequences
is like walking into a forest
and witnessing a lot of leaves
fall on your head.

2246

47 Philip Emeagwali's Cosmic Supercomputer

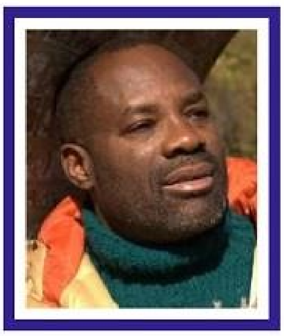

Philip Emeagwali

info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture 170930
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

47.1 My Quest for the Fastest Supercomputer