24 How I Invented a New Internet

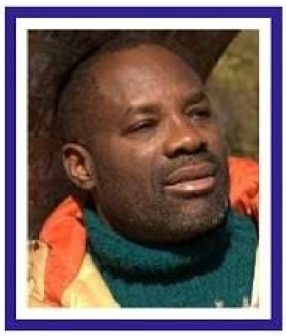

Philip Emeagwali
info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture 180126-2
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

24.1.1 Inventing a New Algebra

To invent
is to see something that was
previously unseen
and to see something
that did not previously exist.

I invented new computational mathematics
that are important
to my oil-producing country of birth,
Nigeria,
and that will enable
the petroleum industry
to use the massively
parallel processing supercomputer
to discover and recover
otherwise elusive
crude oil and natural gas.
I invented new computational mathematics
because I asked
important mathematical questions
that are central to
the toughest problems
in computational physics.
The research pure mathematician
asks questions
that are directly centered on
mathematical knowledge itself.

I am the research
massively parallel processing
computational mathematical physicist
who asked mathematical questions
that had corresponding physical answers.
As a pure mathematician
of the 1970s,
I knew there are no easy [quote unquote]
"new mathematics."
Also, I knew that
it is beyond the intellectual reach
of somebody trained only in
pure mathematics
to possess a deep understanding
of the eighty-one [81] partial derivatives
that represent
the four physical forces
that define
the Second Law of Motion
of physics
and to as well understand

that thirty-six [36] partial derivative terms
are missing
in the mathematical physics textbooks
on the multiphase porous media flows
and missing since Darcy's formula
was **empirically** discovered in 1856
and missing since Darcy's formula
was used *in liu* of
the nine **Emeagwali**'s
partial differential equations
of modern calculus.
That lack of deep knowledge of physics
was the reason, a pure mathematician
could not discover and correct
the mathematical errors
that I discovered
in the central partial differential equations
of the physics of the petroleum industry
and could not invent the new calculus,
the new algebra, and the new algorithms
that I invented

in the early 1980s
and that were the cover stories
of top mathematics publications,
such as the cover story
of the May 1990 issue
of the *SIAM News*.
The *SIAM News*
is the top publication in mathematics.
The *SIAM News*
is written by research mathematicians
and written for research mathematicians.
I was not on the cover
of the top mathematics publications
because I was good looking.
I was on the cover
of the top mathematics publications
because I contributed to mathematics.

24.1.2 Why Algebra is the Recurring Decimal

948

Why is the word “algebra”
the recurring decimal
in school reports about the contributions
of **Philip Emeagwali** to mathematics?
The reason algebra
is my recurring decimal
across school reports
and **across**
my ensemble of 65,536 processors
is that I invented
partial **difference** **algebraic** equations
and that I used those algebraic equations
to approximate
the nine abstract, continuous
partial **differential** equations
of modern calculus
that I also invented.
Those Emeagwali’s equations, in turn,
encoded the Second Law of Motion
of physics
that were discovered three centuries earlier.

In the finite difference discretization
of the continuous
initial-boundary value problem
of mathematical physics,
extreme-scale algebra
is the recurring decimal.
Extreme-scaled algebraic computations
were executed
by the massively parallel processing
supercomputer
that computed
the weather forecast
in your evening television news.
Algebra
is about using the known
to know the unknown.
Today's weather that is a known quantity
is used to forecast tomorrow's weather
that is an unknown quantity.
For that reason, algebra—which is used to
know the unknown from the known—

must be used to know tomorrow's weather
from the known weather of today.
So, every television viewer
that is watching the weather forecast
is a secondary consumer
of the algebraic knowledge
that was massively parallel processed
across the modern supercomputer.

24.1.3 My Initial-Boundary Value Problem

Before my experimental discovery
that occurred at
10:15 in the morning New York Time
of Tuesday the Fourth of July 1989,
only one processor
within an ensemble of
65,536 processors
could be harnessed to solve
the toughest initial-boundary value problems
in computational physics
and computational mathematics.

That singular processor was *de facto*
merely a computer.
After my experimental discovery
of the Fourth of July 1989
of how all those 65,536 processors
could be harnessed
and used to solve the toughest problems,
that global network of processors,
that was also a new internet,
was *de facto* a new supercomputer.
That experimental discovery
is the reason the name
"**Philip Emeagwali**"
came up first
when Google was queried
with the search phrase:
"**contributions** to the development
of the computer."
At the granite core
of my experimental discovery
was how and why
I emailed 65,536
initial-boundary value problems

and why I emailed them
with a **one-to-one** correspondence
to 65,536 processors.
Each email I sent and received
from each processor
contained data and answers
for and from
the **Emeagwali** new system of equations of algebra
that I had to solve.
The **Emeagwali** new system of equations
arose from
my system of coupled, non-linear,
time-dependent, and state-of-the-art
partial differential equations
of a new calculus
that I encoded, in the early 1980s,
into a set of laws
of physics.
The **Emeagwali's** new system of
partial differential equations
contained the most advanced
and the most important expressions

in modern calculus.
The **Emeagwali's** new system of
partial differential equations
is mathematically equivalent
to the system of
partial differential equations
that were used to define
one of the Seven Millennium Problems
of mathematics.
The slight difference
between the **Emeagwali's** new system of
partial differential equations
and the system used to define
the millennium problem
was in the real-world domain
of the initial-boundary value problem
that was governed by the system of
partial differential equations.
The **Emeagwali's** new system
of partial differential equations
governs the subterranean motions of fluids—
such as oil,
water, and gas—that flow underneath

the surface of the Earth.
The millennium equations
govern the motions of fluids
that flow **across**
the surface of the Earth—such as rivers,
lakes, and oceans—and govern
the motions of fluids
that flow above the surface
of the Earth—such as the wind
and the moisture
that condenses to rain and snow.

24.1.4 Wanted: A Polymath

The initial-boundary value problems
that I experimentally discovered
how they could be solved
were called grand challenge problems
for a good a reason.
I experimentally discovered
how to solve the toughest problems
in extreme-scale computational physics
and solve them

across a new internet
that is an ensemble of
commodity-off-the-shelf processors
that are identical
and that are equal distances apart
from each other.
That project was a massively
parallel processing supercomputer research
that should have been conducted
by a large multidisciplinary team.
Or conducted
by a lone wolf parallel processing
supercomputer scientist
who is both a **jack-of-several**-sciences
and a master of **several sciences**.
To succeed
in experimentally discovering
the new massively parallel processing
supercomputer,
was like looking God
in the face.
The lone wolf supercomputer scientist
must be a **polymath**,

instead of a mathematician.
The reason the grand challenge problem
was impossible
for a mathematician
that was trained only as a mathematician
to solve
is that **that** mathematician
could only think on his blackboard.
The **polymath**
visualizes the grand solution
to the grand challenge problem
and visualizes the time-to-solution
across boards,
and visualizes that time-to-solution
from the storyboard
to the blackboard
to the motherboard
and, finally, **across** motherboards.
I—**Philip Emeagwali**—used my
multidisciplinary knowledge
that I acquired over two decades
of specialized training,
that spanned

from geology to meteorology,
from calculus to algebra,
from computer to internet
and I used that **multi**disciplinary training
to conceive and to execute
the crucial supercomputer experiments
that led to my experimental discovery
on the Fourth of July 1989
of how to massively parallel process
across a new internet
that's a global network of
64 binary thousand
commodity-off-the-shelf processors
that are identical
and that are equal distances apart
from each other.

24.1.5 Struggles Against Dogma

Back **in the 1970s and '80s,**
almost every vector processing

supercomputer scientist
believed that parallel processing
is a huge waste of everybody's time.
So, I was executing my massively
parallel processing experiments
and executing them
against the orders
of the leaders of thought
in the world of computing
—such as the **Steve Jobs**
of personal computing—
and against the opinions
of the leaders of thought
in supercomputing,
such as **Gene Amdahl** and **Seymour Cray**.
In the 1970s and '80s,
the *terra incognita*
that was the emerging field of
massively parallel processing
supercomputing
was as empty as a ghost town
that had only one permanent resident.
I—**Philip Emeagwali**—

was that permanent resident
of the farthest frontier of supercomputing
called massively parallel processing.
In the 1970s and '80s,
if you could find your way
to any massively parallel processing
supercomputer,
its administrator will deem you worthy
and grant you a supercomputer account
to become its lone wolf programmer.
Because the internet of the early 1980s
was then uncrowded,
I had an unusual email address
from the early 1980s.
That email address
had no dot com **suffix**.
In the mid-1980s,
I had the email address spelled
emeagwal @ think dot com.
Emeag**wal** was spelled like my last name
without the last letter "i."
Think dot com

was the second registered
dot com **suffix**.

24.1.6 Visualizing a Small Copy of the Internet

That global network of
64 binary thousand
commodity-off-the-shelf processors
that I experimentally discovered
that it could be programmed
to solve the toughest problems
in computational physics
was a new **internet**.
That new internet
was a small copy
of a never-before-understood Internet,
that had only 65,536 processors
around a globe
instead of billions of computers
around a globe.
I visualized each of my
two-to-power sixteen
commodity processors

as equal distances apart
from each other
and around a globe
in a sixteen-dimensional hyperspace.
And I visualized my ensemble
of processors
as evenly distributed **across**
the hypersurface of a hypersphere
in a sixteen-dimensional universe.
I visualized my ensemble
of processors
as outlining a new internet
that I visualized
in my sixteen-dimensional universe.

24.1.7 David Versus Goliath

I—**Philip Emeagwali**—was the **David**
from the world of the massively
parallel processing supercomputer
that was ridiculed and mocked
for challenging the **Goliath**

—named **Seymour Cray**—who designed 962
seven in ten supercomputers
in the world of the vector processing
supercomputer
of the 1970s and ‘80s.
I visualized my massively
parallel processing supercomputer
as my sling shot
that is a small copy of the Internet
that can shoot
65,536 small pebbles
from its as many processors.
Those pebbles were my metaphors
for the as many
initial-boundary value problems
of modern calculus
and computational physics.
I visualized shooting
all the 65,536 small pebbles
at once.
I can only record the fastest computations
and record them **across**
65,536

processors
and record them
by throwing all my rocks
at once,
instead of throwing them
one at a time.
In the 1970s and '80s,
I was the **David** of supercomputing.
I was ridiculed and caricatured
by well-regarded supercomputer scientists.
I was called a "lunatic"
and dismissed from research teams
that believed that
all supercomputers
must do only one thing **at a time**.
Seymour Cray
—the **Goliath** of supercomputing—
believed that all supercomputers
should compute only one thing **at a time**.
Seymour Cray
was armed with one big sword.
Seymour Cray's sword
was my metaphor

for his vector processing supercomputer.
Seymour Cray's most famous quote is this:
[quote]
"If you were plowing a field,
which would you rather use?
Two strong oxen
or 1024 chickens?"
[unquote]
As reported in the June 20, 1990 issue
of *The Wall Street Journal*,
I—**Philip Emeagwali**—
experimentally discovered that
65,536 chickens
that learned to work together,
or work in parallel,
can plow more field
than the strongest ox that works alone.

24.1.8 How I Invented a New Internet

I'm **Philip Emeagwali**.
I'm the subject of school reports
because I invented

a new supercomputer
that was the precursor
to the modern supercomputer.
I invented a new supercomputer
that is a small copy
of a new internet.
The new internet that I invented
is defined and outlined by an ensemble
of 65,536
commodity-off-the-shelf processors
that are identical
and that are equal distances
apart.
That new internet
is complex, abstract, and a mystery.
The 65,536 processors of my new internet
were married together
by 1,048,576 bi-directional email wires
and married together
as a new supercomputer
that computed cohesively
and did so as one new integrated
supercomputer

and communicated seamlessly
as one new internet.
My 64 binary thousand processors
that outlined my new internet
communicated via emails
and did so with a complexity
that I cannot completely describe
in words alone.
Nor can I completely describe
my processor-to-processor email exchanges
and completely describe them as equations
on a blackboard alone
or completely describe them as algorithms
on a motherboard alone.
I began supercomputing at age 19
on June 20, 1974
in Corvallis, Oregon, United States.
I was the lone wolf
and the only full time programmer
of the fastest supercomputer
of the 1980s.
Today, the fastest supercomputer
costs the budget of a small nation.

The fastest supercomputer
is programmed by thousands
of supercomputer scientists.
The fastest supercomputer
occupies the space of a soccer field.
The **Holy Grail** of the fastest possible
supercomputer
is to marry together
all the processors in the world
and marry them
to all the computers in the world
and marry them
to all the supercomputers
in the world
and marry processors and computers
and supercomputers together
and as a never-before-seen internet
that will become a never-before-seen
planetary-sized supercomputer
that will turn our science fiction
to our descendant's non-fiction.

I'm **Philip Emeagwali**.

24.2 How I Named a New Internet

24.2.1 How I Named a New Internet

I'm **Philip Emeagwali.**
I'm the massively parallel processing
supercomputer scientist
that conducted research alone
and conducted it from the age of nineteen
in Corvallis, Oregon, United States
to the age of 35
in Los Alamos, New Mexico, United States.
To the supercomputer scientist,
Los Alamos, New Mexico
is the capital of supercomputing.
Prior to my experimental discovery
of the massively parallel processing
supercomputer
that occurred on the Fourth of July 1989,
it was said that

parallel processing
is a beautiful theory
that lacked experimental confirmation.
Prior to the Fourth of July 1989,
I was the unknown supercomputer scientist
who told his massively parallel processing
supercomputer story alone
and told it to no supercomputer scientist
in particular.
In fact, my 1,057-page report
that I distributed
to vector processing
supercomputer scientists
of the 1980s
and that described my new supercomputer
was, at first, thrown into the trash.
After my experimental discovery
of the massively parallel processing
supercomputer
that occurred on the Fourth of July 1989,
I became a known supercomputer scientist
and those that threw my 1,057-page report
into the trash

wanted to become my new best friend
and clamored
to retell the story
of how I experimentally discovered
the massively parallel processing
supercomputer
that is a new internet.
As I became more known,
I discovered that
many insidious voices
were clamoring to retell my story
and to retell it
in their visions,
rather than in my original vision.
I discovered that
their thousand secondary voices
can drown my primary voice.
I discovered that
the story of my new supercomputer
that is not a computer *per se*
but that is a new internet *de facto*
was reduced to a cacophony
of secondary voices.

I want to redeem my story
and reclaim my voice
and make my voice
the loudest voice
in the world of the modern supercomputer
and make my voice
to be the most continuously heard voice
in the history of the Internet.

24.2.2 The Magic Zone: Naming My New Internet

Each of my processor—within my ensemble
of 65,536 processors—
had its unique name
that's also its unique email address
that's sixteen bits long.
I used a binary reflected code
to generate
my 64 binary thousand unique names
that were each
a unique string of
sixteen zeroes and ones.
With the binary reflected

972

internet naming scheme
that I used,
if two email addresses
differed by only one bit,
then the processors
that corresponded
to those two email addresses
differed by only one bit.
And those two processors
were directly connected.
That connection
allows nearest-neighbor email
communications that maximizes
the speed I could attain
while executing
my floating-point arithmetical operations.
My sixty-four binary thousand emails
travelled across
one binary million,
or one million, forty-eight thousand
five hundred and seventy-six [1,048,576],
bi-directional email wires.

24.2.3 The Fastest Supercomputer

I had theorized that speed increase
of a factor of sixty-four binary thousand
and, therefore, I was expecting
to experimentally discover
the world's fastest speed in computation.
Even though I was expecting
the world's fastest speed in computation
in 1988,
I was shocked
and in disbelief
when world's fastest speed in computation
popped up
on the computer monitor
of my four-brained workstation
that front-ended
my 64 binary thousand-brained
massively parallel processing machine.
Even though my experimental discovery
matched my companion theoretical discovery
that preceded it, a decade earlier,
I developed an **imposter syndrome**.

I developed it, in part, because
I was black and African.
I developed an **imposter syndrome**
because I was ostracized
and only one white computational scientist
attended the first lecture
on massively parallel processing,
that I gave in November 1982
in Washington, DC.
Because I developed an **imposter syndrome**,
I convinced myself
that I made a mistake.
Due to self-doubt,
I wasted half a year—the last half of 1989—
searching for the mistake
in my new massively parallel processing
supercomputer code
that did not exist.
My Eureka moment
felt like a jolt of electricity.
I screamed loud enough for others to hear.
Eiwooooooooooooooooooo
Eiwooooooooooooooooooo

Eiwoooooooooooooooooooo.
I screamed in my ancestral Igbo language
of Onitsha, Anambra,
the former East-Central State of Nigeria,
West Africa.
On the Fourth of July 1989,
I experimentally recorded
a world record computational speed of
twenty thousand calculations
per second
per processor.
Then my screaming continued for half a year
with each small refinement
that led to a new world record.
My computational speed
peaked at forty-seven thousand
three hundred and three [47,303]
calculations per second
for each processor
that I massively parallel processed through.
My supercomputer speed
that made the news headlines
in 1989

976

was achieved
in a high-level programming language.
I discovered that
my new massively
parallel processing supercomputer
could compute four times faster
if I had used a low-level
programming language
that could have enabled me
to use an **assembly-coded** routine
and use it to reduce the execution time
of my most computation-intensive kernels.
The supercomputing world record
that I recorded
in high-level programming language
and recorded in 1989 totaled
3.1 billion calculations per second
across my parallel processing machine
that was my new internet.
That fastest computation was recorded
across a new internet
that's a global network of

sixty-five thousand
five hundred and thirty-six [65,536]
processors
that's *de facto*
a massively parallel processing
supercomputer.

24.2.4 How I Visualized My New Supercomputer Solutions

The eternal quest of humanity
for faster computing aids
that began with the abacus in ancient China
remains the Holy Grail of computing.
Yet, that quest had only one **paradigm shift**.
I **paradigm shifted**
from computing within
only one computer
powered by only one processor
to supercomputing **across**
a new internet that is a global network of
64 binary thousand processors,

or supercomputing **across**
a global network of as many computers,
or supercomputing **across**
a new internet
that computes by massively
parallel processing.
In the 1970s and '80s,
the grand challenge of communicating
across
that new internet
was ridiculed and mocked as follows:
Parallel processing
—or supercomputing many things
at once,
instead of computing
only one thing **at a time**—
is a beautiful theory
that lacks an experimental confirmation.
The June 20, 1990 issue
of the *Wall Street Journal*
mentioned that I—Philip Emeagwali—
experimentally discovered
how to massively parallel process

and how to simulate the multiphase
subterranean motions
of crude oil, injected water, and natural gas
flowing a mile-deep inside an oilfield
and how to simulate
the production reservoir
by dividing the oilfield
into sixty-five thousand
five hundred and thirty-six [65,536]
equal parts
that I emailed
to sixty-five thousand
five hundred and thirty-six [65,536]
processors.
That experimental discovery
was also the cover story
of the May 1990 issue
of the *SIAM News*,
the flagship publication of record
of the research mathematics community
that was published by the Society
for Industrial and Applied Mathematics
and that was written

by research mathematicians
and written
for research mathematicians.

24.2.5 My First Awareness—Eureka! Moment

My first newsworthy speedup result
that occurred on the Fourth of July 1989 was
a quantum leap
in supercomputer performance.
It was either a zero speedup
or a 65,536-fold speedup.
There was no speedup
between zero and 64 binary thousand.
I recorded a 65,536-fold increase
in both the speeds of email communication
and arithmetical computation.
I experimentally discovered
the fastest supercomputer
when I experimentally discovered
the fastest super computations

and discovered them
across
my hyper-global network of
65,536 processors.
My network
networked processors,
or as many computers,
around a hyper-globe
with 1,048,576
bi-directional email wires.

A world record speed
in arithmetical computation
and email communication
puts a man and his machine
inside the *terra incognita*
that is the magic zone
of the super computer
and the new Internet.
It hits the sole programmer
of all those 65,536
processors
as a jolt of electricity.

Right away, my few **confidantes**
said I will be famous.
The **paradigm shifting** technology,
named the massively
parallel processing supercomputer,
that I experimentally discovered
changed the way we do the most
computation-intensive arithmetic
on the modern supercomputer.
The supercomputing paradigm
shifted away from executing
floating-point arithmetical operations
that arose from
extreme-scale algebra
and shifted away from
executing those operations in the singular.
The supercomputing paradigm
shifted to dividing the computation-intensive
initial-boundary value problem
into 65,536 subsets of arithmetical problems
that were equal to each other
—and that were equal in terms of
total floating-point arithmetical

983

operations count—
and sending those 65,536 email messages
that each contained
a subset of arithmetical problems
and following each email
with floating-point arithmetical
computations
that solves the companion
initial-boundary value problem
and, most importantly,
executing the computation
and the communication
in the plural senses.
I realized the profound
and the **paradigm shifting** meaning
of that never-before-recorded
65,536-fold increase in speed.
A grand challenge problem
that took 65,536 days,
or 180 years,
to solve
will now take only one day

to solve.
That experimental discovery
opened the door
to today's world of massively
parallel processing supercomputing
in which
a grand challenge problem
that takes 30,000 years,
to solve on one processor
will now take only one day
to solve **across**
ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
commodity-off-the-shelf processors
that is at the core of a
massively parallel processing
supercomputer.
For those still enmeshed
in the old paradigm of sequential computing,
their 30,000 years
of sequential computer work
is merely one day

in my new paradigm
of massively parallel supercomputing.
I discovered the paradigm
in which the boundary
between the computer and the internet
is blurred.

24.2.6 Emails Appeared Invisible

When I began supercomputing
—on June 20, 1974—I envisioned
a planet-sized global
network of computers
that was the precursor
to the Internet.
In subsequent years, I invented
a new internet
that I called a HyperBall
that was described in the book titled:
"History of the Internet."
I also invented a second new internet

that I called a Cosmic Ball. 986
In the mid-1970s, my new internets
remained science fiction.
But on the Fourth of July 1989,
I constructively reduced
that HyperBall science fiction
to **nonfiction**
and I did so
when I became the first person
to experimentally discover that
an ensemble of the slowest
65,536 processors
in the world
can be harnessed to compute faster than
the world's fastest supercomputer
and do so while solving
the toughest problems
in extreme-scale computational physics.
At its logical core, my technological quest
to record the fastest speed
in supercomputing

and to record it **across**
a new internet
that was my ensemble of 65,536
commodity processors
was **indirectly**
a quest to find the elusive
email communication path
and to know the unknown
communication primitive,
and was a quest to find the
communication path
from each commodity processor
to its sixteen nearest-neighboring
commodity processors.
My email communication breakthrough
occurred when I discovered
a **new paradigm**
in email communication,
namely, how to use a new
addressing scheme
that was counter-intuitive

but self-relative.
In the **old paradigm** of computing,
supercomputer scientists
used an addressing scheme
that was intuitive and absolute.
With my self-relative addressing scheme,
I repeated the email messages
that I sent to and received from
my two-to-power sixteen
commodity processors.
With my mathematical discovery
of the self-relative addressing scheme,
my self-doubt evaporated
like dew in the sun.
After that mathematical discovery,
my confidence soared
and I started seeing myself
as the first
massively parallel processing
supercomputer scientist
that could experimentally discover

how to make the impossible-to-compute
possible-to-compute.
To the 25,000 vector processing
supercomputer scientists
of the 1980s
that lacked my new knowledge
of self-relative addressing scheme,
Philip Emeagwali seemed like a lunatic
that was pursuing the elusive goal
of massively parallel processing
across an ensemble of
commodity processors
that was a global network of
65,536 commodity processors
and that was a new internet.
I invented a precise, minimalist code
with email communication primitives
that belies its power.
My code was so minimalist

that my email messages
to and from sixteen-bit long addresses
appeared invisible.
I wanted my data
to be transported
from each vertex of the hypercube
—that was my metaphor
for a commodity processor—
to its sixteen nearest neighboring
vertices
that had a **one-to-one** correspondence
with sixteen nearest-neighboring
commodity processors.
I visualized my emails
as flowing **quietly**
and **seamlessly**
across the surface of a hypersphere
that was embedded
in a sixteen-dimensional universe
and doing so
with the accuracy a cat deploys

to tiptoe deftly
through a laid dinner table.

24.2.7 64 Binary Thousand Automated Emails

I, Philip Emeagwali,
experimentally discovered
how to email 64 binary thousand
challenging initial-boundary value
problems
of modern calculus
and extreme-scale computational physics
and how to email them
to 64 binary thousand processors.
I, Philip Emeagwali,
experimentally discovered
that those ensemble of processors
define and **outline**
a new internet
that I visualized
as my small copy

or **blueprint** or **prototype**
of the Internet.
I sent each of my primitive email
to a sixteen-bit address,
or a unique string of
sixteen zeroes and ones.
Each email address had no @ sign
or dot com suffix.
Each email contained
five sets of data.
I visualized each email
as having **five subject lines**.
I visualized each emailed processor code
as a computation-intensive
initial-boundary value problem
in modern calculus
or extreme-scale computational physics
that I sent
to each of my two-raised-to-power
sixteen processors.
I visualized each emailed code

as sent along sixteen times
two-to-power sixteen
bi-directional edges
of the cube in the sixteenth dimension
that had a **one-to-one** correspondence
with as many email wires
that **I visualized**
as short wires
printed onto circuit boards,
or as long wires
comprised of fiber optic cables.
I visualized each email as sent
with no message body.
I visualized around a globe
in the sixteenth dimension
because my problems and their data,
or email messages,
were contained in three subject lines
when received by sixty-five thousand
five hundred and thirty-six [65,536]
processors

that **I visualized**
as uniformly distributed
around a globe
in a sixteen-dimensional hyperspace.
After each email was received,
the two subject email lines
that contained the sending
and the receiving
sixteen-bit long email addresses
become useless.
I discarded those two email addresses
after each email was received.
I visualized my new internet
as outlined by sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors
that were **identical**
and that were **equidistant**
from each other.
Each processor
was like a dim light

in a sea of darkness.
But when supercomputing together
as one seamless, cohesive
supercomputer,
those sixty-five thousand
five hundred and thirty-six [65,536]
processors
became as bright as the sun.

24.2.8 Equidistant Computers on a Hypersphere

I communicated
around my new internet
that is a global network of
65,536 processors
and I did so by email
and I did so before each set
of arithmetical computations
and I did so by imagining myself
as 65,536 travelers

in hyperspace.
As an aside,
any global network
of commodity processors
that are identical
and that are equal distances **apart**
can be flattened
from three or sixteen dimensions
onto a two-dimensional motherboard.
Topologically,
and to the supercomputer scientist
or mathematician,
that sixteen dimensional network
of commodity processors
still remains
a global network of processors.
That global network
is electronically soldered
onto a two-dimensional surface.
That global network
on a two-dimensional motherboard

is not renamed
a two-dimensional network.
To rename that sixteen-dimensional
network
as a two-dimensional network
will be akin to renaming
a sixteen-dimensional hypercube
that was drawn on paper
as a square.
Or akin to renaming
a sixteen-dimensional hypersphere
that was printed on paper
as a circle.
Put differently,
I had one neural, human brain
enhanced by 65,536 electronic brains.
My email travel paths
were along communication email wires
that encircled
the surface of a globe in hyperspace
that **I visualized**

as short wires
printed onto circuit boards,
or as long wires
comprised of fiber optic cables.
My destinations
were my **equidistant** processors
that **I visualized**
as distributed
on the surface of that **hypersphere**.

24.2.9 Computing in the Sixty-Fourth Dimension

I imagined myself, **Philip Emeagwali**,
traveling
into higher and higher dimensional
universes.
The farthest I could travel
was to and from
two-to-power sixty-four
equidistant processors.
A new internet

999

connected those processors
and connected them
with sixty-four (64) times
two-to-power sixty-four (64)
identical email wires
that **I visualized**
as short wires
printed onto circuit boards,
or as long wires
comprised of fiber optic cables.
In the sixty-fourth (64^{th}) dimensional universe,
I had an Eureka! moment.
I discovered
it will be presumptuous
of our children's children
to imagine they could construct
a hyper internet
that will talk
in the sixty-fourth dimension
to think

two-to-power sixty-four times faster.

1000

24.2.10 Philip Emeagwali Internet

I was in the news headlines
as a result of my experimental discovery
of the massively parallel processing
supercomputer.
That experimental discovery
occurred on the Fourth of July 1989.
I was in the news headlines
because I brought a new face,
a new voice, and a new vision
to the story of the development
of the modern supercomputer
that is not a computer *per se*
but that is a new internet *de facto*.
That experimental discovery
of the massively parallel processing
supercomputer
that is the precursor

to the modern supercomputer
came from both intellect and intuition.
For me, **Philip Emeagwali**,
that precursor to the modern supercomputer
is the coming together
of my 35-year life journey and experiences.
Prior to my experimental discovery,
the use of the massively parallel processing
supercomputer
and how to use the technology
to solve the toughest problems
in extreme-scale computational physics
was not taught in any university
or tested in any laboratory.
Nor can the technology be learned.
Instead, the new massively
parallel processing supercomputer
that solves the toughest problems
in abstract calculus
and extreme-scale computational physics
and solves them faster
than the vector processing supercomputer
had to be seen for the first time

and seen
on that Fourth of July 1989.
By definition, to discover
is to see something that was previously unseen.
A scientist on a re-search for new knowledge
is a truth seeker, especially new truths.
A supercomputer scientist
on a re-search for a faster supercomputer
is seeking a new supercomputer
and, especially, a new computer.
The discoverer
becomes the first truth seeker.
The scientific re-searcher
is on his or her hero's quest
for the previously unseen truth.
Our never-ending quest
for the fastest possible supercomputer
has become our journey
to the frontier of human knowledge.
That never-ending quest
has become a self-directed evolution
in which we are both the creator

and the created. 1003
That journey to the end of knowledge
will force our post-human descendants
of **Year Million**
to address the larger question
of who we are
and where do we want to go.