26 A Father of the Internet

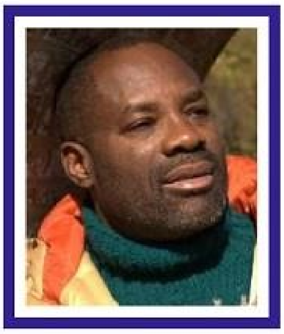

Philip Emeagwali Lecture 180127-1
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

26.1 The Troubadour of Supercomputing

26.1.1 The Troubadour of the Supercomputer

The experimental invention

of the massively parallel processing
supercomputer
—that solves many problems **at once**,
instead of solving only one problem
at a time—
and its absorption into new computers
and into new supercomputers
is one of computing industry's
most hopeful narrative.
Prior to the Fourth of July 1989,
the naysayers
within the supercomputer industry
demanded that
parallel processing adapt to them,
instead of them adapt to
parallel processing.
In the supercomputing community
of the 1970s and '80s,
the massively parallel processing
supercomputer
was greeted with skepticism
and was ridiculed as a huge waste of
everybody's time.

The inventor
is the troubadour
of technological knowledge.
The inventor of parallel processing
is the troubadour of supercomputing.
But unlike the troubadour
that was the medieval lyric poet
of the 11^{th} to 13^{th} centuries
who writes verse to music
the inventor of the modern supercomputer
writes never-before-seen emails
to a never-before-seen
ensemble of processors
that outline a never-before-seen internet
that was never-before-understood
as a global network of processors
that is the fastest supercomputer,
de facto.

26.1.2 Seeing the Big Picture

One unexpected benefit
of being a black and African inventor

and, therefore, forced to invent
as a lone wolf
was that it enabled me to have
a **coherent vision**
that I centered on my new internet.
In the 1970s and '80s,
I developed a body of work
in which the elements were **disparate**
but, yet, fit together
as one cohesive whole
that's a new supercomputer
and that's *de facto* a new internet.
For instance, I discovered that
what I learned or discovered
in previous **boards** translated over,
in whole or in part,
and that far more important,
that I was re-telling
the same story of the motions of fluids
that were governed
by the laws of physics.
I was the **troubadour of supercomputing**
who translated a story

1049

in extreme-scale computational physics 1050
and translated it
from the blank storyboard
to a story in modern mathematics
on the blank blackboard
and translated that story
in a **never-before-seen** calculus
to a story in extreme-scale algebra
and to a story
that resulted in my execution
of immensely computation-intensive
floating-point arithmetical calculations
executed on the motherboard
and translated that story
on the motherboard
to 64 binary thousand stories
across motherboards
and continued to translate those stories
to the boardrooms
to the classrooms
and to your living rooms.

26.1.3 My Supercomputing Style

Amongst research supercomputer scientists,
that style of speaking and thinking
was **distinctly** mine.
An artist often uses the same style
to portray different subjects.
An artist may be an **impressionist**
like the late 19th century Frenchman
Claude Monet
or a **surrealist**
like the early 20th century painter
Salvador Dali
or a **modernist**
like the 1920s and '30s painter
Henri Matisse
or a **sculptor**
like the 20th century
Ben Enwonu
of my hometown of Onitsha, Nigeria.
I'm an extreme scale computational physicist
that uses the laws of physics
to **digitally replicate**

the global motions of fluids
that enshroud the Earth.
I **digitally replicated** those motions
across my new internet
that's a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
equidistant processors
that encircled the globe
that's my metaphor for the Earth.

26.1.4 My Supercomputer Origin Story

My quest
for the fastest supercomputer
was a sixteen-year-long journey
that began on June 20, 1974
and began on
a sequential processing supercomputer
that was at
1800 SW Campus Way,
Corvallis, Oregon, United States.

My experimental discovery
of the fastest supercomputer
ended on a parallel processing machine
in **Los Alamos**, New Mexico, United States.
Los Alamos is a quiet, small town
that's often referred to
as the **capital of supercomputing**
but is better known
as the **birthplace of the atomic bomb**
that was dropped in Hiroshima, Japan.
When I began supercomputing,
on June 20, 1974,
I lacked clarity on what my new internet,
or global network of
64 binary thousand processors,
—that's a parallel processing machine—was.
My early vision of a small copy
of that new internet
was a mere idea
—a small seed of an Iroko tree—
instead of the grown
160-foot tall Iroko tree
that it became sixteen years later.

I began in 1974 with a **semi-abstract**, 1054
theorized HyperBall internet.
That new internet
was a global network of computers.
That new internet
evolved from one CPU,
or processor,
to **across** sixteen years
to become a very realistic
new internet
that used the cube
in the sixteenth dimension
as its metaphor.
I experimentally discovered
that new internet
in the sixteenth dimension
and discovered it
as a global network of
two-raised-to-power sixteen,
or 64 binary thousand,
processors,
or as a global network of
sixty-five thousand

five hundred and thirty-six [65,536]
computers.
I was asked:
"When the games are over,
how do you want to be remembered?"
I want to be remembered
as the supercomputer inventor
that connected those dots (or vertices),
so to speak.
I want to be remembered
as the supercomputer discoverer
that told the **cohesive** story
and discovered those internets
as *de facto* **cohesive** computers
that are the **fastest**.

26.1.5 The Processors of My New Internet

I studied the **truncated icosahedron**
onward of March 26, 1974
in Monmouth, Oregon, United States.
The **truncated icosahedron**

and the cube in the sixteenth dimension
were the geometrical shapes
that inspired my new internets
and that I framed my new internets around.
I used the **truncated icosahedron**
as my inspiration
for my new Cosmic Supercomputer
and for inventing my new internet
and inventing it
with a one-to-one correspondence
between its processors,
or computers,
and the as many vertices
of the **truncated icosahedron**.
I invented that new Cosmic Supercomputer
that was *de facto* the **first** internet
that is a global network of
thousands of equidistant processors.
That **Philip Emeagwali** theorized internet
was inspired by sixteen years of studying
how and why
bees honeycombs
are hexagonally-structured.

I used the cube in the sixteenth dimension
as my inspiration
for inventing the first internet
and inventing it
with a one-to-one correspondence
between its two-raised-to-power sixteen
processors,
or its 64 binary thousand computers,
and the sixty-five thousand
five hundred and thirty-six [65,536]
vertices
of the cube in the sixteenth dimension.

26.1.6 The Philip Emeagwali Internet

Also, I used the geometrical object,
called the **truncated icosahedron,**
as my design inspiration for inventing
a never-before-seen
global network of processors
that is a new internet.
I constructively reduced that new internet
to practice

and to a new supercomputer
that has a **one-to-one correspondence**
between its email wires
and the as many **bi-directional** edges
of that **truncated icosahedron**.
I named that new internet
that gave rise to a never-before-seen
supercomputer
a Cosmic Supercomputer.
The technology is also called
the **Philip Emeagwali** internet
and the **Philip Emeagwali** Supercomputer.

26.1.7 The HyperBall Supercomputer

To invent a second internet,
I used the cube
in the sixteenth dimension,
called the hypercube,
as my design inspiration.
I used the hypercube
to invent a new internet

that is a never-before-understood
global network of processors.
That new internet
outlined and defined
a HyperBall Supercomputer
that has **a one-to-one correspondence**
between its sixteen times
two-raised-to-power sixteen
bi-directional email wires,
or its one **binary million** email wires,
and the 1,048,576
bi-directional edges
of the cube in the sixteenth dimension.
My geometrical clarity
that emanated from both
the truncated icosahedron
and the hypercube
made the two new internets
that I invented
both **visible** and **concrete**.
My geometrical clarity
allowed my contributions
to the development

of the most massively parallel processing supercomputer 1060
that is the fastest computer of today,
or the modern supercomputer,
to stand out on their own.
New inventions yield new words
and newer languages.
The word “supercomputer”
was first used in 1967.
The word “internet”
was not in the *Daily Times* of Nigeria,
that was the newspaper
that I read
while growing up in the early 1970s
in Onitsha, East Central State, Nigeria.
New inventions yield new illustrations.
Therefore, the new
Philip Emeagwali internet,
or the **Philip Emeagwali**
Cosmic Supercomputer,
must yield a **never-before-seen** illustration
of that new internet
and that new supercomputer.

The illustrations of my HyperBall
global network of processors
were **beautiful**.
There were widely copied
without giving credit to **Philip Emeagwali**.
You've seen illustrations
that were inspired by
the **Philip Emeagwali** internet
and seen them on magazines
or television or in schoolrooms.
But you didn't realize that
Philip Emeagwali
first illustrated them for his new internet.
The earliest illustrations
of my small copy of the internet
that was **camouflaged** by
64 binary thousand processors
became like a **wild horse**
that belongs to any horse rider
that finds it.
But like a wild horse,
the new supercomputer
that I illustrated was difficult to ride.

26.1.8 Is a Discovery Made by Luck or Genius?

It's mistakenly believed that
serendipity, or luck,
must always play a role in making
a scientific discovery
or in making a technological invention.
There's nothing **serendipitous**
about my experimental discovery
that occurred on the Fourth of July 1989
of how and why
the massively parallel processing
supercomputer
could be used to execute
the fastest floating-point arithmetical
computations
and record that fastest calculations **across**
an ensemble of millions upon millions
of commodity-off-the-shelf **processor**s
that are identical
and that were equal distances
apart.

Prior to my experimental discovery
of massively parallel processing
that occurred on the Fourth of July 1989,
the **central** debate of supercomputing
was that supercomputer scientists
accepted Amdahl's Law **limit**
that **limited** the maximum speed up
—to a factor of eight—that I could attain
across my global network
of 64 binary thousand
commodity processors.
Supercomputer scientists
accepted **Gene Amdahl**'s hypothesized
but **untrue** maximum speed-up.
That acceptance meant that
I was attempting to do the **impossible**
and to attain a linear speedup
and to do so
across my global network of
64 binary thousand
commodity processors
that computed in parallel,
or solved 64 binary thousand

computational physics problems
at once.
The supercomputing community
uncritically accepted Amdahl's Law
and accepted it
without demanding for the evidence
that it is true,
or that it is even a law.
For me, **Philip Emeagwali**,
my mathematical journey
to the abstract, unknown world
of the partial differential equation
of modern mathematics
began from the **embers**
(that is a small piece of burning **charcoal**)
of the 30-month long
Nigeria-Biafra Civil War.
My journey
was from the war front
to the science front.
My technological quest
for the then unknown massively parallel
supercomputer

—that can be used
as an instrument of physics
and used for solving
64 binary thousand
computational physics problems
and used for solving them **at once**—
ended with 64 binary thousand
startling voices
in Silicon Valley, California.
My experimental discovery
of parallel processing
had to be preceded
by my theoretical discovery
of parallel processing.
My experimental discovery
of the **pre-cursor**
to the modern supercomputer
was the end product
of my mixture of **abstract** calculus
and advanced algebra
and fastest parallel processing
supercomputing.
The experimental discoverer

of parallel processing 1066
must be a **triple threat**
—that is simultaneously
at the frontier of knowledge
in computational physics,
at the frontier of knowledge
in computational mathematics,
and at the frontier of knowledge
in massively parallel processing
supercomputing.
That discoverer must possess
both the intellect and the knowledge,
as well as the financial resources
needed to acquire the most expensive
supercomputer in the world.
The fastest supercomputer
costs the budget of a small nation.
With that price tag, I cannot log into
the fastest supercomputer
by serendipity, or luck.
As an **extreme-scale**
computational physicist,
I had to know what I was doing

to be the sole programmer
of the fastest supercomputer
of the 1980s.
As a **large-scale**
computational mathematician,
I had to know what I was doing
on each of the sixteen supercomputers
that I programmed alone
in the sixteen years
onward of June 20, 1974.
Why was a young, black, and African
supercomputer scientist
given charge of sixteen
massively parallel processing
supercomputers?
The reason was this:
In the 1970s and '80s,
vector processing supercomputers
—that then costs forty million dollars each—
were reserved for only
white supercomputer scientists.
At that time, the parallel processing
supercomputer

was abandoned and dismissed and mocked
as a waste of everybody's time.
In the 1980s, I alone
was logged into
sixteen massively parallel processing
supercomputers.
For me, **Philip Emeagwali,**
and on the Fourth of July 1989,
there was no accidental discovery
of massively parallel processing.
The reason was that I executed
the most computation-intensive
floating-point arithmetical calculations
and calculated how to solve
a grand challenge problem
in extreme-scale computational physics
and solve it
at the fastest recorded speed.
How I accomplished that fastest speed
is too detailed to be described in one hour
and requires at least a **hundred one-hour**
lecture series.
Giving a full-breadth lecture series

on how I experimentally discovered
massively parallel processing
and discovered it **across**
a global network of processors
that is a new internet
is almost like giving a lecture series titled:
"The Complete History of the Universe."
The **high points** of my quest
for how and why
parallel processing makes
modern computers **faster**
and makes the new supercomputer
the **fastest**
include my discovery
of how to solve
the most large-scaled system
of linear equations of algebra
that occurs in petroleum reservoir modeling
of the flow of crude oil and natural gas
within oilfields.
The **high points** of my quest
for what makes
massively parallel processing

supercomputers **fastest** 1070
include my discovery
that I could solve the most abstract
system of coupled, non-linear,
time-dependent, and state-of-the-art
partial differential equations
of modern mathematics,
called Emeagwali's Equations.
The reputation of these equations
earned them a place of honor
in the list of 20 grand challenge problems
that are computational test-beds
for all supercomputers.
In the 1980s,
these grand challenge problems
were impossible
to massively parallel process
across 64 binary thousand
processors.
Each processor
had its own operating system
and memory.
It was also considered impossible

to massively parallel process **across** 1071
64 binary thousand computers
that enshrouded a globe
and encircled it
as a global network of computers
that are identical
and that were equal distances
apart.
Looking back, in 1946
the fastest computer in the world
used only one
scalar processing unit.
In 1988,
the fastest computer in the world
still used only one
vector processing unit.
Shortly after my experimental discovery
on the Fourth of July 1989,
it made the news headlines
that an African supercomputer wizard
in the United States of America
had experimentally discovered
how the most massively

1072

parallel processing supercomputer
ever built
can massively parallel compute
and compute with 65,536
commodity processors
and solve 65,536 computational physics
problems **simultaneously**.
I am that African supercomputer scientist
that was in the news
in 1989.
My discovery opened the door
to the new world of fastest supercomputing
where the large-scale
computational physicist
can massively parallel process
across 10 binary million
commodity processors.
After a 10 binary **million** fold increase
in the speed of computation,
it's not unreasonable to expect another
100 fold increase in speed
that could enable the fastest supercomputer
of tomorrow

to massively parallel process
and to do so **across** one binary **billion**
commodity processors
that each had its own operating system
and memory,
or to do so **across** one binary **billion**
identical computers
on the internet
of maybe
the twenty-second century.
That experimental discovery
of parallel processing
—or parallel computing
many initial-boundary value problems
of calculus and physics
and solving them **at once**,
instead of sequentially solving
one computational problem **at a time**—
is the answer
to the often asked question:
[quote unquote]

"What did Philip Emeagwali
contribute to the development
of the computer?"

26.1.9 Beauty as Driving Force

I was asked:
"Is there beauty
in the mathematics used to
theoretically discover parallel processing?
If so, did that beauty help you
to experimentally discover
parallel processing?"
My set of floating-point arithmetical
operations is detailed
and its beauty is visible only to
the **arithmetician**.
My system of equations of a new algebra
is abstract
and its beauty is visible only to
the **algebraist**.
My system of coupled, non-linear,
time-dependent, and state-of-the-art

partial differential equations
of a new calculus
was used to define
some of the grand challenges
of computational physics
and its beauty is visible only to
the **grand wizards of calculus**.
My geometrical illustrations were described
as beautiful
because the beauty of a geometrical object
is self-evident.
What attracted eye balls
to my experimental discovery
of parallel processing **across**
my new HyperBall supercomputer
was my 1970s geometrical illustrations
of the paths that I parallel processed
my emails through.
Unlike a system of coupled, non-linear,
time-dependent, and state-of-the-art
partial differential equations
of modern mathematics
that's abstract, invisible, and very ugly

1076

my geometrical illustrations
of my new HyperBall
and **new Cosmic Supercomputers**
were concrete, visible, and very beautiful.
It’s beautiful to look at the
bi-directional edges,
to look at the **one-to-one correspondence**
to the as many email wires
that interconnected
64 binary thousand vertices
and did so
with a **one-to-one correspondence**
to 64 binary thousand
commodity processors,
or to 64 binary thousand
identical computers.
I might add that the beauty in mathematics
differs from the beauty
in music or in a novel.
A song and a story
entertain.
However, my mathematical equation,
or my computer algorithm,

is **functional** at the low level
and **intellectual** at the abstract level.
Before my experimental discovery
that occurred on the Fourth of July 1989
the leaders of thought
in the world of supercomputing
ridiculed parallel processing
as a beautiful theory
that lacked **experimental confirmation**.
Parallel processing is beautiful
if and only if
it parallel processed perfectly
and became the driving force
behind all computers
and the fastest supercomputers.