38 Father of the Modern Supercomputer

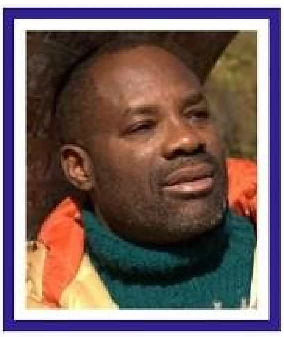

Philip Emeagwali Lecture 180608-2
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

38.1.1 Father of the Modern Supercomputer

The supercomputer
that is not parallel processing,
or computing many things **at once**,
is like a train without its locomotive.
In the early 1980s, I discovered that
it would be theoretically possible
to send emails back and forth
through a new internet
that is a new global network
of commodity-off-the-shelf processors.
That is, commonly-available processors
derived from the same processor,
same memory,
and same board technologies
that were used to power the personal
computers of the 1980s.
It was a paradigm shift because

the new massively parallel processing supercomputer
is several orders of magnitude
more cost-effective
than the prevailing vector processing supercomputer
that was the fastest in the 1980s.

38.2 A Supercomputer Created From Nothing

38.2.1 Contributions of Philip Emeagwali to the Computer

I'm **Philip Emeagwali**.
I'm the first person to ever discover
the fastest computations **across**
the slowest processors.
I invented

how to use 64 binary thousand
computer codes
that executed
floating-point arithmetical operations
that arose while solving
the toughest problems
in modern calculus.
I invented
how to execute 64 binary thousand
excruciatingly-detailed
general circulation models
and how to execute those models
to **foresee** otherwise **unforeseeable**
global climate changes.
I invented
how to execute those models
and execute them
as 64 binary thousand
petroleum reservoir simulators.
I invented

how to **execute** 1732
those reservoir simulators
and execute them to **recover**
otherwise **unrecoverable**
crude oil and natural gas.
I invented
how to execute
computational **aerodynamics** codes
and how to execute them to design
the complex geometric configurations
of aircrafts
that will yield better fuel efficiency.
I invented
how to execute them
to simulate airflows
across aircrafts
that were not confined
to the inside of a **wind tunnel**
instead of simulating airflows
across aircrafts

that were confined
inside a **wind tunnel.**
I invented
how to use them
as 64 binary thousand
computational fluid dynamics codes
and how to use those codes
to **hindcast** the past motions
or **forecast** the future motions
of planetary fluids,
such as the Earth's oceans
and atmosphere
as well as the motions
of water along rivers, lakes,
and groundwater
and to **hindcast** the rotation of fluids
or to **forecast** the rotation of fluids
that arises from planetary rotation
and **stratification.**
I invented

1733

how to investigate the hydrodynamics
and the **magneto-**hydrodynamics
of rotating fluids.
The reason those problems
were called grand challenges
in supercomputing
was that they were considered
too computation-intensive
and, therefore, **impossible** to compute
on existing supercomputers.
I invented
how to make the
impossible-to-solve
that is **impossible**
within a vector processing
supercomputer
and make it **possible-to-solve**
and **possible**
across a new internet
that is a new ensemble of 65,536

tightly-coupled processors
that shared nothing with each other
and that defines a new supercomputer.
I invented
how to make the **impossible-to-compute**
possible-to-compute
and how to execute
the fastest computations
by experimentally reducing
the **time-to-solution**
and reducing that time
from 65,536 days of computing,
or 180 years of computing,
on one processor
[or one computer]
to only one day of the fastest computing,
across a new global network of
65,536 tightly-coupled processors
[or across as many computers].
I invented

that new global network
as a parallel processing machine
that is **not** a supercomputer *per se*
but that is a new internet *de facto*.
In an **echoic retentive** manner
of speaking,
I experimentally discovered
180 years in one day
and I invented
a new supercomputer
that has a better cost-benefit ratio.

38.2.2 Where's Philip Emeagwali?

I'm **Philip Emeagwali**.
My technological quest
was to invent
new ways to increase the speeds
of computers and, ultimately,
to invent the **fastest** supercomputer.

The supercomputer of today
will be the computer of tomorrow.
We create tomorrow
by what we invent today.
Since 1989, school children
doing school reports
on the contributions of **Philip Emeagwali**
to the development of the supercomputer
often ask:
"Where is **Philip Emeagwali**?"
I was in supercomputer centers
across the United States.
I began supercomputing
on June 20, 1974 at age nineteen.
I began supercomputing
in the Computer Center
at 1800 SW Campus Way,
Corvallis, Oregon, United States.
I began supercomputing
when President Richard Nixon

was in the White House. 1738
However, my invention
of how to parallel process
and how to do so **across**
my new internet
that is a new global network of
64 binary thousand
tightly-coupled processors
that were already available
in the market
and how to fast compute
and how to communicate **synchronously**
while computing **simultaneously**
and how to do both **across**
a new massively parallel processing
supercomputer
made the news headlines
in 1989.
I made the news headlines
sixteen years

after **I** **began supercomputing**
in Corvallis, Oregon, United States.

38.2.3 Struggles of Philip Emeagwali

In the 1980s, I used my fingers
to count the number of
part time programmers
of the handful of
massively parallel processing machines
that existed.
In the 1980s,
I was the only fulltime programmer
of the most massively parallel processing
supercomputer in the world
and I was *de facto*
the first supercomputer scientist.
As a lone wolf high-performance
supercomputer programmer,
I learned that it takes time

to make an invention, 1740
such as inventing
the fastest supercomputer.
I failed a hundred times
before I invented
how to harness the slowest 65,536
processors in the world
and how to harness those processors
and how to do so
with a **one-to-one** correspondence
between processors and problems.
I invented
how to harness those tightly-coupled
central processing units, or CPUs,
and how to use them
to simultaneously solve
as many initial-boundary value problems
arising in extreme-scale
computational physics.
I invented
how to solve the toughest problems

in computational mathematics
and how to solve those problems
across processors
that shared nothing with each other.
and solve them faster than
any supercomputer
that only vector processed its
floating-point arithmetical computations
or **across** processors
that shared everything between each other.

38.2.4 Quest of Philip Emeagwali for the Fastest Supercomputer

My supercomputing quest
for the **invention**
of the fastest parallel processing
supercomputer
began on June 20, 1974
in Corvallis, Oregon, United States.

My quest for the fastest supercomputer 1742
began like the dispersal
of the small seed
of the 160-foot-tall **Iroko** tree.
The **Iroko** tree
is venerated in my ancestral Igbo Land,
of the southeastern region of Nigeria
of tropical West Africa.
The small seeds of the **Iroko** tree
are dispersed by bats.
The bat unwittingly chooses
what spot to grow the **Iroko** tree
and choose the destiny
of the **Iroko** tree.
Like the **Iroko** tree,
my **invention**
of the massively parallel processing
supercomputer
seemed to choose where it will **reveal**
itself.
That invention

gradually **revealed itself**
across sixteen years,
and **revealed itself**
across sixteen supercomputers,
and **revealed itself**
across sixteen cities.
My quest for the parallel processing
supercomputer
began as a vague idea.
My quest for the high-performance
supercomputer
began as the small seed
of an **Iroko** tree.

38.2.5 A Black Supercomputer Geek in Oregon

My quest for the parallel processing
supercomputer
began on a sequential processing
supercomputer

in Corvallis, Oregon, United States. 1744
In my first year of programming
sequential processing supercomputers,
I accessed those supercomputers
via a Teletype Model 33 ASR.
The acronym ASR
stands for **Automatic Send-Receive**.
On June 20, 1974,
the Teletype Model 33 ASR
that I was using
could be in my one-room studio
apartment
at 195A Knox Street South,
Monmouth, Oregon, United States.
But the Teletype Model 33 ASR
was inside a computer laboratory.
That computer laboratory
was at three hundred and forty-five
[345] Monmouth Avenue North,
in Monmouth, Oregon, United States.
That computer laboratory

was twenty-two miles away 1745
from the two supercomputers
—both in the Computer Center,
in Corvallis, Oregon, United States—
that I was remotely programming.
The Teletype Model 33 ASR
automatically sent and received
my supercomputer codes and data.
My electronic messages
were sent and received
via a telephone line.
I used that telephone line
to electronically speak
with a Programmed Data Processor
that was called the PDP-8 minicomputer.
Nine months later, in late March 1975,
I relocated from Monmouth (Oregon)
to Corvallis (Oregon)
and I continued programming
the supercomputer
—from 200 feet away—

and from a Teletype Model 33 ASR
that was in Kidder Hall
at 2000 SW Campus Way,
Corvallis, Oregon, United States.

38.2.6 A Black Supercomputer Geek in Nuclear Labs

My **invention**
of the massively parallel processing
supercomputer
revealed itself
within a new internet
that is a parallel processing machine
that resided in **Los Alamos**, New Mexico,
United States.
That parallel processing machine
is the **Iroko** tree
of my world of supercomputing.
That parallel processing machine
that I experimentally discovered
to be a new internet

powered by
a new global network of
65,536
identical processors
that were already available in the market
was the **precursor**
to the fastest, modern supercomputer
of today
that is powered by
a network ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
identical processors
that were already available in the market
anyway.
That parallel processing machine
that I invented
is the **precursor**
to the modern computer
of today
that is powered by

1748

about one hundred
processors
that were already available
in the market anyway.
That parallel processing machine
was the **precursor**
to the high-performance supercomputer
of today
and the **precursor**
to the modern computer
of today.
To invent
is to create **something**
out of **nothing.**
For sixteen years, onward of June 20,
1974,
I—**Philip Emeagwali**—
was the lone wolf
supercomputer programmer
that was **perched**
on the top of that **Iroko** tree

of the forest of fastest supercomputers
that are powered by
billions upon billions
of processors that were already available
in the market anyway.

38.3 How I Changed the Way We Look at Computers

38.3.1 Changing the Way We Look at Computers

In 1946, the programmable
supercomputer
was invented to compute **automatically**
to compute **faster**.
For me, **Philip Emeagwali**,
my technological quest
was for the fastest supercomputer
that I could program to make the
impossible-to-solve problems arising

in extreme-scaled computational physics
possible-to-solve
and possible **across**
a new internet
that is a new global network of
65,536 tightly-coupled processors
with each processor
operating its own operating system
and with each processor
having its own dedicated memory
that shared nothing with each other.
That necessity
to solve the most extreme-scaled
problems arising in computational
physics
took me to the **crossroads**
of mathematics, physics,
supercomputer and internet sciences.
During the sixteen years
onward of June 20, 1974,

I sojourned from a system of equations
of algebra
that I solved on a sequential processing
supercomputer
to the **crossroad**
between the laws of physics
and the partial differential equations
of calculus.
I continued my sojourn
to the **crossroad**
between calculus
and the system of equations
of algebra.
I continued my sojourn
to the **crossroad**
between algebra
and a set of floating-point operations
of arithmetic.
I continued my sojourn
to the **crossroad**

between arithmetic 1752
and my new global network of
processors
that is a new internet.
I experimentally discovered
that the processors
are to the modern supercomputer
what the **elements**
are to the **periodic table**.
I experimentally discovered
that the processors
are the basic building blocks
of the modern supercomputer
just as **atoms**
are the basic building blocks
of **molecules**.
I experimentally discovered
that computational physics
is at the foundation
of the fastest supercomputer

just as **axioms**
are at the foundation
of **axiomatic mathematics**.
I experimentally discovered
that the regular, short, and equidistant
email wires
that I visualized as fiber optic wires
of the **Philip Emeagwali** internet
are its basic building blocks
just as the hypercube
is constructed from the cube
and is constructed
from the regular and equidistant
edges of the hypercube
from lower dimensions.
I visualized
the new **Philip Emeagwali** internet
as a **small copy**
of the planetary-sized Internet
and as a new global network of

64 binary thousand
tightly-coupled, already-available
processors
that is *de facto*
one seamless, cohesive unit
that is the fastest supercomputer
that computed in parallel,
or computed by solving
64 binary thousand
initial-boundary value problems
at once.
My invention
of the massively parallel processing
supercomputer
was independent of processor
technology.
That invention was a blueprint
for a new internet
that, in turn, is a prototype
of the **Philip Emeagwali**

Cosmic Supercomputer.
I invented that Cosmic Supercomputer
in the 1970s and '80s
and I described
my Cosmic Supercomputer
over a series of lectures
that I posted at **emeagwali dot com**.

38.3.2 Changing the Way We Look at Supercomputers

Since 1946,
the year the programmable
and sequential processing
supercomputer
was invented,
the number one complaint
against the supercomputer—that is

the world’s fastest computer—was that
the supercomputer was not fast enough
and could not be used
to accurately forecast the weather.
The supercomputer
is not fast enough to simulate
the long term changes
in the Earth’s atmosphere,
that is also known
as global climate change.
For the four decades onward of 1946,
the sequential processing
or vector processing supercomputer
powered by
only one **isolated** processor
—that was not a member
of an ensemble of processors—
inspired the development of
weather forecasts.
Each weather forecast
is computed from a

supercomputer model
that is **rooted**
in a companion mathematical model
that is **rooted**
on the laws of physics.
At the foundation
of the supercomputer weather forecast
is a large system of
partial **difference** equations
of algebra
that was derived from
finite difference **discretizations**
and approximations
of the primitive equations
of meteorology.
The primitive equations
are a system of coupled,
non-linear, time-dependent,
and state-of-the-art
partial **differential** equations
that is the toughest problem in calculus.

For those four decades, 1758
the weather forecast
was generated within a supercomputer
that used only one
isolated processor
that was not a member
of an ensemble of processors.
One of the reasons
the weather forecast
of today
is more accurate
than the weather forecast
of the first four decades,
onward of 1946,
is that today's weather forecast
is **simultaneously** computed
in parallel,
and **synchronously** communicated
across up to ten million
tightly-coupled already-available
processors

that shared nothing with each other.
My **invention**
of the massively parallel processing
supercomputer
was independent of processor
technology.
That **invention** was a blueprint
for a new internet
that I envisioned
as being a subset of the Internet.
I visualized that new internet
as being at the core
of an emergent Cosmic Brain.

38.3.3 The Modern Supercomputer

I **invented**
how to forecast the weather
and how to do so with greater accuracy

and how to do so **across** 1760
my new internet
that is a global network of
65,536 tightly-coupled processors.
In my invention, the weather is forecast
by executing the set of floating-point
arithmetical operations
that arises
from solving the primitive equations
of meteorology.
The primitive equations
is a system of coupled, non-linear,
time-dependent, and state-of-the-art
partial differential equations
of modern calculus.
I discretized and replaced
the primitive equations
with my system of equations
of extreme-scale algebra
that, in turn, approximated
those partial differential equations.

I invented how to solve 1761
the arising system of equations
of algebra
that arose from the primitive equations
of meteorology and calculus
and how to solve
that extreme-scaled algebraic problem
and solve it **across**
my new internet
that is a new global network of
64 binary thousand tightly-coupled,
processors.
Those already-available processors
were identical
and were equal distances **apart**.
That new global network of processors
was a new internet.
I invented that new internet
to be at the **granite core**
of the high-performance supercomputer
that I used to compute in parallel.

In contrast
to the new massively parallel processing
supercomputer,
the old **sequential** processing
or the old **vector** processing
supercomputer
computed **sequentially**,
or solved only one problem **at a time**.
I **invented**
how to more accurately forecast
the weather
and how to forecast it **across**
the globe
and **across** a new global network of
tightly-coupled processors
that is a new internet.
I programmed that new internet
as a new global network of
two-raised-to-power-sixteen processors.
I visualized that new internet
as a global network of

millions upon millions 1763
and billions upon billions
of identical processors
that had better cost-benefit ratio.
The new supercomputer
that I invented
is less expensive because
its processors were mass produced
and were already available
in the market.
The reason my **invention**
of the massively parallel processing
supercomputer
was reported in the news media
and reported in the June 20, 1990 issue
of *The Wall Street Journal*
was that it was then
a technological breakthrough
to harness a new internet
that is a new global network of
64 binary thousand

identical processors. 1764
It was a technological breakthrough
to **invent** how an ensemble
of the slowest processors
in the world
could be harnessed
and used to compute **faster**
than any vector processing
supercomputer
that computed
with the **fastest** processor.
Back in 1989, the community of 25,000
vector processing
supercomputer scientists in the world
believed in **Seymour Cray**—the
legendary pioneer of vector processing—
who scorned and ridiculed
the massively parallel processing
supercomputer
and dismissed
the unorthodox technology

as a huge waste of everybody's time.
My **invention**
of the massively parallel processing
supercomputer
that is a new internet
enabled me to compress
180 years of **time-to-solution**
on only one processor
and to compress that time
to just one day of **time-to-solution**
across my new internet
that is my global network of
64 binary thousand
processors
that computed in parallel.
My technological breakthrough
opened the path
to the massively parallel processing
supercomputer.
That new path was taken by
Chinese supercomputer scientists

and used to compress
their 30,000 years, or ten million
six hundred and forty-nine thousand
six hundred [10,649,600] **days,**
of their **time-to-solution**
on one processor
to just one day of **time-to-solution**
across ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
processors
that computed in parallel.

38.3.4 The Chinese Are Coming

Three decades later,
China is exploiting the rich
and the fertile consequences
that arose from
my **invention**
of the massively parallel processing

supercomputer
that occurred
on the Fourth of July 1989
in Los Alamos, New Mexico,
United States.
For the past three decades,
the world's fastest supercomputer
was massively parallel processing
and do so either in Japan
or in the United States.
China now has the world's fastest
supercomputer
that is massively parallel processing
across
10.65 million already-available, tightly-coupled processors.
Chinese high-performance
supercomputer scientists
divides computation-intensive problems
in extreme-scale computational physics
to enable them to solve them, in parallel,

and as **10.65 million** problems.
China's use of **10.65 million**
processors
to power its world's fastest
supercomputers
vindicated my use of 64 binary thousand
identical processors
to record the world's fastest computation
that I recorded
on the Fourth of July 1989.

38.3.5 The Supercomputer Shapes the Future of Mankind

I invented
the massively parallel processing
supercomputer
and I discovered the technology
as the bridge

that closed the distance
between the first supercomputer of 1946
and the modern supercomputer of today.
Parallel processing is the bridge
that closed the distance
between yesterday's
sequential processing supercomputer
and today's high-performance
supercomputer.
The massively parallel processing
supercomputer
is used to reduce **times-to-solution**
and used to increase **productivity**
and used to radically reduce
times-to-market.
The massively parallel processing
supercomputer
is used to tackle the toughest problems
arising in computational physics
arising in modern calculus,
and arising in extreme-scale algebra.

I foresee our children's children 1770
massively parallel processing **across**
their internet
that will be their global network of
one trillion processors
that enshroud the Earth.
I foresee the internet
of our children's children
as their planetary supercomputer
that solves our as-yet-unsolved
grand challenge problems.