# 29 How I Invented a New Supercomputer

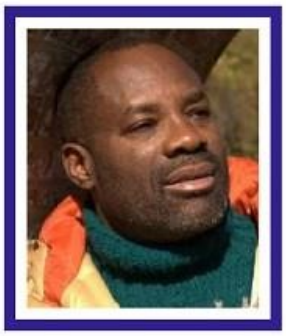

Philip Emeagwali
info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture 180913-2
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

## 29.1.1 Who's Philip Emeagwali, the Discoverer of Parallel Processing?

I invented
a new supercomputer

that could be used
to solve the toughest problems
arising in mathematics and physics.
Such problems are called
the Grand Challenge Problems
of supercomputing.
Back in 1989
the year I completed my invention
and began appearing
in major U.S. newspapers,
**Seymour Cray**
—who was the then leading mind
in the world of vector processing
supercomputers—could not understand
the mathematics of where each of my
64 binary thousand
initial-boundary value problems
should be
and when it should be
at each of my as many processors.
Back in 1989,

the supercomputer scientists
that attempted to parallel process **across**
an ensemble of processors
were processing
by the seat of their pants.
That is, those supercomputer scientists
of the 1970s and ‘80s
did not understand
the complicated mathematics
and did not have the command
of scientific materials
and the subject matter knowledge
that was needed
to solve the Grand Challenge Problem.
The Grand Challenge Problem
was at the crossroad
where calculus, algebra, physics,
computing, and supercomputing
met each other.
Because they did not understand
the Grand Challenge Problem,

the 25,000 vector processing
supercomputer scientists of the 1980s
and earlier
began to hate
the massively parallel processing
supercomputer
that I used
to solve the Grand Challenge Problem.
Back in the 1980s, I was the only
internet scientist
that knew the 64 binary thousand,
or the 65,536, "address tags"
that directed where on my internet
each of my 65,536
initial-boundary value problems
of calculus and physics
was delivered.
To the 25,000 vector processing
supercomputer scientists of the 1980s,
reading my 65,536
email message-passing codes

was as incomprehensible
as reading a Chinese newspaper.
My invention
of how to provide the address tags
was the necessary pre-condition
to the invention
of the massively parallel processing
supercomputer
that is also a new internet *de facto*.
That invention
has rich and fertile consequences
and contributed
to the more complete understanding
of how and why
the technology called
parallel processing—or solving
millions upon millions of problems—
**across** as many processors
and solving them at the **same time**
makes the computer faster
and makes the supercomputer super.

## 29.1.2 Birth of a New Computer

My quest for the precursor
to the modern supercomputer
that is fastest by parallel processing **across**
a new internet
that is a new global network of
processors
began as a vague idea.
That quest began
as the seed of an Iroko tree
and blossomed, sixteen years later,
into the world's fastest supercomputer
that is the Iroko tree
of the unknown forest
named the massively parallel processing
supercomputer.
The Iroko tree
is the tallest tree in Igbo Land
of southeastern Nigeria.
The Iroko tree
grows along the west coast of Africa.

The Iroko tree
can live for up to 500 years.
My quest for the parallel processing supercomputer
began in the early morning of
Thursday June 20, 1974
in Corvallis, Oregon, United States.
My quest for the modern supercomputer
ended at 8:15 in the morning
of Tuesday the Fourth of July 1989
in Los Alamos, New Mexico, United States.
That quest for the parallel processing supercomputer
led to my **deeper** and **surer**
understanding of the Internet
as a planetary supercomputer-hopeful.
I had ideas about parallel processing
since 1974.
But, until the Fourth of July 1989,
I did not experimentally prove
that parallel processing
makes the impossible-to-compute
possible-to-compute.

My experimental discovery
of parallel processing
that occurred **across** my new internet
that is a new global network of
65,536 processors
occurred on the Fourth of July 1989.
That invention
of the parallel processing supercomputer
was my lockdown evidence
and it was the first
experimental confirmation
of the beginning
of new era in the world of supercomputing.
That experimental discovery
was processor agnostic
and was node agnostic.
My discovery
of the parallel processing supercomputer
made the news headlines
because I successfully tested
the new supercomputer
and that I experimentally confirmed it
as the world's fastest computer.

My invention
of the parallel processing supercomputer
went beyond theory
to become an experiment-verified invention
of a new supercomputer
that is not a computer *per se*
but that is a new internet *de facto*.

## 29.1.3 Steve Jobs Questioned Parallel Processing

I—**Philip Emeagwali**—
was the African supercomputer scientist
that was in the news headlines
onward of my invention
of the parallel processing supercomputer
that occurred on the Fourth of July 1989.
I was in the news for inventing
how and why
a new ensemble of the slowest processors
can be harnessed
and be used to solve the toughest problems
arising in both computational mathematics
and computational physics

and how that new ensemble
of the slowest processors
can be harnessed and be used
to solve computation-intensive problems
that even the fastest
vector processing supercomputer
was unable to solve.
Because I was in the news,
**Steve Jobs**
tried to reach me by telephone
in June 1990.
**Steve Jobs**
and his computer programmers
explored how they could increase
the speed of their computers
and increase them by incorporating
the technology of the massively
parallel processing supercomputer
that I invented.
Put differently, **Steve Jobs**
wanted to know
how I reduced 65,536 days,
or 180 years, of **time-to-solution**

on only one processor
and how I reduced that **time-to-solution**
to only one day of **time-to-solution**
**across** a new internet.
That new internet
that I invented
is a new global network of
65,536 tightly-coupled processors
that emulates one seamless,
cohesive computer
that is a new supercomputer, *de facto*.
It was for a good reason
that to massively parallel process
and do so **across**
a new ensemble of 65,536 processors
was called the grand challenge problem
of supercomputing.
It was for a good reason
that the June 14, 1976 issue
of the *Computer World* magazine
carried an article
that ridiculed, mocked, and rejected
the massively parallel processing

supercomputer.
That article in the *Computer World* magazine
was titled:
"Research in Parallel Processing
Questioned as 'Waste of Time'."
Thirty-two [32] years after that negative
article in the *Computer World* magazine,
**Steve Jobs**
also questioned research
in parallel processing, describing it
as a huge waste of everybody's time.
On June 9, 2008,
**Steve Jobs**
told the opening session
of Apple's Worldwide Developers Conference
in San Francisco, California
that his research computer scientists
at his Apple Corporation
questioned research in parallel processing
as a huge waste of their time.
As reported in the June 10, 2008 issue
of the *New York Times*,
**Steve Jobs**

told Apple's Worldwide Developers
that:

[quote]
[PROP ALERT]
"The way the processor industry
is going
is to add more and more cores,
but nobody knows
how to program
those things,"
**Steve Jobs** said.
And **Steve Jobs** continued:
"I mean, two, yeah;
four, not really;
eight, forget it."
[unquote]

I invented
the massively parallel processing
supercomputer
that **Steve Jobs**
and his research computer scientists

declared impossible to invent.
The massively parallel processing technology
is at the heart of fastest supercomputer
of today
that, hopefully, will become the
everyday computer of tomorrow.
I experimentally discovered
that the technology
of massively parallel processing
is a necessary condition
for creating
both the fastest computers
and the fastest supercomputers.
Since the programmable supercomputer
was invented in 1946, we never invented
a new supercomputer
without, first and foremost,
experimentally discovering
faster supercomputer speeds.
By the definition of the word "computer,"
we will never invent
a new supercomputer
without, first and foremost,

experimentally discovering
a new fastest supercomputer speed.
To achieve that grand wizardry
in fastest massively parallel processing
supercomputing
requires the **visceral** understanding
that the massively parallel processing
supercomputer
is not a computer, *per se*.
I experimentally discovered
that my massively parallel processing
supercomputer
that I visualized
as a small global network of
65,536 commodity processors
that were identical
and that were equal distances
**apart**
is a small internet, *de facto*.

## 29.1.4 An Invention is a Black Box in a Dark Room

It's been said that:

"Out of the heart, the mouth speaks."
I have spoken out of my heart.
I have spoken a lot about supercomputers.
I have spoken about the massively
parallel processing supercomputer
that I invented.
I have spoken since the 1970s
and spoken about how I was
the first eye witness
to the experimental discovery
of how and why
a new ensemble of the slowest processors
that computes together
as one seamless, cohesive
massively parallel processing supercomputer
is a new internet, *de facto*.
The words I spoke
came from my discoveries,
heart, and brain.
Those words will define me for posterity.
To witness a discovery
that has some rich, fertile,
and far-reaching consequences

is like walking into a forest
and witness a lot of leaves
fall on your head.
In a nutshell,
this is my supercomputer story,
and my contribution
to the development of the fastest
supercomputer.
I began programming supercomputers
on June 20, 1974
in Corvallis, Oregon, United States.
Two years later
and in nineteen seventy-six [1976]
and at the National Computer Conference
in New York City,
a panel of supercomputer experts
warned that
parallel processing machines
are [quote unquote] "large and clumsy."
That panel of supercomputer experts
ridiculed, mocked, and rejected
the then unproven technology
of the massively parallel processing

supercomputer.
That panel of supercomputer experts
offhandedly dismissed
the technology of parallel processing
as a huge waste of everybody's time.
In the nineteen eighties [1980s],
I was dismissed
from my research teams
and dismissed
for advocating the massively parallel
processing supercomputer.
In the nineteen eighties [1980s],
I could count the number of programmers
that were massively parallel processing
and count them on my fingers.
But I needed only one finger
to count them because I was the only person
that was logged on
twenty-four seven [24/7]
and logged on the most massively
parallel processing supercomputer
ever manufactured.

Today, the market value
of that massively parallel processing
supercomputer
that I programmed alone
has appreciated by a factor of
one thousand.
That new supercomputer
that I programmed alone
now costs
the budget of a small nation.
I experimentally programmed
the new massively parallel processing
supercomputer
and I programmed it alone
and I did so because
the 25,000 vector processing
supercomputer scientists
of the 1980s
were heeding the textbook warnings that
it will forever remain
experimentally impossible
to harness the total theorized
supercomputing power

of an ensemble of thousands of processors.
The supercomputing community's
abandonment
of the massively parallel processing
supercomputer
unintentionally created
a certain cachet of exclusivity
and did so because
they were *de facto* excluded
from the frontier of human knowledge
that was defined
by the massively parallel processing
supercomputer
of the 1980s and earlier.
The supercomputer textbooks,
of the 1970s and '80s,
cited **Amdahl's Law**
that was erroneously formulated
in April 1967
and cited that incorrect law to argue that
massively parallel processing
**across** 64 binary thousand
processors

will forever remain impossible,
or at least remain impractical.

## 29.1.5 Wizardry is Making the Impossible Possible

I'm **Philip Emeagwali**.

I began supercomputing
on June 20, 1974 at age 19
in Corvallis, Oregon, United States.
Back in the 1970s and '80s,
the core technology
that drives both the modern computer
and the massively parallel processing
supercomputer of today
and drives them at their fastest speeds
was ridiculed, mocked, and rejected.
In scientific research,
a scientist that discovered something new
that other scientists did not know
asserts his or her authority
over his or her discovery
and do so by authoring

an in-depth research report
on his or her discovery
and then giving a companion
series of lectures that are shared
with everybody in the world,
including those on YouTube.
My series of lectures on my contributions
to the development
of the modern supercomputer
are posted at YouTube dot com
slash [/] **emeagwali**.
The vector processing
supercomputer scientists of the 1980s
that I asked to review my
1,057-page research report
on the new massively parallel processing
supercomputer
that I was programming
ridiculed, mocked, and rejected
my experimental discoveries and inventions.
In late 1989,
when my experimental discovery
won the top prize

in the field of supercomputing
and began to make the news headlines
across major U.S. newspapers
those supercomputer scientists
that ridiculed, mocked, and rejected
my research report
confessed to news journalists
that they threw my 1,057-page
supercomputer research report
into the trash.
After my experimental discovery
of the massively parallel processing
supercomputer
that occurred
on the Fourth of July 1989,
those supercomputer scientists
that ridiculed, mocked, and rejected
my discovery of the precursor
of the modern supercomputer
realized that the history
of the supercomputer
is unfolding right before their eyes
and that they were not

a part of the story
of the development of the modern
supercomputer.
In 1989
and after I won the top prize
in supercomputing,
those supercomputer scientists
that ridiculed, mocked, and rejected
my experimental discovery
of parallel processing
became envious of my new fame
and ran back to me
to beg for a second copy
of my 1,057-page
supercomputer research report
that they previously threw into the trash.
The news headlines
described me as the
“African Supercomputer Wizard”
and did so because
my experimental discovery
of the massively parallel processing
supercomputer

that occurred **across** a new internet
that is a new global network of
65,536 tightly-coupled, commonly available
processors
and that occurred
on the Fourth of July 1989,
was described as the wizardry act
of making the previously
impossible-to-compute possible-to-compute.
What was erroneously written
in the supercomputer textbooks
of the 1940s through ‘80s,
was that it will forever remain impossible
to harness the potential power
of parallel processing.
The June 14, 1976 issue
of the *Computer World* magazine
interviewed supercomputer experts
that were attending
the 1976 National Computer Conference
in New York City.
Those supercomputer experts were asked
if it will ever be possible

to invent
how to use parallel processing,
or how to use an ensemble
of thousands of processors,
and how to use them to execute
the fastest computations **across**
that ensemble.
The unanimous opinion
of those supercomputer scientists
was summed up in an article
in the June 14, 1976 issue
of the *Computer World* magazine.
**E. Drake Lundell Jr**,
who was the computer industry editor
for the *Computer World* magazine,
wrote an article titled:

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

Looking back retrospectively,

the reason I was **not** discouraged
from massively parallel processing
**across** an ensemble of 65,536 processors
was that I was then only twenty-one [21]
years old and, therefore,
I had the **years to waste.**
It seemed like I was wasting fifteen years,
onward of age nineteen,
and wasting that prime decade and half
pursuing parallel processing that then
seemed impossible.
I spent the thirteen years,
onward of 1976,
conducting my supercomputer research
and researching as a lone wolf
at supercomputer centers that were
**across** the United States.
In the 1970s and ‘80s,
I had a presence in U.S.
supercomputer centers.
I programmed supercomputers
in Corvallis, Oregon.
I programmed supercomputers

in Pasadena, California.
I programmed supercomputers
in Chicago, Illinois.
I programmed supercomputers
for sixteen years
and I was the first person
to be referred to
as a supercomputer scientist.
As an extreme-scale computational physicist
at the frontier of the supercomputer,
my research focus
was on how to invent
a massively parallel processing
supercomputer
that I defined and outlined
by a new ensemble of 65,536
tightly-coupled
commodity-off-the-shelf processors
and how to invent
how to reduce 180 computing-years
on only one processor
that is solving the toughest problem
in extreme-scale computational physics

to just one supercomputing-day
**across** my new internet
that is a new global network of
65,536 tightly-coupled processors
that are harnessed
to solve the same
initial-boundary value problem
in modern calculus
and/or mathematical physics.

## 29.1.6 Eleven Out of Ten

Eleven out of ten people
did not understand
how I invented
the precursor to the modern supercomputer.
It was not I—**Philip Emeagwali**—
that first reported
my invention
of the massively parallel processing
supercomputer
that computes faster than
any vector processing supercomputer.

My experimental discovery
of how to parallel process
an initial-boundary value problem
of calculus and physics
—called extreme-scale
petroleum reservoir simulation—
occurred on the Fourth of July 1989.
That experimental discovery
of massively parallel processing
was first reported by
The Computer Society of the IEEE.
The IEEE is the acronym
for the Institute
of Electrical and Electronics Engineers.
The contributions
to computational mathematics
that arose from my experimental discovery
of the parallel processing supercomputer
was reported in the May 1990 issue
of the *SIAM News*.
The *SIAM News*
is written by research mathematicians
for research mathematicians.

The *SIAM News*
is the flagship publication of SIAM,
the acronym for the Society
for Industrial and Applied Mathematics.
However, my new mathematical knowledge
that the *SIAM News* described
as my mathematical contributions
to computational mathematics
was how the *SIAM News* understood
my mathematical discovery
of nine new partial **differential** equations
of modern calculus
and understood my mathematical invention
of the companion
nine new partial **difference** equations
of extreme-scale algebra
that approximated my nine new
partial **differential** equations.
What the *SIAM News* understood
as my contributions
to computational mathematics
was not how I understood my invention.
A discovery is like the moon

that has two parts:
the visible part
and the hidden part.
In 1989,
the news media was reporting
the concrete and the visible part
of my experimental discovery
of the massively parallel processing
supercomputer
and was ignoring
the abstract and the invisible part
of that discovery
that occurred on the Fourth of July 1989.

## 29.1.7 Diary of a Black Cowboy Physicist

Back in 1977,
I worked in the civil engineering field
of highway construction.
I did so for the State of Maryland
and I helped expand portions
of the federal highway

between Baltimore (Maryland)
and Washington, D.C.
Nine years later, I worked
for the United States government.
I worked as an engineering physicist
that was responsible for updating
nine Standard Operating Procedures.
I used each Standard Operating Procedure
to safely operate each dam
and I used the procedure to safely control
the water level of the reservoir that is
upstream of the dam
and I used the procedure to operate
the hydroelectric powerplant
within the dam.
All the nine dams that I worked on
were located within the state of Wyoming.
Wyoming
is a state in the western region
of the United States.
Wyoming
is defined by vast plains
and by the Rocky Mountains.

Wyoming
is best known for epic Cowboy
and western movies
and for its Yellowstone National Park.
One of those nine dams
that I operated
was the imposing
214-feet high Pathfinder Dam.
**Pathfinder** was an arch dam
that was a triumph
of early 20th century design.
Water cascaded from
the mid-level outlet tunnel
of the Pathfinder Dam.
The Pathfinder Dam
on the North Platte river of arid Wyoming
is listed on the National Register
of Historic Places.
The North Platte
is a river that rises in the snowmelt
of the Colorado Rocky Mountains.
The nine reservoirs in the state of Wyoming
that I operated

stored water that served
the parched farm lands of Wyoming.
Glendo Dam in the state of Wyoming
is an earth fill dam
that is 190 feet high.
Glendo Dam has a crest length
of 2,096 feet.
As a practicing engineering physicist
at the Glendo Dam,
I visualized the total amount of water
that I controlled as follows:
I had 800,000 acre feet of water
stored upstream of Glendo Dam.
One acre foot of water
covers an acre of land
to a depth of one foot.
The term "SOP"
is the United States Bureau of Reclamation's
acronym
for its Standard Operating Procedure.
In my SOP for Glendo Dam,
the water storage capacity
and the water redistribution

is divided as follows:
100,000 acre feet
for farmland irrigation;
115,000 acre feet
to control sediment deposition;
275,000 acre feet
for controlling floods
and avoiding dam break;
and 310,000 acre feet
for producing hydroelectric power.
Those engineering experiences aside,
I was primarily a research engineer,
not a practicing engineer.
The difference is this:
the practicing engineer
likes to solve problems
that are considered solveable
while the research engineer attempts to
solve the toughest problems
that were considered unsolveable.

## 29.1.8 Solving the Toughest Problem in Physics

I was in the news headlines
because I experimentally discovered
that the impossible-to-compute
is, in fact, possible-to-compute.
A scientist becomes famous
when he or she
creates new knowledge,
or makes a discovery or an invention,
that creates new wealth
and that makes the world a better place.
For me, **Philip Emeagwali**,
I defined solving the toughest problems
arising at the frontier
of engineering knowledge
as the science and technology
of using my knowledge of
extreme-scale algebra, abstract calculus,
computational physics,
and the massively parallel processing
supercomputer
and using that knowledge
in a never-before-seen way
and using that knowledge

to solve extreme-scale
initial-boundary value
mathematical problems
and using that knowledge
that help build better bridges **across**
my ancestral hometown of Onitsha (Nigeria)
or to help design faster airplanes,
safer ships, safer nuclear powerplants,
and even more fuel efficient cars.
The now ubiquitous technology
of the massively parallel processing
supercomputer
that was scorned and rejected
in the 1940s through ‘80s
is used by practicing engineers
and used to increase their productivity
and used to reduce their **time-to-market**.
A teacher asked her students:

“What is **Philip Emeagwali** famous for?”

I answered:

“The riddle of parallel processing
was experimentally solved
by **Philip Emeagwali**
in Los Alamos, New Mexico, United States,
and experimentally solved
at 8:15 in the morning
of Tuesday the Fourth of July 1989
that was the U.S. Independence Day.”
That invention
was my rock-solid proof—and not a hunch—
that parallel processing
makes computers faster.
Doing many things **at once**,
or supercomputing in parallel,
became a sure-fire way
to increase the speed
of all modern supercomputers.
My invention
of the parallel processing supercomputer
was my first major breakthrough
in the world of science and technology.
That invention was the milestone
that put the name **Philip Emeagwali**

into school reports.
I experimentally discovered
all the parallel processed speedup
they was to discover.
It is often said that
the invention
of the massively parallel processing
supercomputer
is the single most transformative technology
and the biggest advance in physics
since Newton, Galileo.

## 29.1.9 To Invent Demands Genius, Not Luck

To invent a new computer
or a never-before-seen internet
demands genius, not luck or serendipity.
I had no serendipitous invention
of how to massively parallel process
and how to compute **across**
a new internet

that is a new global network of
millions upon millions
of tightly-coupled processors
that are equal distances apart
from each other.
**I visualized**
my small copy of the Internet correctly.
**I visualized**
that never-before-seen internet
*a priori*.
My **theoretical visualization**
enabled me to experimentally discover
that the shape of the cube
in the sixteenth dimensional hyperspace
will give my new internet
regular **form** and **freedom**.
Each processor
within my small copy of my new Internet
communicated via emails
and along sixteen
mutually orthogonal directions

and along as many dimensions
of my imagined
sixteen dimensional universe.
**I visualized** those sixteen directions
as mutually perpendicular
and embedded
within my imaginary
sixteen-dimensional universe.
Each processor
within my small copy of the Internet
communicates in sixteen directions
and communicated
by sending and receiving emails
to and from
its sixteen nearest-neighboring
processors.
My **epiphany** was my discovery
that to execute the fastest computation,
the deepest source of the
massively parallel processing
supercomputer's computing power

was not in its
two-raised-to-power sixteen,
or 64 binary thousand,
tightly-coupled processors
that actually computed.
I theoretically and experimentally
discovered that
the deepest source of the power
of the massively parallel processing
supercomputer
was in its sixteen times
two-raised-to-power sixteen,
or one binary million,
short and regular email wires
that did not compute.
It took me sixteen years
of massively parallel processing
and supercomputing **across**
a new global network of processors
to fully understand
that my new frontier

of supercomputer knowledge
was a new internet *de facto*,
not a new computer *per se*.

## 29.1.10 Sixteen Years to Overnight Success

I'm **Philip** **Emeagwali**.

I'm not an overnight success
in supercomputing.
I began programming
scalar processing supercomputers
and began on Thursday June 20, 1974
and began at 1800 SW Campus Way,
Corvallis, Oregon, United States
and began at age nineteen.
Three weeks after I began
supercomputing,
I was on the front-page
of a local newspaper

that was on the newsstands
of cities of Monmouth (Oregon)
and Independence (Oregon).
During my subsequent sixteen years,
I was ridiculed and abandoned
in my supercomputer research.
My quest was to invent
a new massively parallel processing
supercomputer
that is a new internet.
During the sixteen years onward of
June 20, 1974,
I was shunned for challenging
the dominant sequential processing
supercomputer paradigm.
It was on the sixteenth anniversary
of my first foray
into the world of the supercomputer
—or on June 20, 1990—
that *The Wall Street Journal*
reported that I—**Philip Emeagwali**—

has figured out
how to make the impossible-to-compute
possible-to-compute, namely,
I figured out
how to massively parallel process
and figured out how to compute **across**
a new ensemble of processors
that shared nothing with each other.
That invention
of a new supercomputer
made the news headlines because
it was transformative and translational
and **opened the door**
to a new era in computer science.
For the forty-three [43] years
onward of 1946,
the massively parallel processing
supercomputer
was ignored and misunderstood
by the computer science community.
After my invention

of the Fourth of July 1989,
the enabling technology
of parallel processing
is embodied into every modern computer
and is universally used
to increase the speed
of the fastest supercomputers
in the world.
I figured out
how to synchronously communicate
**across**
a new global network of
1,048,576 email wires
and I figured out
how to simultaneously compute
on 65,536 processors.
But it took me sixteen years
to become the overnight success
that I was portrayed
in major U.S. newspapers,
such as the June 20, 1990 issue

of *The Wall Street Journal*.
Back in the 1980s,
there was, arguably, only one
massively parallel processing
supercomputer in the world.
That massively parallel processing
supercomputer
was powered by
64 binary thousand tightly-coupled
processors
that were identical
and that shared nothing
between each other.
That massively parallel processing
ensemble
of 64 binary thousand processors
only allowed
one supercomputer programmer
to lock all its 65,536
tightly-coupled processors
and lock those two-raised-to-power

sixteen processors simultaneously.
So, I was logged onto
my massively parallel processing
supercomputer
and logged on sixteen hours a day.
I was even running supercomputer codes
while I slept.
I have been married
since August 15, **1981**.
In my marriage, the supercomputer
was the other woman.
In the 1980s, it was impossible
to have a second
modern supercomputer scientist
that was massively parallel processing
**across** an ensemble
of 65,536 processors.
I knew that I—**Philip Emeagwali**—was
the only full-time programmer
of the most massively parallel processing
machine ever built.

In the 1980s, I was standing alone
at the farthest known frontier
of supercomputing
and I had a sense
that I was onto something that was
bigger than myself.
I was confident
that I had the best chance
of **figuring out**
how to massively parallel process
**across** a new internet
that is a new global network
of 65,536 processors.
I was confident
because I had sixteen years of
supercomputing experience
that began on June 20, 1974
in Corvallis, Oregon, United States.
I was confident
because a few insiders believed that
I could invent

how to massively parallel process,
or how to solve many problems
**at once**
instead of solving only one problem
**at a time**.

## 29.1.11 The Impossible is, Sometimes, Possible

The reason every supercomputer
programmer-hopeful, except I
**—Philip Emeagwali—**
abandoned
the massively parallel processing
supercomputer-hopeful
was that naysayers said that the
message-passing programming
of an ensemble of millions upon millions
of processors
was akin to looking at God in the face.
During my quest
for the fastest supercomputer,

I felt like I was walking alone
along a small road with a small lamp.
In the 1970s and ‘80s,
it was often said that
parallel processing
is a huge waste of everybody’s time.
As a lone wolf supercomputer scientist,
my grand challenge was to draw
the massively parallel processing
supercomputing power
that I needed to record
the fastest computational speed
in the history of the computer.
That fastest speed, that I recorded
at 8:15 in the morning
of the Fourth of July 1989
in Los Alamos, New Mexico,
United States,
was the reason major newspapers
called me the
African Supercomputer Wizard.

The wizard
in the unknown world
of the massively
parallel processing supercomputer
must command all 65,536
tightly-coupled processors.
That wizard must control
all 1,048,576
bi-directional email pathways
that married those processors together
as one seamless, cohesive supercomputer
that is a new internet *de facto*.
Back in the 1970s and ‘80s,
I was excited
because I was the unknown
supercomputer wizard.
I was excited
because I locked all 64 binary thousand
commodity processors
and locked all one binary million
commodity email wires.

I was excited
because I locked the entire parallel
processing machine
and locked them at all times.
The year
that I locked my massively parallel
processing supercomputer
was 1989.
The place
that I locked my massively parallel
processing supercomputer
was Los Alamos, New Mexico,
United States.

## 29.1.12 Overcoming Obstacles to Modern Supercomputing

Looking back
to the National Computer Conference
that took place in June 1976

in New York City,
a panel of supercomputer experts
ridiculed the massively parallel
processing supercomputer-hopeful
and dismissed the technology
as **large** and **clumsy**
and dismissed it
as a huge waste of everybody's time.
So my research in parallel processing
was science fiction
to the attendees
of that **National Computer Conference**
of 1976.
In the 1970s and '80s,
I drew my inspiration from half-human
cyborgs
that I imagined as
**super-intelligent** lizards
of Year Million
that will colonize outer space
and that are half Year Million computers

and half Year Million post-humans—
and that will make our science fiction
their Year Million non-fiction.
I theorized that
the technology of massively
parallel processing
will be at the core of the brain power
of the half-human cyborgs
of the fourth millennium.
I theorized that each half-human cyborg
will talk to a trillion cyborgs
and will communicate in real-time
and will compute together
to make our science fiction
their non-fiction.
You hear about a half-human cyborg,
but I had always imagined
that—by Year Million—
trillions upon trillions
of **immortal** post-human **cyborgs**
could roam our Milky Way galaxy

and do so to accomplish what seems
impossible to humans.
Those post-humans of Year Million
will be half-humans because
they could compute and communicate
in parallel.

## 29.1.13 How I Invented a New Supercomputer

The sequential processing
supercomputer
that is programmable was invented
in 1946.
The massively parallel processing
supercomputer
that is programmable **across**
an ensemble of processors
became faster than any supercomputer
after my experimental discovery
that occurred

on the Fourth of July 1989.
The modern supercomputer
that occupies the space of a soccer field
and costs the budget of a small nation
is a completely different beast
from the everyday computer.
The universe is huge
and is 13.82 billion years old.
The supercomputer is small
when compared to the universe.
When modeling the universe,
we cannot squeeze the universe
into a supercomputer.

I'm **Philip** **Emeagwali**.

I am a supercomputer scientist
with the spirit of a mathematician
and the soul of a physicist.
That is, I encode some laws of physics
into some equations of calculus

that I, in turn, discretize
into equations of algebra
that I further converted
into an equivalent set of
floating-point operations of arithmetic.
On the Fourth of July 1989,
I figured out
how to speed up
an ensemble of processors
that outlined a new internet
and how to speed it up
by a factor of 65,536, or more.
That invention
is my contribution
to the development of the computer.