# 22 Philip Emeagwali Supercomputer

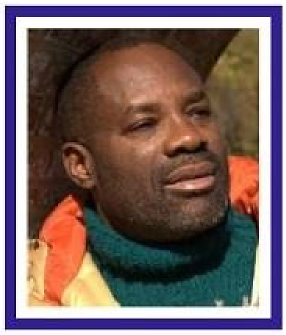

Philip Emeagwali Lecture 180125-1
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

## 22.1 Discovering What Makes Computers Faster

### 22.1.1 How I Discovered What Makes the Computer Faster

At 10:15 in the morning New York Time
Tuesday the Fourth of July 1989
I experimentally discovered

how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**
and I invented
how and why to use
that **new supercomputer knowledge**
to build a **new supercomputer**
that encircled the globe
in the way the internet does.
The reason my invention
made the news headlines
was that for the four decades
onward of 1946
the parallel processing machine
was a supercomputer-hopeful
that no supercomputer scientist
understood
what made it super.
The new supercomputer, in turn,
gave birth to
the new field of computational science.
And a new supercomputer

gives birth to
a new computational science.
The importance of computational science
was underscored in an article
that was in the May 8, 1987 issue
of *The Chronicle of Higher Education,*
the flagship newspaper
that presents news to universities.
That article was written by
computer and information technology writer
**Judith Axler Turner.**
The article was titled:
[quote]
"Some Hail 'Computational Science'
as Biggest Advance Since Newton, Galileo."
[unquote]
My Fourth of July 1989
experimental discovery
of how and why
to use massively parallel processing
to solve initial-boundary value problems
of a new calculus
and of the fastest computational physics

made the news headlines
as the biggest advance
in computational science.
My world’s **fastest**
supercomputer algebraic calculations
made the news headlines
and entered
into the June 20, 1990 issue
of the *Wall Street Journal*.
A week later,
the computer and information technology
writer, **Judith Axler Turner,**
wrote in the June 27, 1990 issue
of *The Chronicle of Higher Education*
that I—**Philip Emeagwali**—
[quote]
"took on an enormously difficult problem…
solved it alone,
has won computation's top prize,
captured in the past
only by seasoned research teams.”
[unquote]
My discovery

that made the news headlines
back in 1989
was the tipping point
of the increasing speeds
of the supercomputer.
That turning point
in massively parallel processing
supercomputing
of the Fourth of July 1989,
led to a new computer.

## 22.1.2 School Reports on Philip Emeagwali

I was asked to explain
why American children
are writing school reports
on my early life in Nigeria, Africa
and on my contributions
to the development
of the fastest supercomputer.
Teaching the groundbreaking discovery
of any historical scientist

is not mandated in U.S. schools.
However, it's included in the guidelines
known as the Core Knowledge Series.
It's included in social studies standard.
Each teacher decides
how to incorporate stories
about scientists into her curriculum.
The computer
was not invented by super-intelligent aliens
from the moon
that are disguised as humans.
Therefore, the fathers
of the modern computer
should be studied in schools
that use computers.
**Philip Emeagwali**
is studied in American schools
because I experimentally discovered
how and why
parallel processing **across** a new internet
is **faster** than computing
within any vector processing supercomputer

that was the state-of-the-art technology
of the 1980s.
Ironically, I am mostly studied in schools
in the United States,
not in my country of birth,
Nigeria (Africa).
Historically, if a scientist—such as
**Albert Einstein** or **Thomas Edison**
or **Alexander Graham Bell**—is studied
in American schools,
that scientist will later be studied
in schools all over the world.
But if a scientist is only studied in schools
in Nigeria,
that scientist will only be studied
in schools in Nigeria.
My experimental discovery
of how and why
parallel processing makes computers **faster**
first made the news headlines
in 1989
in the United States

and my discovery story
spread to American schools
and to newspapers in other countries.

### 22.1.3 A World Without Supercomputers

The human species
evolved from Africa
and evolved about two hundred thousand
years ago.
The reason our human ancestors
discovered
was to make their world
a more knowledgeable place
Our ancestors invented
to make their world a better place.
Fire
is man's first invention,
or rather man's first discovery.
Our ancestors did not discover fire
to make the news headlines
but discovered it
to make their world better.

We discover
not to make the news headlines
but to contribute to human progress.
For two hundred millennia,
we discovered
to make the world
a more knowledgeable place.
We discovered
to discover new fields of study.
The new field that I discovered
in the 1970s and ‘80s
is what is now described as
modern parallel processing supercomputing.
The supercomputer is a witness
to humanity’s most computation-intensive
problems.
The supercomputer doesn’t just solve
the toughest problems.
The supercomputer
is the modern **diving rod**
for discovering crude oil and natural gas.
The supercomputer
is the **crystal ball**

for foreseeing otherwise **unforeseeable**
global warming.
The supercomputer
is an instrument for telling the future.
I experimentally discovered
that the global circulation model
with rigorous reproducibility requirements
running **across** a new internet
that's a global network of
commodity processors
that emulates a new supercomputer
can be used to gaze across the centuries.
In my home country of Nigeria,
that is a member of OPEC
—the acronym
for the Organization
of Petroleum Exporting Countries???—
their petrodollars
is their instrument of national development
and poverty alleviation.
Crude oil and natural gas
are at the core essence
of Nigeria's **sovereignty** and identity.

## 22.1.4 My Quest for the Rosetta Stone of Computing

My experimental discovery
of massively parallel processing
that occurred
on the Fourth of July 1989
made the news headlines.
That discovery
of the parallel processing supercomputer
were highlighted in the June 20, 1990
issue of *The Wall Street Journal*
and entered
as the new supercomputer knowledge
of how to manufacture
faster computers
and the fastest supercomputers.
That discovery of the precursor
of the modern supercomputer
made the news headlines
because
it was akin to the decipherment
of the Rosetta stone

of the unknown world of supercomputers
that, in turn, will be used to
discover and recover **otherwise** elusive
crude oil and natural gas.
The Rosetta stone was discovered
in Rosetta, Egypt in 1799.
The decipherment of the writings
on the Rosetta stone
enabled historians
to decipher the previously undecipherable
writings of ancient Egyptians,
and the writings of Africans
that lived along the Valley
of the River Nile.
The Rosetta stone
enabled us to know **Imhotep**
as the father of medicine.
The Rosetta stone
enabled us to know that
the Pyramid of Giza
was a tomb
for the Egyptian Pharaoh Khufu,
who was the second ruler

of the Fourth Dynasty.
The Rosetta stone
enabled us to know the Pharaohs,
or the kings of ancient Egypt.
My technological quest
for the fastest computation
began on a sequential processing
supercomputer.
My supercomputing began
in the early morning of Thursday
June 20, 1974.
My supercomputing began
in the Computer Center
at 1800 SW Campus Way,
Corvallis, Oregon, United States.
My supercomputing
was my technological quest
for the Rosetta stone
that will enable me
to experimentally discover
how and why
massively parallel processing
must be **embodied** within

the fastest supercomputer.
In 1989, it made the news headlines
that I experimentally discovered
that parallel computing
is faster than serial computing,
and, in particular, faster
when applied to solving
the most computation-intensive
initial-boundary value problems
of a new calculus
and of the fastest computational physics.
Such problems arise
and are at the core
of computational mathematics
and computational physics.
In a sense, solving difficult problems
in parallel is observed
when dogs or lions or humans
cooperatively hunt a bigger game,
or hunt in groups.
Three thousand seven hundred [3,700]
years ago,
the Pyramid of Giza

in Africa
was cooperatively constructed **in parallel**.
That Pyramid of Giza
remains the oldest
and the only remaining
of the Seven Wonders
of the Ancient World.
Human parallel processing computing
could have been used
to solve
the grand challenge problem
that was posed four millennia ago
by the African mathematician Ahmes
and posed in his papyrus
that's the oldest mathematics literature.
So, the idea of harnessing
the power of several commodity processors,
or computers,
had been around
since the nineteen forties [1940s].
However, the massively parallel processing
supercomputer
was science fiction

in the 1940s.
In nineteen forty-six [1946],
there was only one programmable computer
in the world.
The December 13, 1947 issue
of the *New York Times*
described that first programmable computer
of 1946 as
[quote]
“the only electronic computer
among the four ‘mathematical brains’
now in use.”
[unquote]
That first supercomputer
was at Aberdeen Proving Ground,
outside Baltimore, Maryland, United States.
Fast forward four decades
from that sequential processing
supercomputer of 1946,
I was an expert
in parallel processing supercomputers
who declined a job offer
as a vector processing

supercomputer scientist
at Aberdeen Proving Ground,
Aberdeen, Maryland, United States.
That first programmable computer
was the fastest computer,
or the supercomputer,
of nineteen forty-six [1946].
In 1946, the massively parallel processing supercomputer
was dismissed as science fiction.
The January 11, 1946 issue
of *The New York Times*
wrote:
[And I quote]
“Meteorologists contemplate that
with enough of these machines
(100 was mentioned as an arbitrary figure)
area stations could be set up
which would make it possible
to forecast
the weather all over the world.
The United States would be divided
into “blocks”

penetrating into the stratosphere,
and every condition
which would have any bearing
on the weather would be analyzed.
Plans for the machine even provide
for ‘alarm’
which would give a warning
if any error in calculation occurred.”
[end of quote]
The fastest supercomputer in the world
costs the budget of a small African nation.
So manufacturing
sixty-five thousand
five hundred and thirty-six [65,536]
programmable supercomputers
and manufacturing them
in nineteen forty-six [1946]
was as laughable as spending
**65 trillion** dollars
to build a mega supercomputer today
that’s a global network of
sixty-five thousand
five hundred and thirty-six [65,536]

supercomputers.

### 22.1.5 My Discovery of the Rosetta Stone of Computing

In the 1980s,
I was the sole full-time
supercomputer programmer
of the first and the only ensemble
of sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors,
that I called a primordial internet.
I visualized my new internet
as my global network of
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors
that were identical
and that were equal distances **apart**.
That internet
that I visualized

in a sixteen-dimensional hyperspace
became my Rosetta stone.
I deciphered that Rosetta stone
at 10:15 in the morning New York Time
Tuesday the Fourth of July
of nineteen eighty-nine [1989].
That experimental discovery
of how and why
massively parallel computing works
made the news headlines
and added a new zest
—or a faster speed
and a higher speedup—
to the quest for the modern
massively parallel processing
supercomputer.
That experimental discovery
opened the door
to the vast and dynamic field of research
called massively parallel processing
supercomputing.
That modern supercomputer
is now powered by

more than ten million
commodity-off-the-shelf processors
that cooperatively solves
the most computation-intensive
mathematical problems
arising in extreme-scale
computational physics.
I stopped the full-time programming
of the most massively parallel processing
supercomputers in 1989.
I stopped supercomputing,
and doing so in parallel,
in 1989.
I stopped because I had reached
the theoretical maximum speed increase
that could be experimentally discovered
and recorded
**across** an ensemble of processors
that were the building blocks
of a new supercomputer.
My invention of a new
parallel processing supercomputer
wasn't finished by 1986.

But by 1989,
my parallel processing supercomputing
was too finished.
I wasn't getting any supercomputer
speed-up increases
after the maximum speed-up
that I experimentally discovered
on the Fourth of July 1989.
And I realized that
I have over-programmed the internet
that I visualized
as my global network
of 65,536 commodity processors.
I visualized my ensemble of processors
as supercomputing together
to emulate one seamless, cohesive internet
that is a not a new computer *per se*
but that is a new supercomputer *de facto*.

## 22.2 Changing the Way We Look at the Computer

## 22.2.1 Changing the Way We Look at the Computer



As a research supercomputer scientist,
my focus was on answering
the big, **unanswerable** questions
of the 1970s and '80s.
And, in particular, to answer
the grand challenge question
of massively parallel processing
supercomputing.
That grand challenge question was:
"How do we change the way
we look at the supercomputer
of tomorrow?"
The June 20th, 1990 issue
of the *Wall Street Journal*
reported that my experimental discovery
of the massively parallel processing
supercomputer
will change the way
we look at the supercomputer.
**In the old way**, we thought about
a conventional supercomputer

as powered by one strong ox.
That strong ox
was a metaphor for one **powerful** processor.
**In the new way**, we think about
a modern supercomputer
as powered by sixty-five thousand
five hundred and thirty-six [65,536]
chickens.
Those chickens were my metaphors
for sixty-five thousand
five hundred and thirty-six [65,536]
**weak** processors.
In that *Wall Street Journal* article
it was noted that
I experimentally discovered
that sixty-five thousand
five hundred and thirty-six [65,536]
commodity-off-the-shelf processors
that worked together
as one cohesive computing team
computed several orders of magnitude faster

than one supercomputer
powered by
only one fast vector processor.
In the 1980s,
I was the strongest **proponent**
for the chickens,
or for the modern massively
parallel processing supercomputer.
**Seymour Cray**—who made his name
in the 1980s
as the strongest **proponent**
for conventional vector processing
supercomputers—
was the strongest **opponent**
of the modern massively parallel processing
supercomputer.

### 22.2.2 A Scientific Discovery Represents the Truth

The lesson that I learned
from my quest for the fastest
massively parallel processing supercomputer
was this:

The success of a scientific discovery
is not dependent
on a [quote unquote] “not guilty” verdict
from every notable scientist.
The science fiction writings
of creative writers
are different
from the factual writings
of research scientists.
As a research supercomputer scientist,
I could not create the fastest computation
and create it
in the manner a creative writer
creates her science fiction novel.
I discovered, not created,
the fastest computation.
I experimentally discovered
the fastest computation
**across** my new internet.
My new internet was my new computer
as well as my new supercomputer.
My new internet
was a global network of

sixty-five thousand
five hundred and thirty-six [65,536]
commodity-off-the-shelf processors.
It's said that a science fiction novelist
is born to tell tales.
I say that the scientific discoverer
is born to tell truths.
Any scientific discovery
must be reproducible in a laboratory.
My experimental discovery was reproduced
by polymaths at home with
physics, calculus, and parallel processing
supercomputing.
My experimental discovery was
and can be reproduced because
it represented the truth.
It's been said that art
is what we can get away with.
I say that not discovering
is what we can't get away with.

### 22.2.3 Discovery Follows Vision

For sixteen years,
onward of June 20, 1974,
my technological vision
followed sixteen
mutually orthogonal dimensions
in hyperspace.
I followed sixteen directions.
That vision led me **across**
sixteen times
two-to-power sixteen,
or 1,048,576,
commodity-off-the-shelf processors.
Each processor communicated via email
in sixteen directions
and communicated to send and receive
initial and boundary conditions
for my 65,536 initial-boundary value
problems
and to share those intermediate answers
with its sixteen
nearest-neighboring processors.

It’s by **indirection**
that we discover new directions
for scientific progress.

### 22.2.4 Increasing the Speed of the Modern Computer

In the 1970s and ‘80s,
it was anticipated that **Moore’s Law**
will come to an end.
That means that the anticipated
speed increases
of processors and computers
will not continue to double
every two years,
as predicted by **Moore’s Law**.
With the anticipated end of **Moore’s Law**,
I anticipated that
doubling the number of computer cores
will be the only way
to double the speedup
of the modern parallel processing
supercomputer.

In the mid-nineteen seventies,
supercomputer pioneers
—such as **Seymour Cray**
and **Gene Amdahl**—
ridiculed and mocked
my parallel processing theory.
I theorized that
I could use the new internet
that I visualized
as a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors
and that I could use that new internet
to solve computation-intensive
grand challenge problems.
The poster boy
of the twenty grand challenge problems
of supercomputing
was the global circulation model
that was used to foresee
otherwise unforseeable
global warming.

In the nineteen eighties [1980s],
no automation tools existed
for automatic message passing
**across** processors or computers.
For that reason, I had to explicitly
email each processor
that I harnessed
to experimentally discover
the fastest computation.
The supercomputer textbooks
of the 1970s and '80s
wrote that harnessing
the massively parallel processing
supercomputer
to solve one of the twenty grand challenge
problems of supercomputing
is impossible.
Before my experimental discovery
that occurred on the Fourth of July 1989,
it was impossible
to synchronously email 65,536
commodity processors

and command via emails
them to compute together
as one seamless, cohesive supercomputer
that is not a new computer *per se*
but that is a new internet *de facto*
and that solved a grand challenge problem
in extreme-scale computational physics.
Since it was believed to be impossible
to parallel process,
manufacturers of vector processing
supercomputers
dismissed my parallel processing theory
as a huge waste of everybody's time.
I theorized that I could
massively parallel process
by programming a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors.
I theorized that
I could use those commodity processors
to communicate and execute

extreme-scale
petroleum reservoir simulators
and to compute and communicate them
**faster** and do so by a factor of
sixty-five thousand
five hundred and thirty-six [65,536]
increases
in the speeds of both
email communication
and arithmetical computations.
It made the news headlines
when I experimentally confirmed
my theory
and confirmed massively parallel processing
on the Fourth of July 1989.
That experimental discovery
of the massively parallel processing
supercomputer
now helps petroleum geologist
to discover and recover
otherwise elusive
crude oil and natural gas.

That's one of the practical technologies,
as well as the rich and fertile consequences
that came out of my invention
of massively parallel processing.
That experimental discovery
of massively parallel processing
is the reason
**one in ten** supercomputers
are purchased by the petroleum industry.

## 22.2.5 My Biggest Obstacle

Back in the 1970s and '80s,
my unorthodox parallel processing approach
to supercomputing
met a lot of resistance.
I was rejected and mocked
whenever I proposed that
parallel processing will work.
In those two decades,
my massively parallel processing
supercomputing premise
was that

the logic of the grand challenge problem
should determine how the problem
should be solved,
not vice-versa.
That is, it's only the laws of logic
and physics
that are **sacrosanct**,
not the technology
that, in the first place,
must bend for the laws of logic and physics.
In an article dated June 14, 1976,
the *Computer World* magazine
interviewed the foremost
supercomputer experts
that were attending
the National Computer Conference
in New York.
Those supercomputer experts
unanimously told
the *Computer World* magazine
that the supercomputer theory
of parallel processing will be
[quote]

**"a waste of time."**
[unquote]
In 1989, twenty-five thousand [25,000]
research supercomputer scientists
logged on each day
onto conventional vector supercomputers.
Due to that **skepticism** and negative press,
it was widely believed that
parallel processing
is a huge waste of everybody's time.
For that reason, I was the only person
that was logged on each day
onto the most powerful
and the most massively parallel processing
supercomputer
in the world.
I visualized my modern massively
parallel processing supercomputer
as my new internet
powered by a global network of
64 binary thousand processors.
I visualized my new internet
as married together

as one seamless, cohesive
whole supercomputer.
I visualized my new supercomputer
as an ensemble of
64 binary thousand processors
that were married together
by one binary million email wires.
In 1989, I was researching alone
on how to use
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors
and how to use them
to solve
one grand challenge problem.
In nineteen eighty-nine [1989],
it made the news headlines
that I—**Philip Emeagwali**,
an African Supercomputer Wizard
in the United States
has experimentally discovered
how to use a new internet
that's a global network of

sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors
and how to use that new internet
to solve the toughest
initial-boundary value problems
in calculus and physics.
I experimentally discovered
how to use my internet
as a massively parallel processing
supercomputer
and use that internet
to reduce the **time-to-solution**
of the most computation-intensive
grand challenge problems.
I experimentally discovered
how to speed up
from one hundred and eighty [180] years,
or sixty-five thousand
five hundred and thirty-six [65,536] days,
within only one processor
to just one day

**across** one internet.
I visualized that new internet
as a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors.
The two leading lights
of sequential and vector processing
supercomputing paradigms
—namely, **Gene Amdahl** and **Seymour Cray**,
respectively—
argued that it will be impossible
to experimentally record
the speed increase in supercomputing
that I recorded in 1989.
The 25,000 [quote unquote] "hot brains,"
or conventional supercomputer scientists
at National Science Foundation
supercomputer centers stayed with
conventional vector processing
supercomputers.
The reason those 25,000
supercomputer scientists

stayed away
from the massively parallel processing
supercomputer
was that each believed that
it will be impossible
to use 64 binary thousand processors
to solve
one grand challenge problem.
In contrast, I believed that it will be possible
to parallel process
and to do so when it seemed impossible
to do so.
Beyond faster computation speeds,
using several cores
—in both computers and supercomputers—
has other rich consequences.
One such advantage
of multicore processing technology
is that it increased
the **reliability** of the modern computer
and improved
the **fault-tolerance**

of the modern massively parallel processing supercomputer.

### 22.2.6 How to Increase the Speed of Quantum Computers

A 12-year-old writing a school report
on the contributions of **Philip Emeagwali**
to the development of the computer
asked me:
"How do we increase
the speed
of quantum computers?"
I answered:
In classical parallel computing,
I experimentally discovered
how to solve
all sixty-five thousand
five hundred and thirty-six [65,536]
challenging problems
and how to solve them
at the **same time**

and how to solve them **across**
a global network of one million
forty-eight thousand
five hundred and seventy-six [1,048,576]
commodity email wires
that fed data and answers
from initial-boundary value problems
and fed them to and from
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors.
I discovered
how to solve extreme-scale problems
in computational physics.
My physical surroundings entered into
my initial-boundary value problems
of a new calculus
and of the fastest computational physics.
I’m surrounded
by the air and the water
that entered into
my general circulation models
that I executed **across**

my ensemble of 64 binary thousand commodity processors.

## 22.3 Changing the Way We Look at the Computer

### 22.3.1 The Inside of a Quantum Computer

I was asked:
"What does a quantum computer look like?"
The inside of a quantum computer
is one of the coldest places
in the known universe.
The inside of a quantum computer
is minus 273 degrees Celsius.
The inside of a quantum computer
is 150 times colder than
**interstellar** space.
The first quantum computer

is not quite a quantum computer.
That first quantum computer
is a monolithic black box
that’s 12 feet by eight feet by ten feet tall.
That first quantum computer
fills a small bedroom.
The quantum computer
will not make the massively
parallel processing supercomputer
obsolete.
The reason is that a quantum computer
will not be a general-purpose computer.
The quantum computer
might look like a refrigerator
because it needs to be cooled.
In quantum computing,
the computer memory
and the processor
must be isolated.

## 22.3.2 How to Reduce 180 Years to One Day

I experimentally discovered

how to reduce the **time-to-solution**
from one hundred and eighty [180] years,
or sixty-five thousand
five hundred and thirty-six [65,536] days,
within one processor
to only one day of **time-to-solution**
**across** a new internet
that's powered by
a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors.
Those processors were used to
solve all sixty-five thousand
five hundred and thirty-six [65,536]
challenging problems
and solve them at the **same time**.
I began sequential processing
supercomputing
in the summer of
nineteen seventy-four [1974]
and I began by wanting to discover

the massively parallel processing
supercomputer
in nineteen seventy-four [1974].
I began parallel processing supercomputing
without being able to visualize
the modern supercomputer
and visualize it
in nineteen seventy-four [1974].
I began modern supercomputing
without being able to even articulate
the modern supercomputer
and to do so back in nineteen seventy-four
[1974].
In the 1970s, my grand challenge
was to visualize the shape
of my internet
and to visualize it
as a 7,918-miled diameter internet.
And, most importantly,
articulate that internet
as the source of the fastest computations,
both present and future.
But back in nineteen seventy-four [1974],

or even in the late nineteen seventies, 866
I wasn't sure how my
experimental discovery
of massively parallel processing
will be contextualized
with calculus, algebra, arithmetic,
codes, and emails.
The reason the speedup of
sixty-four binary thousand
that I experimentally discovered
made the news headlines
in nineteen eighty-nine [1989]
was that
the new knowledge
that parallel processing **works**
could not be proven wrong.
Like any scientific discovery,
my experimental discovery
was one hundred percent doubt-free.
That experimental discovery
was the end-product
of an acid test type experiment
that I conducted

**across** a new internet
that's a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors.
The supercomputer
that is sixty-four
binary thousand times faster
than the computer
is immensely more complex
than the computer.

## 22.3.3 One Thing Supercomputers Can Do Now That They Couldn't Do 30 Years Ago

I'm often asked:
"What's the one thing supercomputers
can do now
that they couldn't do
30 years ago?"
My answer was this:

The fastest supercomputer of today
is one hundred million times faster
than the fastest supercomputer
of 30 years ago.
That fastest supercomputer
is powered by ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
commodity-off-the-shelf processors.
Each of those ten million processors is proof
that parallel computing
is not a huge waste of everybody's time,
as was alleged
in the June 14, 1976 issue
of the *Computer World* magazine.
Thirteen years after that negative article
in the *Computer World*,
I saved everybody's time by
experimentally discovering
—on the Fourth of July 1989—
how to reduce 180 computer-years

of **time-to-solution**
to only one supercomputer-day
of **time-to-solution**.
I experimentally discovered
how to put
65,536 processors, or as many computers,
inside one supercomputer
That is, I experimentally discovered
a new internet.
My **experimental discovery**
was the new knowledge
of how computers can compute **faster**
and how supercomputers
can compute fastest.
I experimentally discovered
how to compute **faster**
and how to do so by a factor of 65,536.
That experimental discovery
of the massively parallel processing
supercomputer
made the news headlines.

That experimental discovery
was reported in the June 20, 1990 issue
of *The Wall Street Journal*.
That experimental discovery
is the reason American students
are writing school reports
on the contributions of
**Philip Emeagwali**
to the development of the computer.

### 22.3.4 Who Stole My New Supercomputer?

Back in 1974, parallel processing
was a future computer industry.
For that reason,
my parallel processing
supercomputer research
was largely ignored and I worked alone
for the subsequent sixteen years.
The inventor

hid the struggles and obstacles
he overcame
and hid them
to make his invention appear
as if it was invented
in one brilliant Eureka Moment!
It took me fifteen years
to invent the new supercomputer
that was read in fifteen minutes.
During those fifteen years,
I was ridiculed, rejected, and mocked
as the bush fowl
that crowed
in the language of another village.
For me, **Philip Emeagwali**,
I defined massively parallel processing
as my technological quest
for the speed of now
from the supercomputer of tomorrow.
I did so because
the supercomputer of today

872

will become the computer of tomorrow.
In my quest
for the fastest supercomputer,
the details did not matter
as long as the ending is happy.
After 1989, my massively parallel processing
supercomputer discovery
became politicized.
The obsession of haters
was to knock me off my perch.
After 1989,
a team of American supercomputer scientists
tried to steal the credit
for the new supercomputer
that I invented alone.
For the four decades onward of 1974,
I held those **new supercomputer** thieves
**at bay**.
After forty-four years
of their rampant criticisms and derisions
and mockeries

of the new supercomputer
that I invented alone,
that team of thieves
started claiming the credit
for the new supercomputer
that I invented **alone**.
Those thieves told everybody
that they invented the new supercomputer
that encircled the globe
in the way the internet does.
They stole my original drawings
and renamed my new supercomputer
from [quote unquote]
"Philip Emeagwali Supercomputer."
They put their names on my invention.
My name was completely erased
as the **sole inventor**
of my new supercomputer.
My invention of a new supercomputer
that was comprised of
an ensemble of commodity processors

that were connected
and that were connected
as a new internet
was stolen by supercomputer scientists
who publicly condemned my new internet
and mocked it since 1974.
To avoid legal prosecution
and public shaming,
they stole my new supercomputer
under a pseudonym, or false name.
That was the most audacious theft
in the history of the computer.
That theft
brings into question
their honesty
and moral character.
They stole the credit
for the invention
of my new supercomputer
that was a new internet
and stole my new technology

under a number of false identities
and fluid pseudonyms
and stole my new computer
under a number of throw-away
email addresses
and throw-away
mobile telephone numbers.
A paradigm shifting discovery
goes through three stages of rejection
before its final acceptance.
My quest for the fastest supercomputer
began on June 20, 1974
at 1800 SW Campus Way,
Corvallis, Oregon, United States.
Two years later, the June 14, 1976 issue,
of the *Computer World*
—the flagship publication
of the computing industry—
wrote an article that was titled:
"Research in Parallel Processing
Questioned as 'Waste of Time???? FC'"

Back in the 1980s, my 1,057-page |
research report
on massively parallel processing
was ridiculed, mocked, and rejected
and was thrown into the trash.
That mocked report
was the first stage of rejection.
On the Fourth of July 1989,
I experimentally discovered
how and why
an ensemble of the slowest processors
in the world
could solve the toughest problems
in extreme-scale computational physics
and solve them faster
than the fastest
vector processing supercomputer
in the world.
I was mocked by vector processing
supercomputer scientists
and mocked because the internal timer

of my front-end computer
had a timer resolution
of only one millionth of a second.
For that inadequate timer resolution,
I could not reproducibly
time my one billionth of a second
that, in turn, was needed
to time my world record computation speed
of billions of calculations per second.
It took me six months
to resolve my hardware timer issue,
and indirectly resolve it via algorithms.
That internal timer error
made my experimental discovery
to seem irreproducible
which, in turn, gave the naysayers
a pretext to reject
the massively parallel processing
supercomputer.
In the 1960s through ‘80s,
research supercomputer scientists

cited **Amdahl's Law**
and supercomputer textbooks
when explaining
why the massively parallel processing
supercomputer
cannot be harnessed
and why it cannot be used
to solve the toughest
initial-boundary value problems
arising in modern mathematics
and in extreme-scale computational physics.
But on the Fourth of July 1989,
I—**Philip Emeagwali**—
experimentally discovered
that those titans
of the world of vector processing
supercomputers—such as,
**Seymour Cray**—were wrong.
The supercomputer textbook
was wrong because
I discovered that

the impossible-to-compute
is, in fact, possible-to-compute
and I did so
by using an ensemble of the slowest
65,536 processors
and using that ensemble
to solve the toughest problems
in extreme-scale computational physics
and solve those problems faster
than the fastest supercomputer
could solve them.
In 1989, that discovery
of the massively parallel processing
supercomputer
won me the **top prize**
in the field of supercomputing
and made the news headlines
and made vector processing
supercomputer scientists
that ridiculed, mocked, and rejected me
to beg me to become my co-inventor

and to become my new best friend.
Some went as far as claiming my invention
as their invention.
In the 1970s and '80s,
I was disowned
by four American universities
who each objected to my pursuing
my research on the massively
parallel processing supercomputer.
Today, each of those universities
have claimed me as their [quote unquote]
"famous alumni."
Rejection is the common denominator
of any breakthrough scientific discovery.
If the supercomputer scientist
of the 1970s and '80s,
understood the new supercomputer
that computes **across**
an ensemble of processors
that is a new internet
then, that new supercomputer

is not new to him or her.
And if my massively parallel processing
supercomputer research discovery
was accepted in the 1980s,
then I was not the inventor
of that new supercomputer.
A discovery
that is accepted, at first,
is not a groundbreaking discovery.
A discoverer
that is accepted, by all,
did not make a groundbreaking discovery.
When my experimental discovery
of the massively
parallel processing supercomputer
made the news headlines,
I—**Philip Emeagwali**—
the new supercomputer inventor
had no research home
other than my apartment.

## 22.3.5 A Supercomputer Gold Medal for Nigeria



Before my 1989 experimental discovery
of how and why
parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**,
my country of birth, Nigeria,
did not have a frame of reference
in the textbook on the history of science.
That historical omission
is striking
when you realize that technology
has improved the standard of living
in Nigeria and Africa.
My two grandparents—at 17 Mba Road,
Umudei Village, Onitsha (Nigeria)—were
never vaccinated.
My two grandparents—at 6 Wilkinson Road,
Onitsha (Nigeria)—never spoke

into the telephone.
My four grandparents
never saw the Atlantic Ocean
that was 190 kilometers away
from Onitsha (Nigeria).
It's technology that made it possible
to travel farther than Lagos (Nigeria)
and to send and receive an email
to and from the United States to Nigeria.
Without that frame of reference
of the Nigerian supercomputer scientist
that contributed
to the development of the computer
studying science and technology
in a Nigerian school
will be like the Nigerian child
watching the Olympics Games
for contributions to human progress
and then not seeing
the Nigerian contingent
at the opening ceremony.

At 10:15 in the morning New York Time
Tuesday the Fourth of July 1989,
the U.S. Independence Day,
I experimentally discovered
the massively parallel processing
supercomputer
and I discovered the technology
as faster than
the sequential processing supercomputer.
That supercomputer discovery
was equivalent
to the Nigerian national anthem
being played
during the Olympic gold medal ceremony
for contributions
to the development of the computer.
After my 1989 experimental discovery
of the massively parallel processing
supercomputer,
I became the Nigerian
that was most searched for

on the Internet.
My experimental discovery
of the massively parallel processing
supercomputer
made the news headlines
and caused a sensation
within the supercomputing community.
Every supercomputer scientist
knew that, in theory,
an ensemble of processors
could be **fabricated** or **manufactured**.
My experimental discovery
of the massively parallel processing
supercomputer
made the news headlines because
it was the first time ever
that an ensemble of the slowest processors
computed faster than
the fastest processor
that was not a member
of an ensemble of processors.

886

For that reason,
my experimental discovery
changed the way
we look at the computer
and the supercomputer.
In the old way, or in the 1980s,
the fastest 500 supercomputers
in the world
were powered by
one powerful, custom-made
vector processing unit.
That is, the old supercomputer
used only one processor.
In the new way,
or after my 1989 experimental discovery,
the supercomputer is supercomputing
with an ensemble of
commodity-off-the-shelf processors
that are identical.
My experimental discovery
of the massively parallel processing

supercomputer
is a milestone
that made the news headlines
because
it has rich and fertile consequences.
My experimental discovery
of the massively parallel processing
supercomputer
was recorded in the June 20, 1990 issue
of *The Wall Street Journal*
and recorded because
it opened a new window
in science and technology.
Parallel processing
is the most important discovery
in the history of computational science.
The importance of computational science
was underscored in an article
that was in the May 8, 1987 issue
of *The Chronicle of Higher Education,*
the flagship newspaper

888

that presents news to universities.
The article was titled:
[quote]
“Some Hail ‘Computational Science’
as Biggest Advance Since Newton, Galileo.”
[unquote]
In most of the twentieth century,
parallel processing
was *terra incognita*
on the map of the supercomputer.
Since the 1940s,
supercomputer scientists
speculated that parallel processing
could play a role
in the supercomputer of the future
but had no idea
how to actually parallel process
an extreme-scale problem
in computational physics.
It was on the Fourth of July 1989,
I experimentally discovered

how to massively parallel process
the toughest problems
arising in extreme-scale algebra,
modern calculus, and computational physics.
Parallel processing
has taken the computer
into a new era,
making the impossible-to-compute
possible-to-compute.

### 22.3.6 More Information

I'm **Philip Emeagwali**.

I have only presented to you
a teaspoon full of ocean water.
Trying to acquire in sixty minutes
the knowledge
that took me sixty years
to acquire

is like taking a teaspoon full of water from the Atlantic Ocean.

890