27 Philip Emeagwali Computer: How I Invented the Fastest Supercomputer

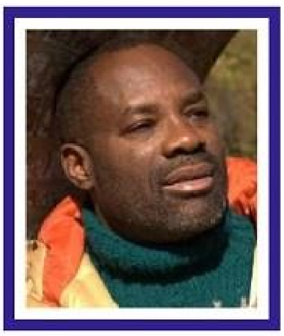

Philip Emeagwali
info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture 180127-2
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

27.1 How I Discovered What Makes the Computer Faster

27.1.1 What is Philip Emeagwali Known For?

I'm Philip **Emeagwali**.
at **emeagwali** dot com.
I know a lot about **Philip Emeagwali**
because I'm **Philip Emeagwali**.
To discover
is to hear the unheard,
to see the unseen,
and to know the unknown.
But we can only discover something
if and only if
that thing pre-existed.
What's the point in searching
for a black goat in a dark room
if there was no goat in the room?

At 10:15 in the morning New York time
Tuesday the Fourth of July 1989,
the US Independence Day,
I saw something
that was previously unseen
by any human being.
I saw the slowest processors
in the world
compute together
to compute faster than the fastest
supercomputer
in the world.
I saw that fastest computation **across**
a new internet
that is a global network of
65,536 processors.
I experimentally discovered
how and why
parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest.**
I **experimentally invented**

how to use that **new**
supercomputer knowledge
to build a **new supercomputer**.

27.1.2 School Reports on Philip Emeagwali

In the United States,
students are writing school reports
on my early life in Nigeria, Africa,
and on my contribution
to the development of the supercomputer.
The writing of school reports
on the life and discoveries
of role models in science—such as
Archimedes, Galileo Galilei, Isaac Newton,
Charles Darwin, Thomas Edison,
and Albert Einstein—is not a school
requirement.
However, stories about scientists
are included in the US elementary
and middle school teaching guidelines
that are called the **Core Knowledge Series**.

Scientific biographies 1083
are studied in US social studies standard.
Each US school teacher
decides how to incorporate
stories about scientists into her curriculum.
And stories about how
the modern computer was invented
is a favorite in US schools.
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.
I'm studied in American schools because
I experimentally discovered
how and why
the technology of massively
parallel processing
makes modern supercomputers

that are computing **across** 1084
a global network of processors,
or **across** a new internet,
fastest.
I am the computer scientist
that is most studied in elementary,
middle, and high schools
in the United States.
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.
Mungo Park,
the Scottish explorer of the River Niger
is only studied in schools in Nigeria.
Mungo Park
is not studied in schools in Scotland

or United States.
My experimental discovery
that occurred
on the Fourth of July 1989
and that discovery
of how and why
parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**
first made the news headlines
in 1989
and in the United States.
That discovery story
spread into school reports
in elementary, middle,
and high schools in the United States.
That discovery
not only redefined how others saw me
but, also, how I saw myself.

27.1.3 A World Without Supercomputers

The human species 1086
evolved from Africa
and evolved nearly two hundred thousand
years ago.
The reason our human ancestors discovered
was to make their world
a more knowledgeable place.
Our human ancestors invented
their hunting tools
to make their world a better place.
Fire is man's first invention,
or rather man's first discovery.
Our human ancestors did not discover **fire**
to make their news headlines.
They discovered **fire**
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 our world
a more knowledgeable place.

We discovered
to discover new fields of study.
The new field of study
that I discovered,
in the 1970s and ‘80s,
is called massively parallel processing
supercomputing.
I discovered that new field of study
by conducting my sixteen-year-long quest
between fields,
between classical physics
and modern mathematics,
between abstract calculus
and extreme-scale algebra,
and between
the most computation-intensive
floating-point arithmetical operations
and the largest ensemble of processors.
The supercomputer
is a witness to humanity’s
most computation-intensive problems.

The supercomputer
doesn't just solve the toughest problems
in extreme-scale computational physics.
The supercomputer
is the modern **diving rod**
for discovering more 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 discovered that
the global circulation model
that is executed **across** a new internet
that's a global network of
commodity processors,
or a global network as many computers,
and that emulates a new supercomputer
can be used to gaze **across** the centuries.
My country of birth, Nigeria,
is the sixth largest crude oil and natural gas

producer
within the Organization of Petroleum
Exporting Countries, called OPEC.
For Nigeria,
oil is the instrument of development.
Oil is at the core essence
of Nigeria's **sovereignty** and identity.
Since one in ten
massively parallel processing
supercomputers
are purchased by the petroleum industry,
the modern supercomputer
is an instrument of economic development
in OPEC nations.

27.1.4 My Quest for the Rosetta Stone of Computing

The reason my experimental discovery
of the massively parallel processing
supercomputer
that occurred
on the Fourth of July 1989
made the news headlines in 1989

was because
knowing for the first time
how and why
parallel processing works
was akin
to my **decipherment** for the first time
of the Rosetta stone
for the fastest supercomputer
and for my discovering
of how to use the modern supercomputer
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 **decipherment** of the Rosetta stone
enabled us to know **Imhotep**
as the father of medicine.
The **decipherment** of the Rosetta stone
enabled us to know that
the **Pyramid of Giza**
was a tomb
for the Egyptian **Pharaoh Khufu**.
The **Pharaoh Khufu**
was the second ruler of the Fourth Dynasty.
The **decipherment** of the Rosetta stone
enabled us to know the Pharaohs,
or the kings of ancient Egypt.
My technological quest
for the **decipherment**
of the fastest computation
began on a sequential processing
supercomputer
and ended
on the Fourth of July 1989
on a massively parallel 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
the fastest supercomputer.
That technology is called
massively parallel processing.
I experimentally discovered
that massively parallel supercomputing
is **faster** than serial supercomputing,
and, in particular, it is faster
when applied to solving
the most computation-intensive problems
arising in extreme-scale

computational physics.

27.1.5 History of Parallel Processing

The concept of hunting together,
or in parallel,
is seen when a pack of lions
attack a lone elephant.
Ancient Egyptians implicitly used
human parallel processing,
or harnessed the power of
many Egyptian human computers,
to design the Pyramid of Giza.
The Pyramid of Giza
is the oldest
and the only remaining
of the Seven Wonders
of the Ancient World.
Human parallel processing
could have been used
to solve the grand challenge problem

that was posed 1094
by the African mathematician,
named Ahmes,
and posed in his papyrus.
Ahmes' papyrus
is the oldest literature in mathematics.
The thought of parallel processing,
or how to harness
the power of many computers,
had been around
since the nineteen forties [1940s].
But parallel processing
was science fiction
in the 1940s.
In nineteen forty-six [1946],
there was only one programmable computer
in the world.
That computer was the supercomputer
of nineteen forty-six [1946].
So manufacturing sixty-five thousand
five hundred and thirty-six [65,536]
supercomputers
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.
The massively parallel processing
mega supercomputer
has no upper speed limit.
The supercomputer's practical speed limit
is determined by economics,
not science or technology.

27.1.6 My Discovery of the Rosetta Stone of Computing

In the 1980s, I was the sole full-time
programmer
of the first and only ensemble
of sixty-five thousand

five hundred and thirty-six [65,536] 1096
commodity-off-the-shelf processors
that I visualized as the first internet.
I visualized that internet
as my global network of
sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors.
That internet became my Rosetta stone.
I **deciphered** that Rosetta stone
in nineteen eighty-nine [1989].
That experimental discovery
added a new zest
—or a faster speed
and a higher speedup—
to the quest for the modern supercomputer.
That modern supercomputer
is now powered by
more than ten million commodity
processors.
I stopped full-time programming

of massively parallel processing
supercomputers
on the Fourth of July 1989.
I stopped supercomputing
in 1989
because I had experimentally discovered
the theoretical maximum speed increase
that could be extracted
from supercomputing
in parallel.
My invention wasn't finished by 1986.
But by 1989, my supercomputing
was too finished.
I wasn't getting any supercomputing
speed increases
after my experimental discovery
of massively parallel processing
that occurred
on the Fourth of July 1989.
And I realized that
I have over-programmed my new internet
that's a global network of

65,536 processors,
or a global network of as many computers,
computing together
to emulate one seamless and cohesive
massively parallel processing
supercomputer.

27.1.7 Changing the Way We Look at the Computer

The first supercomputer
that I programmed
was marketed as the fastest computer
in the world
and was manufactured in December 1965.
For the two and half decades,
onward of December 1965,
the supercomputer lost its direction.
The leaders of thought
in supercomputing
didn’t have a clear, long-term strategy
of how to leapfrog the technology
from being powered by
a singular sequential processing unit

that only executed
one million instructions per second
to millions
of central processing units
that could execute the set of
floating-point arithmetical operations
that was needed to solve
the toughest problems in modern calculus,
extreme-scale algebra,
and computational physics
and that was needed
to solve them
a **million** billion trillion times
faster.
My focus
in massively parallel supercomputing
was on **answering**
the **unanswerable** questions
of the 1970s and '80s.
And in answering
the biggest question
in the field of massively parallel
supercomputing.

That grand challenge question was: 1100
"How do we change the way
we look at the supercomputer?"
My supercomputer discovery
that occurred on the Fourth of July 1989
was summarized
in the June 20th, 1990 issue
of the *Wall Street Journal*.
The *Wall Street Journal*
reported that my experimental
supercomputer discovery
of how to massively parallel compute
and how to record
the fastest speeds in computation
will change the way
we look at the modern supercomputer.
In the old way, we thought about
the conventional supercomputer
as powered by one strong ox.
That strong ox
was a metaphor
for one powerful vector processing
supercomputer.

In my new way, we think about
the modern supercomputer
as massively parallel processing
and as powered by
sixty-five thousand
five hundred and thirty-six [65,536]
chickens.
Those chickens were my metaphors
for my sixty-five thousand
five hundred and thirty-six [65,536]
weak processors.
In that *Wall Street Journal* article
it was noted that I—**Philip Emeagwali**—
experimentally discovered
that an ensemble of sixty-five thousand
five hundred and thirty-six [65,536]
processors
that computed together
as one cohesive computing team
computed several orders of magnitude faster
than one powerful vector processing
supercomputer
that was not a member

of an ensemble of computers
and/or supercomputers.
In the 1980s,
I was the lone wolf programmer
of sixteen massively parallel processing
ensemble of processors
and I was the strongest **proponent**
for the chickens,
or for my ensemble
of the 65,536 slowest processors
in the world.
Seymour Cray—who made his name
in the 1980s
as the strongest **proponent**
for conventional vector processing
supercomputers—
was the strongest **opponent**
of modern parallel processing
supercomputers.

27.1.8 A Scientific Discovery Represents the Truth

1103

With my speed increase
of a factor of 65,536,
you'd think the whole United States,
much less the supercomputer community,
would have crowned me, in 1989,
“the father of the modern supercomputer”
or something.
Yet, my experimental discovery
of the massively parallel processing
supercomputer
was, at first, controversial.
The lesson that I learned
from supercomputing since June 20, 1974
was this:
The success of a scientific discovery
is not dependent
on a “not guilty” verdict
from every notable scientist.
I learned that
the fiction writings of creative writers
are different
from the factual writings
of research scientists.

1104

As a supercomputer scientist,
I could not create the fastest computation
and create it
in the manner a creative writer
creates her novel.
I discovered, not created,
the fastest computation.
I discovered the fastest computation
across the first internet.
My new internet was my new computer.
My new internet
was a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
processors.
It's said that a 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 discovery was reproduced
by polymaths at home with

physics, calculus, and supercomputing.
My 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.

27.1.9 Discovery Follows Vision

For the sixteen years
that followed June 20, 1974,
my supercomputing vision
followed sixteen
mutually orthogonal directions
in hyperspace.
That supercomputing vision
led me **across**
sixteen times
two-to-power sixteen
commodity-off-the-shelf processors
that were identical

and that were equal distances
apart.
Each of my 65,536 processors
communicated via emails
that I sent to
and that I received from
my sixteen-bit long
binary identification numbers.
Each binary number was my metaphor
for an email address
that had no @ sign,
that had no dot com suffix,
and that travelled
along sixteen mutually orthogonal directions
and travelled
to communicate with its sixteen
nearest neighboring
processors.
I discovered that
it’s by **indirection**
that we discover new directions
for scientific progress.

27.1.10 Increasing the Speed of the Modern Computer

1107

In the 1970s and '80s,
I anticipated that **Moore's Law**
will end.
That means that
the anticipated speed increase
of 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 processors
will be the only way
to double the speed
of the modern computer.
In the mid-nineteen seventies,
supercomputer experts
—such as **Seymour Cray**
and **Gene Amdahl**—
ridiculed the theory of parallel processing
and dismissed it

by mocking the technology
as a beautiful theory
that lacked experimental confirmation.
In contrast, I theorized that
I could use a new internet
that I visualized
as my global network
of sixty-five thousand
five hundred and thirty-six [65,536]
commodity-off-the-shelf processors,
and use them to solve
computation-intensive
grand challenge problems,
such as those extreme-scale
global circulation models
that are used
to foresee otherwise unforeseeable
global warming.
In the nineteen eighties [1980s],
no automation tools existed
for **automatic message passing**
across processors.

For that reason, I had to explicitly 1109
email each processor
before I could harness that processor
and harness it
to experimentally discover
the fastest computation
that could be recorded
across an ensemble of processors.
In the 1970s and ‘80s,
the vector processing
supercomputer scientist
believed that it would be impossible
to **synchronously** email
65,536 processors,
or to email them in parallel.
Seymour Cray, who designed
seven in ten vector processing
supercomputers sold in the 1980s,
believed that it would be impossible
to simultaneously command
all those processors
and to harness them
to compute together

and do so as one seamless, cohesive supercomputer 1110
supercomputer
that will be faster than
any of his vector processing supercomputers.
For those reasons, my research reports
—of the 1970s and ‘80s—
in massively parallel processing
supercomputing
were ridiculed, mocked, and rejected.
The manufacturers of supercomputers,
such as IBM Corporation,
as well as publications
in the world of computing,
such as *Computer World* magazine,
ridiculed, mocked, and rejected
the parallel processing supercomputer
and dismissed the technology
as a huge waste of everybody’s time.
The first time that I programmed
a supercomputer
was on June 20, 1974
at age nineteen [19].
That supercomputer

was the first
to compute at the speed of
one million instructions per second.
That supercomputer
was at 1800 SW Campus Way,
Corvallis, Oregon, United States.
When I was nineteen years old [19],
the sequential processing supercomputer,
that was invented in 1946
was twenty-eight [28] years old.
I visualized the **CPU**
of that sequential processing supercomputer
to be a network of only one processor.
In mathematical lingo
and within a zero-dimensional space,
the supercomputer that I programmed
at age nineteen [19]
is a global network
of only one **CPU**.
From that zero-dimensional space,
I progressively visualized
my global network of processors
through sixteen dimensional space

and even through sixty-four
dimensional space.
In a parallel processing supercomputer
research report
that was rejected and mocked
in September 1983,
and that was rejected
after I had paid
the seventy-five dollar publication fee,
I theorized that
a global network of
two-raised-to-power sixteen processors
should be 64 binary thousand times
faster than one supercomputer
that is powered by only one processor.
I theorized that global network of processors
as married together
as one seamless, cohesive supercomputer
and married together
by sixteen times
two-raised-to-power sixteen,
or one binary million email wires.

I theorized that
I could program that global network of
sixty-five thousand
five hundred and thirty-six [65,536]
processors
and that I could program them
to **emulate** a new supercomputer
that is faster than any
vector processing supercomputer.
I theorized that
I could harness those 64 binary thousand
processors
and use them to communicate
and to execute
extreme-scale, multi-physics
petroleum reservoir simulators
and use them to compute
and to communicate
and compute faster by a factor of
sixty-five thousand
five hundred and thirty-six [65,536].
I experimentally confirmed
my theory of massively parallel processing

and I confirmed that new supercomputer
on the Fourth of July 1989,
the US Independence Day.
For me, the Fourth of July 1989
was the day of fire, the day the massively
parallel processing supercomputer
became the fire we can't put out.
After that experimental discovery
of the Fourth of July 1989,
trying to stop the acceptance of
the massively parallel processing
supercomputer
became like trying to stop midnight.
That experimental discovery
made the news headlines
and was in the June 20, 1990 issue
of *The Wall Street Journal*.
That experimental discovery
is the reason **one in ten** supercomputers
are purchased
by the petroleum industry alone.
That experimental discovery
has rich and fertile consequences.

27.1.11 My Biggest Obstacle

Back in the 1970s and '80s,
my massively parallel processing
supercomputer research
was ridiculed, mocked, and rejected.
I conducted my research alone because
the likes of **Seymour Cray**
offhandedly dismissed
the parallel processing supercomputer
as a huge waste of time.
As the principal investigator, **Seymour Cray**,
alone, had a thousand assistants,
or a thousand actual investigators.
I was an actual investigator
that did not have any research assistant.
Because I did not have
any research assistant,
I had to know more
about parallel processing supercomputing
in order to invent
a new supercomputer

1116

and to invent it
without any research assistant.
My supercomputing premise
was that the logic
of the grand challenge problem
should determine how the problem
should be solved,
not vice-versa.

27.1.12 My New Thinking

In supercomputing, it's only the laws of logic and physics
that are **sacrosanct**,
not the technology.
To me, that supercomputing premise
was a more intuitive approach
because our universe
is 13.7 billion years old
while our abstract mathematical knowledge
is not even
13.7 thousand years old.

That is, our mathematical knowledge
has been around for merely
one millionth of the age of our universe.
And yet, our knowledge of the computer
is only half of 137 years.
That is, the human knowledge
of the computer
is half of one hundredth
the time of our mathematical knowledge
that, in turn, is one millionth
the age of our universe.
And I—**Philip Emeagwali**—was 37 years old
when I figured this timeline.

27.1.13 Parallel Processing Mocked

In an article dated June 14, 1976,
the *Computer World* magazine
interviewed supercomputer experts
that were attending
the National Computer Conference
in New York City.
Those supercomputer experts

unanimously told the *Computer World*
that the supercomputer theory
of parallel processing
will be
[quote]
“**a waste of time**.”
[unquote]
In 1989, twenty-five thousand [25,000]
supercomputer scientists
logged on each day
onto conventional vector supercomputers.
I had sixteen
parallel processing supercomputers
that were the **precursor**
to the modern supercomputer
that nobody else, except I,
was willing to log into.
Due to that **skepticism** and negative press,
only I—**Philip Emeagwali**—
logged on each day
onto the most powerful
and the most massively
parallel processing supercomputer

in the world.
I visualized the **precursor**
to the modern supercomputer
as my new internet
powered by a global network of
64 binary thousand
commodity-off-the-shelf processors
that were identical
and that were equal distances **apart**.
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 harness the total
supercomputing power
of those sixty-five thousand

five hundred and thirty-six [65,536]
commodity-off-the-shelf processors
and how to use them
to solve
one grand challenge problem
in computational physics.
In nineteen eighty-nine [1989],
it made the news headlines
that I—**Philip Emeagwali**,
the African Supercomputer Wizard
in the United States
has experimentally discovered
how to program a new internet
that's a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
processors
and how to program that new internet
to solve the toughest problems
in physics and calculus.
I experimentally discovered
how to program that new internet
as a new supercomputer

and how to program it
to reduce the **time-to-solution**
of computation-intensive
grand challenge problems
in science, technology, engineering,
and mathematics.
I experimentally discovered
how to speed up computations **across**
that new internet
and speed it 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 that new internet
that's a global network of
sixty-five thousand
five hundred and thirty-six [65,536]
processors.
The two leading lights
of conventional supercomputing
—namely, **Seymour Cray**

and **Gene Amdahl**— 1122
wrote that it will be **impossible**
to experimentally record
the speed increase in supercomputing
that I recorded in 1989.
The reason 25,000 in inverted commas
“hot brains,”
or conventional vector processing
supercomputer scientists
that were at the U.S.
National Science Foundation
supercomputer centers
stayed with conventional, vector processing
supercomputers
was that they believed that
it will be **impossible**
to harness the total computing power
of an ensemble of
65,536 slowest processors
that defined
the precursor to the modern supercomputer
that computes in parallel.

The reason I conducted
my massively parallel processing
supercomputer research
as a lone wolf
was that I believed it will be possible
to do what seemed impossible,
namely, solve 65,536
initial-boundary value problems
of modern calculus
and computational physics
and solve them **at once**,
instead of solving them
one problem **at a time**.
Beyond faster computation speeds,
using several processors, or 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 supercomputer.

1123

27.1.14 How to Increase the Speed of Quantum Computers

A young supercomputer scientist
asked me:
“How do we increase
the speed
of the quantum computer?”
I explained that
in classical parallel processing
supercomputing,
I experimentally discovered
how to solve
sixty-five thousand
five hundred and thirty-six [65,536]
challenging initial-boundary value problems
of calculus and physics
and how to solve **simultaneously**
and solve them **across**
one million forty-eight thousand

five hundred and seventy-six [1,048,576]
email wires.
Those email wires
fed data and answers
to and from
my ensemble of sixty-five thousand
five hundred and thirty-six [65,536]
commodity processors.
I experimentally discovered
how to solve
the toughest problems in physics
and solve them
by using the modern parallel processing
supercomputer
as my instrument of extreme scale
computational physics.

27.1.15 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 twelve 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 new supercomputer
that solves
the toughest problems

in computational physics,
such as general circulation modeling
to foresee otherwise unforeseeable
climate changes.
The quantum computer
might look like a refrigerator
because it must be cooled.
In quantum computing,
the computer memory and processor
must be isolated.

27.1.16 How to Reduce 180 Years to One Day

The Holy Grail
of the 21st century supercomputing
is the "billion-fold speedup challenge"
of parallel processing **across**
one billion processors
that could encircle the Earth
and encircle it as a new internet.
On the Fourth of July 1989,
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 just 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
that solved all
sixty-five thousand
five hundred and thirty-six [65,536]
challenging initial-boundary value problems
of physics and calculus
and solves them **simultaneously**.
I began supercomputing in the summer of
nineteen seventy-four [1974]
and began by wanting to discover
the modern supercomputer
in nineteen seventy-four [1974]
and began without visualizing

the modern supercomputer
in nineteen seventy-four [1974]
or even articulating
the modern supercomputer
in nineteen seventy-four [1974].
In the 1970s, my grand challenge
was to visualize the shape
of the modern parallel processing
supercomputer
and visualize the technology
as a new internet
that is a global network of processors,
or a global network of computers,
that has a diameter of 7,918 miles.
And, most importantly,
my grand challenge was to articulate
that global network
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,
I wasn't sure how my
experimental discovery

will be **contextualized** 1130
with calculus, algebra, arithmetic, codes,
and emails.
The reason the speedup of
sixty-four binary thousand
that I experimentally discovered
on the Fourth of July 1989
made the news headlines
in nineteen eighty-nine [1989]
was that it could not be proven wrong.
Like any scientific discovery,
my discovery
was one hundred percent **doubt-free**.
That 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 modern supercomputer

that is sixty-four
binary thousand times faster
than the conventional supercomputer
is immensely more complex
and far more difficult to program.

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

https://www.fastcompany.com/3040378/9-things-computers-can-do-now-that-they-couldnt-do-a-year-ago

I’m often asked:
“What’s the one thing supercomputers
can do now
that they couldn't do
30 years ago?”
I answered:
The fastest supercomputer of today
is one hundred million times faster
than the fastest supercomputer

of 30 years ago. 1132
That fastest supercomputer
is powered by ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
computing cores.
Each of those ten million computing cores
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.
On the Fourth of July 1989,
and thirteen years after that negative article
in the *Computer World*,
I saved everybody's time
by experimentally discovering
how to reduce
180 years of **nonstop computing**
within only one processor

to only one day of **nonstop supercomputing across** 1133
across
an ensemble of 65,536 processors.
I experimentally discovered
how to put
65,536 commodity processors
inside one new supercomputer
that is a new internet
that is a global network of those
65,536 commodity processors.
My experimental discovery
was how supercomputers
can compute **faster**.
I experimentally discovered
how to compute **faster**
and how to do so
by a factor of 65,536.
That experimental discovery
that occurred
on the Fourth of July 1989

made the news headlines. 1134
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
or googling
for the contributions of
Philip Emeagwali
to the development of the computer.