# 28 How I Invented a New Supercomputer

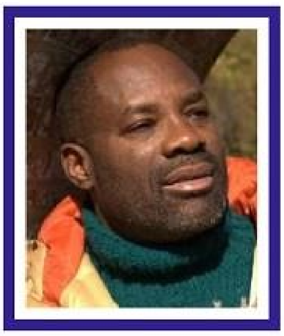

Philip Emeagwali
info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture 180913-1
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

On the Fourth of July 1989,
the U.S. Independence Day,
I experimentally discovered
a new supercomputer
that could solve 65,536 problems

**at once**,
or that could process information
in parallel.
Before that invention
of the Fourth of July 1989,
parallel processing
was dismissed as science fiction
by the authors of textbooks
on supercomputing.
Before that invention,
solving many problems **at once**
was ridiculed
as a beautiful theory that lacked
an experimental confirmation.
On the Fourth of July 1989,
the technology
of the parallel processing
supercomputer
became a tested reality.
Before the Fourth of July 1989,
the massively parallel processing

supercomputer
was like an elephant
with a super body
and the brain of an ant.
The massively parallel processing
supercomputer
was a transformative technology
that moved detailed data modeling
from dream to reality.
The massively parallel processing
supercomputer
is the technology
that enabled
precision petroleum reservoir simulation
of the Niger-Delta oilfields
of the southeastern region
of my country of birth, Nigeria.
The invention
of the massively parallel processing
supercomputer
opened the door

for air-cooled supercomputers
with no requirement for liquid cooling.

### 28.1.1 Philip Emeagwali Supercomputer

I'm **Philip Emeagwali**.
I made headlines in major U.S.
newspapers
for discovering
the massively parallel processing
supercomputer,
such as the June 20, 1990 issue
of the *Wall Street Journal*.
I am a supercomputer scientist
who began programming
supercomputers,
exactly sixteen years earlier,
on June 20, 1974
in Corvallis, Oregon, United States.

The supercomputer 1139
is the world's fastest computer.
The supercomputer
is a living machine that grows
with each increase in speed.
At 8:15 on the morning
of the Fourth of July 1989
in Los Alamos, New Mexico, United States,
I experimentally discovered
the fastest supercomputer.
I invented
a new supercomputer
that computes **across** a new internet
that is a new global network of
65,536 tightly-coupled processors.
I mathematically and experimentally
invented
how to tackle 65,536 challenging
initial-boundary value problems
arising in mathematics and physics.
I invented

how to solve those grand challenge problems 1140
and how to solve them
in a one-to-one corresponded manner.
I invented
how to use emails
to and from sixteen-bit long addresses,
each with no @ sign or dot com suffix,
and how to use those emails
to stitch those 65,536 problems **together**.
I invented
how to stitch problems together
as the original grand challenge problem.
To reach that new frontier
of human knowledge
demanded new techniques and technologies,
such as a new arithmetic,
a new algebra, a new calculus,
a new computer,
and, most importantly, a new internet.
The massively parallel processing
supercomputer

was not invented
by the team of 25,000
vector processing supercomputer scientists
of the 1980s.
I conducted
the parallel processing experiment
that led to the discovery
on the Fourth of July 1989
of the massively parallel processing
supercomputer.
I—**Philip Emeagwali**—was the only person
that invented
how to harness
the total supercomputer power
of 65,536 separate processors.
After my invention, the fastest one thousand
supercomputers in the world
are supercomputing **across**
thousands or millions of
commodity-off-the-shelf processors.
That shift

from one processor
to one million processors
is the biggest paradigm shift
in the history of the computer.

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

Since 1989, school children
were asked to write a school report
on the contributions of **Philip Emeagwali**
to the development
of the modern supercomputer.
Back in 1989, it made the news headlines
that a lone wolf
African supercomputer wizard
the United States
had invented
how to solve
the toughest problems
arising in modern calculus
and computational physics,

and mathematically invented
how to solve
65,536 initial-boundary value problems
of modern mathematics
and invented
how to solve them **at once**.
That invention
occurred on the Fourth of July 1989
and is called the
"massively parallel processing
supercomputer."
I—**Philip Emeagwali**—
was that African supercomputer scientist
that was in the news
back in 1989.
I was in the news because
I experimentally discovered
that the fastest computing speeds
in modern supercomputing
must always be recorded
with parallel processing technology,

rather than with vector processing technology.

### 28.1.3 Group Thinking is the Enemy of Progress

For the fifteen years,
onward of June 20, 1974,
I conducted my supercomputer research alone.
I did so alone
because I was ridiculed, mocked,
and rejected
by all-white research teams
that were exclusively programming
only sequential and vector processing supercomputers.
As a black African-born
supercomputer scientist in the United States,
I felt like I was in exile
wherever I am.
I'm in exile in the United States.
I was in exile in Africa.

I was in exile 1145
in the then uncharted territory
of the massively parallel processing
supercomputer.
A multidisciplinary
supercomputer research team
could comprise of one thousand
scientists and engineers.
Each member
of that supercomputer research team
was at the frontier of knowledge
of physics.
Or at the frontier of knowledge
of mathematics.
Or at the frontier of knowledge
of computer science.
To discover parallel processing
required both theory and experiments
and required a polymath,
rather than a mathematician.
To invent
the massively parallel processing
supercomputer

required a polymath 1146
that was simultaneously at home
at the frontiers of physics, mathematics,
and computer science.
It took me sixteen years
of advanced training,
onward of March 25, 1974,
in Oregon (United States)
as well as weekly attendances
at 500 research seminars
of the 1980s in the District of Columbia
and Maryland (United States),
to become that triple threat
and that polymath
that is at home
at the frontiers of knowledge
in physics, mathematics,
and computer science.
Most importantly,
I was the only research scientist
that gave massively parallel processing
research lectures
to audiences of research

computational physicists 1147
at the United States national laboratories.
I gave research lectures
to research mathematicians
at the international congress
of mathematicians.
I gave research lectures
to research computer scientists
of the two premier computer societies
in the world, namely,
The Computer Society of the IEEE
and the Association for
Computing Machinery.
In the late 1970s and early ‘80s,
I was rejected
because white research scientists
dismissed me
before they heard me
give my research lectures
on how I invented
the massively parallel processing
supercomputer.

The audio and video recordings
of my lectures on the new supercomputer
that I invented
are posted at emeagwali dot com.
To work cohesively
as a supercomputer research team
demands that each team member
follow the team leader.
The supercomputer research teams
of the 1970s and '80s
were coerced to group think
and were technologically brainwashed
to **group think** only in the direction of
conventional vector processing
supercomputing.
The leading proponents
of vector processing supercomputers
were the leading opponents
of parallel processing supercomputers.

### 28.1.4 Sometimes, The Impossible is Possible

In 1989,

there were 25,000 users
of vector processing supercomputers.
I was the only fulltime programmer
of the handful of
massively parallel processing
supercomputers
of the 1980s.
**Gene Amdahl** and **Seymour Cray**,
the two leading opponents
of the parallel processing supercomputer,
argued that it will forever
remain impossible
to parallel process through as many as
eight processors or computer cores.
In the 1940s through '60s,
the group thinkers
in the field of supercomputing
focused only on the
sequential processing supercomputer
technology.
In the 1970s and '80s,
the group thinkers in supercomputing
focused only on the

vector processing supercomputer 1150
technology.
In those two decades,
I was forced to work as a lone wolf
supercomputer scientist
that was not a member
of a 400-person research team.
For that reason,
I wasn't indoctrinated
into group thinking that
vector processing
must always be superior to
parallel processing.
Like other black African scientists
of the 1970s United States,
I wasn't accepted
into any supercomputer research group.
That rejection
forced me to forge a different path
to the modern parallel processing
supercomputer.
That rejection
forced me to think individually

on how to harness the power of 1151
the massively parallel processing
supercomputer
and how to invent the technology
and know it, for the first time,
as the engine that drives
the modern parallel processing
supercomputer.
The reason my invention
of the massively parallel processing
supercomputer
made the news headlines
and was recorded in the June 20, 1990 issue
of *The Wall Street Journal*
was that the parallel processing
supercomputer technology of today
was then dismissed and abandoned
by the leaders of thought
in supercomputing—namely,
**Gene Amdahl** and **Seymour Cray**—
and was then rejected by their followers
who described parallel processing
as a huge waste of everybody's time.

## 28.1.5 Seeing the Unseen

What kept me moving forward
towards the first fastest computation
that I executed
on the Fourth of July 1989
and executed **across**
my ensemble of 65,536
tightly-coupled processors
was a back-of-the-envelope,
theoretical calculation
that I did in the 1970s.
From that back-of-the-envelope calculation,
I theoretically discovered that—**in theory**—
two-to-power sixteen processors
that could only calculate
forty-seven thousand
three hundred and three [47,303]
floating-point arithmetical operations
per second per processor
can be integrated **across**
a small internet

that is a new global network of
65,536 tightly-coupled
commodity-off-the-shelf processors
with each processor
operating its own operating system
and with each processor
having its own dedicated memory
that shared nothing with each other.
That is, I discovered, *a priori*,
that my new internet
will become my new supercomputer
that will be faster than
the old vector processing supercomputer
that was the industry's
state-of-the-art technology
and that computes less than
3.1 billion calculations per second.
In the 1980s,
I was the lone full time programmer
of the only
massively parallel processing supercomputer
that was an ensemble of
65,536 processors.

I was a lone wolf
research supercomputer scientist
that was not a member
of a 400-person
supercomputer research team.
Such research teams were funded
by the United States government.
Such research teams were funded
to bring the best brains
from United States national laboratories
and use that intellectual capital
to bear on a grand challenge problem,
or the toughest problem
in computational physics.
Such problems were described
as grand challenges
because they were perceived
to be otherwise unsolveable.
What made the news headlines
was that I—**Philip Emeagwali**—had
invented
how to harness those computing units,
namely, a new internet

that is a new global network of 1155
65,536 processors.
I invented
how to harness those ensemble of processors
to process simultaneously
and to process together
and to do so as one cohesive whole unit,
or to process in parallel,
or to process as one
integrated super processor
that is a never-before-seen supercomputer
that is the precursor
to the modern supercomputer
that I expect to become
the computer of tomorrow.
That is, in the 1970s and ‘80s,
my technological quest
was for how to massively
parallel process **across** a new internet
and how to massively parallel process
and do so at a time
theorists were theorizing their way
through parallel processing.

A theory is not a discovery.
A theory is an idea
that is not positively true.
In the 1970s and '80s,
parallel processing
was ridiculed, mocked, and rejected.
Parallel processing was scorned
as a beautiful theory
that lacked experimental confirmation.
It's not possible to experimentally discover
the fastest speeds in supercomputing
and discover that fastest speed
by merely and only theorizing about
how to achieve the fastest speeds
via the massively parallel processing
supercomputer-hopeful.
In my experimental confirmation
of the new fastest speed in supercomputing
that occurred on the Fourth of July 1989,
and occurred in Los Alamos, New Mexico,
United States,
I confirmed that the massively
parallel processing supercomputer

can be programmed
to increase productivity
and to reduce the **time-to-solution**
of the toughest problems
in extreme-scale computational physics
and to reduce that **time-to-solution**
from 65,536 days, 180 years,
on only one processor
to just one day
**across** a new internet
that is a new global network of
65,536 commodity-off-the-shelf processors
that were identical
and that were equal distances
**apart**.
After my invention
of the massively parallel processing
supercomputer
the technological progress that followed
in the subsequent three decades
was a series of cleanups and refinements
and re-discoveries.
Three decades later,

my invention 1158
of the massively parallel processing
supercomputer
enabled China
to copy that massively parallel processing
supercomputer
and to use the technology
to massively reduce
their **time-to-solution**
and reduce it from
thirty thousand [30,000] years,
or ten million
six hundred and forty-nine thousand
six hundred [10,649,600] days,
of **time-to-solution**
on only one processor
to just one day **across** ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
commodity-off-the-shelf processors.
That was how China
—that did not invent
the massively parallel processing

supercomputer—
massively parallel processed its way
to the world's fastest supercomputers.

### 28.1.6 My Quest for a New Supercomputer

To enter into the unknown world
of the massively parallel processing
supercomputer,
required that I begin sequential processing
supercomputing
and begin it
on June 20, 1974
in Corvallis, Oregon, United States.
I began supercomputing
when I was only nineteen years old.
I began supercomputing
at 1800 SW Campus Way,
Corvallis, Oregon, United States.
When I began supercomputing,
I was new in the United States,
having arrived on March 24, 1974
after receiving a scholarship letter

1160

that was dated September 10, 1973.
Then and now,
and excluding athletes
only a dozen or so Nigerian teenagers
were invited each year
to study in the United States
and offered a four-year academic scholarship
to do so.
As an aside,
many of the famous computer pioneers
of today
were also teenagers in 1973,
and were born
at about the same time
—of August 23, 1954—
that I was born.
I was born in Akure
in the heart of Yoruba Land
in the Western Region
of the British West African colony
of Nigeria.
I was born in Yoruba Land
to Igbo parents

who grew up in Onitsha, Igbo Land, 1161
of south eastern Nigeria
and lived in Kano, Hausa Land,
in the late 1940s.
My father was educated,
for the six years
inclusive from 1942 to ’47,
at Christ the King College,
Onitsha, in the southeastern region
of Nigeria.
My father left Kano, at age 29,
to come to Akure in 1950 to work as a nurse
in the small hospital in Akure.
In the early 1950s, my father’s salary
of five pounds a month supported
a dozen extended family members
both in Akure and Onitsha, Nigeria.
I was born in a Boy’s Quarter
that was at the junction of Eke Emeso Street
and Oba Adesida Road, Akure, Nigeria.
I was born on Monday
under the sign of Virgo.
The U.S. president, on the day I was born,

was Dwight D. Eisenhower, 1162
a Republican.
Since Nigeria was then a British colony,
Queen Elizabeth the Second,
was the Head of State
of Nigeria.
The Governor-General of Nigeria
on the day that I was born,
Sir John Macpherson, represented
Queen Elizabeth.
When I was born,
there were five adults sharing
a tiny room and a tiny parlour,
namely, my father
Nnaemeka James **Emeagwali**,
my mother
Iyanma Agatha **Emeagwali**,
my aunt
Nkemdilim Grace **Azuokwu,**
my cousin Vincent **Emeagwali,**
and my cousin Charles **Emeagwali**.
Aunt Nkemdilim came from Onitsha
three days after I was born

and on August 26, 1954. 1163
Aunt Nkemdilim came for three-month
*ine omugor*.

**Steve Jobs**
was nineteen years old then
and lived an hour's drive
from Corvallis, Oregon,
and lived in Portland, Oregon.
**Bill Gates**
was nineteen years old then
and lived three hour's drive
from Corvallis, Oregon
and lived in the border state of
Washington, United States.
The difference
between the other research physicists,
research mathematicians,
and research computer scientists,
and myself—**Philip Emeagwali,**
that was a massively parallel processing
supercomputer scientist—
was that those researchers were plowing

the **frontier** of computational physics
or the **frontier** of modern calculus
or the **frontier** of abstract algebra
or the **frontier** of the vector processing
supercomputer.
Most research supercomputer scientists
of the 1970s and '80s
were plowing frontiers of knowledge
that had already been plowed.
As a massively parallel processing
supercomputer scientist
of the 1970s and '80s,
I did not believe in re-plowing
the frontier of knowledge
of the sequential processing supercomputer
or in re-plowing
the frontier of knowledge
of the vector processing supercomputer
that had already been plowed.
Re-plowing
the frontiers of scientific knowledge
that had already been plowed
makes as little difference

as searching for new crude oil and natural gas 1165
in the Oloibiri Oil Field
of Bayelsa State of Nigeria.
The Oloibiri Oil Field
was the first oilfield discovered
in West Africa.
The Oloibiri Oil Field
dried up after twenty years
of oil exploration
and was abandoned back in 1978.
Comparing the new massively
parallel processing supercomputer technology
to the old vector processing
supercomputer technology
was like comparing
constructing a brand new highway
from Cairo (**Egypt**, North Africa)
through Lagos (**Nigeria**, West Africa)
that is 95 hours of non-stop driving
of six and half thousand kilometers
and constructing that brand new highway

to Johannesburg (**South Africa**) 1166
that is 98 hours of non-stop driving
of nearly 7,000 kilometers
from Lagos (Nigeria)
and comparing
that super highway construction project
to the superficial re-paving
of the existing half an hour drive,
45 kilometer highway
between my ancestral hometown
of Onitsha (**Nigeria**) and Awka (**Nigeria**).
That is the reason
the massively parallel processing
supercomputer
costs the budget of a small nation.
Since the first sequential processing
supercomputer
was invented in 1946,
the price-performance of the supercomputer
dropped continuously and exponentially.
If that pace of technological progress
upholds,
the supercomputer of today

will become the computer of tomorrow. 1167
On the Fourth of July 1989,
I—**Philip Emeagwali**—entered into
the history book.
I was profiled in books
such as the one that was titled:
"History of the Internet."
I am the subject of school reports
because I experimentally discovered
a new way of looking at the
modern computer.
The June 20, 1990 issue
of *The Wall Street Journal*
recorded that I experimentally discovered
a new paradigm,
called massively parallel processing
supercomputing.
That new paradigm in supercomputing
changed how we compute
and changed how we solve
the toughest problems
in modern calculus
and extreme-scale computational physics.

1168

That new paradigm
changed how we solve
the system of partial differential equations
that governs
initial-boundary value problems
of modern mathematics,
such as general circulation modeling
to foresee otherwise unforeseeable
climate changes.
That new paradigm
changed how we solve
the toughest problems
in extreme-scale computational physics
and changed how we solve
those problems in parallel
and changed how we solve
those problems
at the fastest supercomputer speeds.
I was asked:
"What makes a discovery or an invention
newsworthy?"
I answered that,
first and foremost, the new knowledge

that is embodied
within the scientific discovery
or within the technological invention
must compete with new celebrity gossips
and new hot button political
and religious issues.
For those reasons, a newsworthy
contribution to human knowledge
must be bold and strange,
or be a new paradigm,
such as be the invention
of the massively parallel processing
supercomputer
that solves the toughest problems
in extreme-scale computational physics
and solves them
in an unorthodox manner
that challenges expectations.
Discovering a new paradigm
in supercomputing
prompts the leaders of thought
in the world of computers

to ask for the discoverer's 1170
telephone number.
That's how and why **Steve Jobs**
got my telephone number
and contacted me in about June 1990.
Back in the 1970s and '80s,
parallel processing was ridiculed, mocked,
and rejected
as a beautiful theory
that lacked experimental confirmation.
Today, we take it for granted
that the modern massively
parallel processing supercomputer
harnesses the total computing power
of up to ten million
six hundred and forty-nine thousand
six hundred [10,649,600]
commodity-off-the-shelf processors
that are identical.
But to **Steve Jobs** my invention
of the parallel processing supercomputer
was like science fiction
becoming non-fiction.

## 28.1.7 The Polymath vs. The Mathematician

A research mathematician
that is trained only in mathematics
or a research physicist
that is trained only in physics
or a research computer scientist
that is trained only in computer science
cannot [cannot, cannot]
theoretically and experimentally
discover
the massively parallel processing
supercomputer
that is the **precursor**
of our modern supercomputer.
Only a **polymath**,
that is simultaneously at home
at the frontier of the system of
coupled, non-linear, time-dependent,
and state-of-the-art
partial differential equations
of modern calculus

and a **polymath** that is at home 1172
at the frontiers of extreme-scale algebra
and computational physics
and a **polymath** that is at home
at the frontier of the massively
parallel processing supercomputer
and only that **polymath**
can invent
how to harness
64 binary thousand tightly-coupled
commodity-off-the-shelf processors
that shared nothing with each other
and invent
how to solve a grand challenge problem
in extreme-scale
computational fluid dynamics.
The *SIAM News*
is the flagship bi-monthly news journal
of record
of the Society of Industrial
and Applied Mathematics.
The articles in the *SIAM News*
are about new mathematics

that are yet to enter
into any mathematics textbook.
The articles in the *SIAM News*
are written by research mathematicians
and are written
for research mathematicians.
The reason my new mathematics
for computational mathematicians
computing **across**
an ensemble of processors
was the cover story of the May 1990 issue
of the widely-read *SIAM News*
was that I was a research
extreme-scale computational mathematician
who pushed the **frontiers**
of modern calculus
and extreme-scale algebra
and computation-intensive arithmetic.
I pushed the **frontier** of the modern calculus
by a distance of thirty-six [36]
partial derivative terms
that were not in any calculus textbook
that was used by the petroleum industry.

I invented
how to solve the most large-scale system
of equations of algebra
that must be solved
to recover otherwise unrecoverable
crude oil and natural gas.
I invented
how to execute
the fastest floating-point
arithmetical operations
for extreme-scale computations
in science and engineering.
The reason computational mathematicians
call them floating-point operations
is that the position of the decimal point
is constantly tracked.
The reason I won the highest award
in supercomputing
was that I was a supercomputer scientist
who pushed the **frontiers**
of the most massively parallel processing
supercomputer.
And the reason

research computational physicists
invited me
to give lectures at NASA,
Los Alamos National Laboratory,
and various U.S. research laboratories
was that I pushed the **frontiers**
of extreme-scale computational physics.
The reason I programmed
supercomputers alone
was that I discovered that
—as a polymath
that is a multi-disciplinary scientist—
I could do research alone.
After sixteen years of sequential processing
supercomputing
and vector processing supercomputing
and parallel processing supercomputing,
I realized that I had more
hands-on programming experience
on the first massively parallel processing
supercomputer
that parallel processed **across**
a new internet

that I envisioned as a new global network of 
65,536 processors
and that I knew how to solve
extreme-scale systems of equations
of modern algebra
and that I was the first person to understand
how to solve them at
unheard of speeds.
Looking back, I devoted
two hundred thousand [200,000] hours
of my life thinking about
the massively parallel processing
supercomputer
that is a new internet,
which, in turn, is more than
anybody else did.
My command of materials
and my possession
of the interdisciplinary fluency
that made it possible
for me to conduct my research alone
and for me to deliver my lectures
on my contributions to the development of

the massively parallel processing
supercomputer
show in my lectures
that are posted at emeagwali dot com.
My Eureka Moment
during which I invented
the precursor to the modern supercomputer
occurred in Los Alamos, New Mexico,
United States
and occurred at 8:15 in the morning
of Tuesday the Fourth of July 1989,
the U.S. Independence Day.

## 28.1.8 Philip Emeagwali Supercomputer

At age nineteen
and on June 20, 1974,
in Corvallis, Oregon, United States,
where I started programming
the first supercomputer
that was rated at
one million instructions per second,
there were hardly any

supercomputer scientist
in the world
that was interested in programming
the parallel processing supercomputer.
In 1974 and earlier,
the technology of parallel processing
was science fiction
and was scorned as a beautiful theory
that lacked an experimental confirmation.
Sixteen years later, I was the only full time
supercomputer programmer in the world
that was programming
the most parallel processing supercomputer
ever built.
Today, all supercomputer scientists
are massively parallel processing
**across** millions upon millions
of tightly-coupled
commodity-off-the-shelf processors
that shared nothing with each other.
The reason my experimental discovery
of the massively parallel processing
supercomputer

that occurred 1179
on the Fourth of July 1989
made the news headlines
was that the supercomputing community
said that
I made the impossible-to-compute
possible-to-compute
and did so
by giving me the top prize
in the field of supercomputing.
When I won that top award in 1989,
the 25,000 supercomputer scientists
in the world
had little faith
in the parallel processing supercomputer
that is the precursor
to the modern supercomputer.
Those 25,000 supercomputer scientists
were deeply entrenched
behind their vector processing
supercomputers,
and, for that reason, they avoided
programming the ensemble of

65,536 tightly-coupled processors 1180
that outlined and defined
the precursor
to the modern supercomputer.
My invention
of the parallel processing supercomputer
opened the door
to the manufacturing
of hybrid supercomputers
that comprised of hundreds of cabinets
and tens of thousands of computer nodes
and as many **GPU** accelerators.
The **GPU**
is the acronym for graphics processing unit,
just as the **CPU**
is the acronym for central processing unit.
The **GPU**
is the soul of the computer
while the **CPU** is the brain of the computer.

Today, the parallel processing technology
is used to reduce
30,000 years of **time-to-solution**

on only one processor
to only one day of **time-to-solution**
on the fastest supercomputer.
That extraordinary reduction
of **time-to-solution**
on the modern supercomputer
translates to increased productivity.
In the new parallel processing
supercomputer paradigm,
the **time-to-market**
was reduced from decades to months
and the **time-to-discovery**
was reduced from years to days
and the **time-to-forecast**
was reduced from months to minutes.
The most important question
in computer science
or in extreme-scale computational physics,
called the Grand Challenge Problem
by the United States government
is this:
"How do we increase
the speed of computers?"

I answered that grand challenge question 1182
with my experimentally-verified invention
of how to massively parallel process
and how to do so **across**
a new internet
that is a new global network of
65,536 commodity processors.
Those processors
were identical.
Those processors
were equal distances apart
from each other.
Those processors
shared nothing with each other.
That invention
of the parallel processing supercomputer
is my contribution
to the development of the computer.

I'm **Philip Emeagwali**.
at emeagwali.com.
Thank you.