40 Father of the Modern Supercomputer

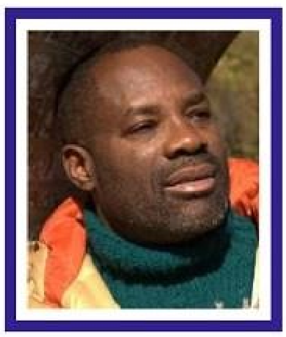

Philip Emeagwali
info@emeagwali.com 202-203-8724

Philip Emeagwali Lecture 170618
Visit http://emeagwali.com for complete transcripts of 100+ lectures.
Video: https://YouTube.com/emeagwali
Podcast: https://SoundCloud.com/emeagwali

40.1 How I Changed the Way We Look at the Supercomputer

40.1.1 Who's Philip Emeagwali, the Father of the Modern Supercomputer?

One million years ago,
our human ancestors looked like apes.
In one million years, or Year Million,
our human descendants
will ridicule us
as looking like humans.
In one million years,
we might only have living Silicon
as our post human descendants.
In one million years,
our post human descendants
could live forever.
In Year Million
the aliens on Earth will be us.
I envision post human aliens

of Year Million 1832
as thinking across a
10,000-mile diameter
Cosmic SuperBrain
that will sprawl across an
epic landscape of their
eighth super-continent
that will enshroud
our seven land continents
and enshroud the Earth
with their Year Million electronic wires.
I foresee our descendants
to be half-humans
and half thinking machines,
as well as being futuristic
technological clutter.
I believe that the grandchildren
of our grandchildren
will not use their internet
the way we use our internet.
Their internet

1833

will be within them
while our internet
is around us.
They will not need supercomputers
because they will be
the fastest supercomputers.
In 1989,
it made the news headlines
that I—**Philip Emeagwali**—
had experimentally discovered
how to execute
the fastest computations
and how to do so **across** a new internet
that is a global network of
65,536 tightly-coupled processors
with each processor
operating its own operating system
and with each processor
having its own dedicated memory
that shared nothing with each other.
Or a new internet

1834
that is a new global network of
as many tiny computers.
I invented
how to parallel process
or how to execute a **billion** set of
computer instructions
and how to execute them **at once**
or how to execute them in parallel
and across a **billion processors.**
That invention
enables the modern supercomputer
to compute a **billion** times **faster**.
I began my quest for that new internet
in 1974 in Oregon (United)
and as a then janitor-mathematician
who put away his **slide rule**,
called an analog computer
that he brought from Onitsha (Nigeria).
Back in 1970 Onitsha (Nigeria)
I was the only person that
had a slide rule.
I also put away my **log table**,

or computing aid, 1835
of 1967 of Gbenoba Road,
Agbor (Nigeria).
I put away both my slide rule
and log table
to learn how to compute fastest
when solving
a large system of equations
of algebra
and how to solve those equations
on a sequential processing
supercomputer
and to learn how those equations
arose from calculus and physics.
I began programming that
supercomputer
on Thursday June 20, 1974
in Monmouth, Oregon,
in the Pacific Northwest region
of the United States.
I entered my programs

into a time-shared supercomputer
that was at 1800 SW Campus Way,
Corvallis, Oregon.
I entered my programs
through my remote job entry terminal.
However, my supercomputer-related
knowledge, discoveries, and inventions
grew over the decade and half
that is onward of June 20, 1974.

40.1.2 Early Obstacles in Supercomputing

By the early 1980s,
I was ahead
in the supercomputer race
for the fastest calculations.
But as a lone wolf
black supercomputer scientist,
I was perceived as a threat,
and as a loose cannon

and was denied access
to vector processing supercomputers.
I was forced to back off
just before I could make a
supercomputing breakthrough.
In 1989, it made the news headlines
that a lone wolf
African supercomputer wizard
won a top US prize
and won it for
inventing
how to solve a set of
65,536 equations
each a system of
366 equations of algebra
that was the world record
in large-scale algebra.
That African supercomputer wizard
invented

how to solve those world record
algebraic equations
and solve them
across a new internet
that is a new global network of
65,536 tightly-coupled processors
that shared nothing with each other.
I—**Philip Emeagwali**—
was that African internet scientist
that was in the news
back in 1989.
That experimental discovery
of massively parallel processing
was my signature discovery
as a computational mathematician
and my Eureka invention
as an internet scientist.
That was my Eureka moment
because I saw two new internets
—one I theorized in 1974,

the other I experimented with 1839
and constructively reduced to practice
in 1989.
Both new internets
were previously unseen.
That invention
of how to do many things (or processes)
at once
and do them within processors
that are within a new **internet**
is the reason American school children
are doing school reports titled:
[quote]
"The Contributions of **Philip Emeagwali**
to the Development of the Computer."
[unquote]

40.1.3 Paradigm Shifting From Computer to Supercomputer

I first discovered parallel processing
as a theorized internet
that is a new global network of
64 binary thousand,
or two-raised-to-power sixteen,
identical computers.
I visualized and theorized
65,536 identical computers
as evenly distributed around the Earth.
I visualized those two-raised-to-power
sixteen identical computers
as being equal distances **apart**
and with as much uniformity as possible
in computers
and regularity as possible
in email wires interconnections.
Over sixteen years,

my theorized internet evolved
towards a global network of
64 binary thousand computers
that encircled a ball
in a sixteen-dimensional space.
I called that ball a HyperBall
and that name gave rise to the term
"Philip Emeagwali's HyperBall."
In the early 1980s,
I re-shaped and re-invented
my hyper-ball
as a roundy hyper-cube
that is still a hyper-ball.

40.1.4 Making the Impossible Possible

In the 1980s,
I—like any other black African-born
but naturalized American scientist—
couldn't conduct supercomputer research

in the U.S. national laboratories 1842
that was the primary place
to conduct research in supercomputing.
Viscerally, I felt that I was on a
hot track to invent
how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**
and that I was pursuing
a supercomputer invention
that couldn’t be invented
under the vision
of any U.S. national laboratory
or be invented
as part of a supercomputing
research team.
In my unsuccessful hiring talks
I provided broad brushstrokes

to research computational physicists 1843
and to research
computational mathematicians
that were unfamiliar
with a hyper-global network of
65,536 tightly-coupled processors
that shared nothing with each other.
My broad brushstrokes
sounded like science fiction
and an empty pipe dream.
In June 1974,
I made a leap of the imagination
when I leaped
from a theorized global network of
64 thousand human computers
around the Earth
to a theorized global network of
64 thousand electronic computers
around the Earth.
I made that leap of the imagination
because that's what humans do:

extrapolate from the **known**
to the **unknown**.
The genius is the ordinary person
that found the extraordinary
in the ordinary.

40.2 Father of the Fastest Supercomputer

40.2.1 Lone Wolf at the Farthest Frontier

At first, it seems hard to select
the **father of the fastest**
supercomputer.
And select him—never her—
from the field of the 25,000
supercomputer scientists
of the 1980s.
In the 1970s and '80s,

the **upper echelon** of the field of
massively parallel supercomputing
was very **sparsely populated**.
In the 1980s,
I used my fingers
to count the number of programmers
of the few massively
parallel processing supercomputers
that were powered by
thousands of processors
that were ever built.
For example, in the 1980s,
I, alone, controlled sixteen
massively parallel supercomputers.
The reason each
massively parallel supercomputer
was idle and available to me alone
was that no programmer
in the world
then knew how to harness it’s up to

1846

64 binary thousand
processors
and, in particular, harness them
to compress the **time-to-solution**
of the most extreme-scale problem
in computational physics.
Each massively parallel supercomputer
was waiting for me, **Philip Emeagwali**—
its then sole fulltime programmer—
to log into it.
And after I logged into each
massively parallel supercomputer,
I felt like I was home alone.

40.2.2 Father of the Fastest Supercomputer

So, who is the **father**
of the fastest supercomputer?
Each of the few dozen

1847

supercomputer scientists
with the courage
to occasionally log into
a massively parallel supercomputer,
and do so back in the 1970s and 80s,
was in his or her own way
a **father** or a **mother**
of the modern supercomputer
that he or she was programming.
That programmer
was a **son** or a **daughter**
of the massively parallel supercomputer
powered by
64 processors
that he or she
was programming in the early 1980s.
That programmer
was a **nephew** or a **niece**
of the massively parallel supercomputer
powered by four processors

that he or she
was programming in the late 1970s.
However, only one programmer
could be at the **farthest frontier**
of the most massively parallel
supercomputer.
In the late 1980s, that **farthest frontier**
was outlined by a new global network of
64 binary thousand
commodity-off-the-shelf processors.
I—**Philip Emeagwali**—invented
a new internet
and I did so
at the farthest frontier of computing,
where the fastest computations occur.
My **new internet**
was powered by a new global network of
64 binary thousand
commonly-available processors,
or a global network of

as many identical computers.
I'm giving this lecture
as the first eyewitness
from that farthest frontier
of massively parallel supercomputing.
I'm giving this lecture
because I was the lone wolf
at that farthest frontier.
I'm giving this lecture
because I was at that farthest frontier
and was at that **uncharted territory**
when everybody else said that
parallel processing
is a huge waste of everybody's time.

40.2.3 Retrospective on Fastest Supercomputing

To experimentally discover
parallel processing

and to discover it in 1989 1850
was to make the impossible-to-compute
possible-to-compute.
In 1989, I invented
how to parallel process.
I demonstrated
how to parallel compute
and **how to** synchronously communicate
and **how to** simultaneously compute
and **how to** do both across
a new internet.
First, I invented that new internet
as my new global network of
65,536
commonly available processors.
Second, I also invented that new internet
as my new global network of
65,536 identical computers.
In 1989, it made the news headlines
that I—**Philip Emeagwali**—

an African Supercomputer Wizard
in the United States
had experimentally discovered
how to make
the impossible-to-compute
possible-to-compute.
I experimentally discovered that
it is possible
to solve extreme-scale problems
in computational physics
and solve them across
a massively parallel processing
machine
that I invented as a new internet.
My experimental breakthrough
laid the foundation for the **pre-cursor**
to the fastest computers of today.
My invention
is **embodied** in the fastest computers
that are now powered by

hundreds of processors. 1852
My invention
is **embodied** in the fastest
supercomputers
of today.
The fastest supercomputers
are powered by
millions of processors.
My invention
is the reason millions of school reports
have been written on the
contributions of **Philip Emeagwali**
to the development of the computer.
My invention
is the reason it is no longer said that
parallel processing
is a beautiful theory
that lacks experimental confirmation.

40.2.4 Theory Versus Experimental Confirmation

A theory is an idea
that is not positively true.
In the 1970s and '80s,
my quest was for the solution
of the toughest problem
of high-performance supercomputing.
In retrospect
and in the language of modern
supercomputing,
I define that grand challenge question
as this:
How can we use **10.65 million**
processors
and use them to invent
how to compress
10.65 million days, or 30,000 years,

of **time-to-solution** 1854
on one processor
to just one day
of **time-to-solution** across
a global network of
10.65 million
processors?
The news media,
such as the June 20, 1990 issue
of the *Wall Street Journal*
reported that I—**Philip Emeagwali**—
experimentally discovered
how to use a new internet
that I invented
as a new global network of
65,536
of commonly available processors,
or a global network of as many
identical computers.
I **invented**

how to use that new internet 1855
to compress 65,536 days, or 180 years,
of **time-to-solution**
on one processor.
I **invented**
how to compress that 180 years
of **time-to-solution**
to just one day
of **time-to-solution** across
a new global network of
65,536 processors
that was **my new internet**.
The reason my experimental
breakthrough
made the news headlines
in 1989
was that I metaphorically opened
65,536 doors to the unknown world
of parallel supercomputing.
That **invention**

was a quantum leap 1856
in **times-to-solution**
of sixteen orders of magnitude
and a speed increase of a factor of
two-raised-to-power-sixteen,
or a 65,536-fold increase.
That **invention**
opened doors
to the then **uncharted territory**
of massively parallel supercomputing.
That experimental breakthrough,
opened ten million
six hundred and forty-nine thousand
six hundred [10,649,600] doors
to the world's fastest supercomputer
of today
that is powered by
as many processors.
I contributed to the development

of the computer and the supercomputer 1857
by moving massively parallel processing
both forward and faster.
I **experimentally discovered**
how to move the modern supercomputer
forward
and how to do so by supercomputing
faster
and doing so by a factor of
65,536.
I moved the pre-cursor
of the modern supercomputer forward
and moved it
from the theoretical level of
“what if it can be done”
to the practical level of
“how to do it.”
Our ancestors used computing aids
that is faster than their toes
and fingers.

1858

Their computing aids
dates to the counting board
and the abacus
of the last four millennia.
Parallel processing
is the biggest paradigm shift
in the history of computing.
Parallel processing
is computing's defining
technical achievement.
Parallel processing
will be around
as long as the river flows
and the grass grows.
For my experimental discovery,
our distant descendants
will no longer ridicule parallel
processing
as a beautiful theory
that lacks experimental confirmation.

40.3 How I Turned Supercomputer Fiction to Non-Fiction

40.3.1 Turning Science Fiction to Non-Fiction

A science fiction writer
is a story teller
that solves the toughest scientific problems
and solves them
by merely waiving his pen
and declaring the **impossible-to-solve**
possible-to-solve.
In contrast, a computational mathematician
cannot solve the toughest problem
in calculus
by merely waiving his hand.
As a computational mathematician,
I can only discover the solution
to the toughest problem in calculus
and only discover it

if and only if
such a solution exists
but was previously not understood.
And I can only invent those things
that are possible to invent.
A science fiction writer
can write about cars
that run only on water.
In contrast, a scientist
must develop a prototype
of at least one car that he claims
only runs on water.
It's possible
for a science fiction writer
to write one hundred science fiction books.
In contrast, **it's impossible**
for a supercomputer scientist
to make two groundbreaking discoveries
in his lifetime.
It's impossible
for one supercomputer scientist

to invent the modern supercomputer 1861
that computes in parallel
and then, subsequently, invent
the hoped-for quantum supercomputer.
The invention of these supercomputers
demand big ideas,
billions of dollars,
and decades of hard work.
The parallel and quantum supercomputers
are each paradigm shifting
and each changed the way
we look at
the computer of tomorrow.
Nature does not give up its secrets
without a fight.
To parallel process
is to compute many things (or processes)
at once.
The technique of computing many things
at once
was known to the census board

that employed thousands of 1862
human computers
to execute billions of
arithmetical computations.
My contributions
to the development of the computer
was my experimental discovery
that supercomputers
powered by 64 binary thousand
commonly available processors
—each akin to a tiny computer—
can be harnessed
to solve many computation-intensive
problems
at once.
I visualized my new internet
as my new global network of
64 binary thousand
commonly available processors,
or a global network of
as many tiny identical computers.

I visualized that new internet 1863
as tightly encircling my room-sized globe.
I visualized my new internet
as two-raised-to-power sixteen,
or 65,536, tiny computers
that were equidistantly distributed
around the surface of a globe
that **I visualized**
in a sixteen-dimensional **hyperspace**.
Because **my visualization**
of my new internet
was new, and because the word “internet”
was not in my vocabulary
in the mid-1970s
I had to coin the term “HyperBall”
to describe the global network of
computers that I theorized.
In 1989, The Computer Society
of the Institute of Electrical
and Electronics Engineers (or IEEE)
issued a press release

that I—**Philip Emeagwali**—has achieved 1864
a technological breakthrough
in massively parallel processing
supercomputing.
That IEEE press release
had an impact because
the Institute of Electrical
and Electronics Engineers
was the world's largest
technical society.
In the May 1990 issue
of the academic journal
named "Software,"
The Computer Society of IEEE
described the economic benefits
of my experimental discovery
of massively parallel processing
and described it as follows:

[quote]
"The amount of money at stake

is staggering.
For example,
you can typically expect to recover
10 percent of a field's oil.” 1865

The Computer Society of IEEE
continued.

“If you can improve your production
schedule to get just 1 percent more oil,
you will increase your yield
by $400 million."
[end of quote]

That 1989 press release
issued by The Computer Society
that announced my
technological breakthrough
in massively parallel processing
supercomputing
as well as the companion articles

published by The Computer Society
in IEEE publications
led to cover stories
in many trade publications
and led to front page stories typically titled:
"African Supercomputer Wizard
Wins Top US Prize,"
and, in particular, it led to stories
on my contributions
to modern large-scale algebra
and new calculus
and my experimental discovery
of parallel processing
and to how I discovered
how and why parallel processing
makes modern computers **faster**
and makes the new supercomputer
the **fastest**,
namely, **the Philip Emeagwali formula**
that then United States President
Bill Clinton described

in his White House speech of August 26, 2000. 1867

My technological breakthrough
that opened the door
to the modern supercomputer
was my invention
of how to execute
floating-point arithmetical calculations
and execute them
at the **fastest** speeds ever recorded.
I visualized my experimental discovery
of the world's fastest calculations
as **occurring** across a new internet.
I visualized that new internet
as defined as a new global network of
65,536
commodity processors.
I invented
how to use that new internet
to send and receive emails
and do both at the fastest bandwidths

ever recorded. 1868
I invented
how to parallel program my new internet
that **I visualized**
as a new global network of
65,536, or 64 binary thousand,
tiny identical computers
and how to parallel program those
already-available processors
to communicate across
another new global network of
1,048,576, or one binary million,
regular and short
email wires
that were equal distances **apart**.
I mathematically and experimentally
invented
how to solve 64 binary thousand
initial-boundary value problems
of calculus and physics
and how to solve them **at once**

and how to email and solve them **across**
64 binary thousand tiny computers
that define a new internet
and how to compress
65,536 days, or 180 years,
of **time-to-solution** on one computer
and compress it
to only one day of **time-to-solution**
across a new internet
that is a new global network of
65,536 identical computers
and that is a supercomputer, *de facto*.
Because parallel processing
was then believed to be **impossible**,
it made the news headlines
when I **experimentally discovered**
that the im**possible-to-compute**
is **possible-to-compute**.
Yet, understanding how I made the
impossible possible
was not what made the news headlines

in 1989. 1870
What made the news headlines
was that I did the impossible.
I turned science fiction
into non-fiction.
In the old way of building computers,
one processor
is connected to one memory.
That processor
executes one instruction **at a time**.
In the new way of building computers,
that made the news headlines
when I discovered it to be faster
and did so on the Fourth of July 1989,
64 binary thousand
processors
are assigned
with a **one-to-one** correspondence
to solve 64 binary thousand
mathematical problems.
Those 64 binary thousand

processors
de facto synchronously executed
64 binary thousand instructions,
or as many floating-point
arithmetical operations,
at a time.
That invention
of the massively parallel processing
supercomputer
can be extended
to a billion commonly available processors
that encircle an internet
that is a seamless, cohesive supercomputer.
School reports on the contributions
of **Philip Emeagwali**
to the development of the computer
highlights his invention
of massively parallel processing
supercomputing.
My invention
made the news headlines

in 1989
because it heralded the end of the era of
vector processing supercomputers
that used only one **isolated**
vector processing unit.
That invention
of the massively parallel processing
supercomputer
represented real, measurable progress
in the development of the computer.

40.4 How I Crossed the Frontier of Supercomputing

40.4.1 The Uncharted Territory

Back in the 1940s,
the technology of the
massively parallel processing supercomputer
was an **uncharted territory.**

For four decades, 1873
the parallel processing supercomputer
only existed as a theory,
or as an idea not positively true.
It was a theory because
the fastest supercomputers
in the world
were not parallel processing **across**
several processors.
That theory
gave rise to the saying that
parallel processing
is a beautiful theory
that lacked experimental confirmation.
In the late 1940s,
the technology of parallel processing
was as distant
from the technology of the supercomputer
just as the planet Mars
is as distant from the planet Earth.
Technologically, to parallel compute

in the 1960s
was like the difference between
going to planet Mars
and reading a science fiction novel
on how **intergalactic space travelers**
travelled to Mars.
The world's fastest supercomputer
costs more than the spacecraft
that sent men to the moon.
Technologically, the supercomputer
that computed the trajectories
of the spacecraft that **travelled to the Moon**
is far more complex
than the spacecraft itself.
So programming a supercomputer
that's *de facto* a new internet
that's a global network of
65,536
processors
and programing those processors
as a lone wolf supercomputer programmer

and programming those processors
to cooperatively solve
an **initial-boundary value problem**
that is classified as the toughest problem
in calculus
that arose from
extreme-scale computational physics
was a grand challenge problem
that was only solved
two decades after the last man returned
from the Moon.
In the 1980s, twenty-five thousand
researchers
embarked on the epic quest
to discover
what will make the supercomputer
fastest.
In contrast, one science fiction writer
alone
can write about a fictional supercomputer
of Year Million

that is infinitely fast.
So discovering what makes
the supercomputer **fastest**
is far more difficult than
writing a piece of science fiction
on how our post-human descendants
of Year Million
could perform computations
that are **infinitely fast**.
A science fiction writer
can spend only a year
to write a bestselling fictional novel.
Her science fiction novel
could be on how fictional tourists
travelled to the planet Mars.
But that science fiction writer
cannot turn her fiction to fact.
She cannot write
the non-fiction analogue
of tourists traveling to the planet Mars.
She cannot write the factual analogue

simply because
we cannot yet travel to planet Mars.
The science fiction writer
is a storyteller
that told imaginative—but untrue—stories.
The science fiction writer
told stories about fictional heroes
that slayed fictional dragons.
But the factual quest
for the fastest supercomputer
takes the hero
—who must be a real life
supercomputer scientist—
and takes him into the **uncharted territory**
of massively parallel processing.
In that **uncharted territory**
the fruits from the tree of knowledge
is metaphorically guarded
by a fire-breathing dragon.
The hero's quest
is to slay the ferocious dragon

and his heroic deed
is to kill that dragon
and kill it with a short sword.

40.4.2 School Reports on Philip Emeagwali

I think of science fiction writings
as pink narratives.
I think of science writings
as blue narratives.
I think of school science report writings
as blue narratives.
The reason the science fiction writer
has the freedom
to tell a thousand lies
is that her lies are central to her fiction.
The reason the science **non-fiction** writer
does not have the freedom
to tell a single lie
is that one lie reduces his non-fiction

to fiction. 1879
The scientific discoverer
cannot tell a lie
because his experimental discovery
is the new truth,
or the new knowledge
about how our universe works.
That freedom to tell a lie
is denied to the scientist
because his **experimental discovery**
must be reproducible
to other experimental scientists
before it enters into science textbooks.
Since 1989,
millions of school science reports
have been written
on the **experimental discoveries**
of **Philip Emeagwali**.

40.4.3 What Makes Supercomputers Fastest

1880

I **invented**
how to solve the most extreme-scale
problems
in computational physics
and how to solve them in parallel,
or **across**
a new internet that I invented
as a new global network of
64 binary thousand
commonly available processors.
I discovered that
the root of the problem
was that **everybody else**
was talking
about parallel processing,
and **nobody else**
was trying to experimentally confirm
parallel processing.
It was like living in a nation of

1881

supercomputer scientists
where **everybody else**
was talking
about losing weight,
and **nobody else**
was trying to exercise
to lose weight.
The difference between
the theoretical discovery
and the experimental discovery
of massively parallel processing
is akin to the difference between
talking about losing weight
and doing something about losing weight.
A million theoretical papers
written onward of the 1940s
and on the topic of parallel processing
gave rise to the ridiculing saying
that parallel processing
is a beautiful theory that lacked
experimental confirmation.

1882

I did not **experimentally confirm**
massively parallel processing
by merely reading about parallel processing
and then writing about parallel processing.
I **experimentally confirmed**
parallel processing
by digging deep into all the
two-raised-to-power-sixteen, or 65,536,
already-available
processors
that outlined my new internet
that I invented
as a new global network of
sixteen times
two-raised-to-power sixteen,
or 1,048,576
regular and short email wires
that were equal distances **apart**.
In the 1980s, I—**Philip Emeagwali**—
commanded and controlled
more processors

than any person in history, 1883
both then and now.
As the lone wolf parallel processing
internet scientist
of the 1970s and ‘80s,
my grand challenge was to
dig deep into myself,
dig deep into my new internet,
and harness the magic
and the potential
of my new internet
and do so by inventing my new internet
as my global network of
processors
that were equal distances **apart**
and that were on the
fifteen-dimensional surface
of a globe
in a sixteen-dimensional universe.

40.4.4 Contributions of a Polymath

1884

I had to look beyond mathematics
and to become the polymath
that visualized the laws of physics
as having analogues
across different boards,
and having analogues
from the **story**board
to the **black**board
to the **mother**board
and across **mother**boards.
The blackboard is a limited medium
that limits the mathematician
to his partial differential equations.
The mathematician
cannot experimentally discover
a law of physics
and discover a law
on the blackboard alone.
The mathematical physicist

cannot experimentally discover 1885
the string theory of physics
and prove string theory
on the blackboard alone.
The theoretical supercomputer scientist
cannot invent
the massively parallel processing
supercomputer
and perform the fastest computation
on the blackboard alone.
The mathematician
cannot solve exactly
a partial differential equation
of calculus
that is defined at **infinite points**
in space and time
and, therefore, that requires
an **infinite number** of calculations
and solve that partial differential equation
with a **finite number** of
identical processors

and solve that partial differential equation
in **finite time**.
The exact solution
of an initial-boundary value problem
of calculus
that is governed
by a partial differential equation
transcends the blackboard.
That exact solution
is defined across four boards, namely,
the **story**board, the **black**board,
the **mother**board,
and across **mother**boards that I named
the **communication** board.
Because the grand challenge problem
transcended one board and one discipline,
I had to search for the fastest supercomputer
in the **shadows**
between physics and algebra.
I sought for the fastest supercomputer
in the **shadows**

between algebra and calculus. 1887
I sought for the fastest supercomputer
in the **shadows**
between calculus and large-scale algebra.
I sought for the fastest supercomputer
in the **shadows**
between algebra and arithmetic.
I sought for the fastest supercomputer
in the **shadows**
between arithmetic
and the processor.
And I sought for the fastest supercomputer
in the **shadows**
between the processor
and my new internet.
I invented that new internet
as a new global network of
65,536
commodity processors,
or a global network of
as many identical computers.

1888

40.4.5 Crossing the Frontier of Supercomputer Knowledge

The reason I looked beyond mathematics
was that I can be mathematically correct
and yet, lack the vision of the polymath,
that will enable me
to cross the frontier
of the then **uncharted territory**
of the pre-cursor
to the modern supercomputer
of today.
I had to cross that frontier
to conquer
the most vexing grand challenge problem
of massively parallel supercomputing.
Without that vision of a new internet,
I—**Philip Emeagwali**—will be standing at
the **uncharted territory**
of massively parallel processing

1889

supercomputer technology
and standing in that territory
like a magician without his magic.