

Analysis of Energy Consumption in Fabs

Report Submitted By

Unni Pillai

College of Nanoscale Science and Engineering

To

ISMI

12 November 2011

Abstract

The report presents an analysis of factors that drive energy consumption in semiconductor fabs. The analysis was based on data collected by ISMI during a survey of its member companies. Owing to the limitation on data available from the 2011 survey, the analysis was done on data pooled from two previous surveys (1997 and 2007) along with the 2011 survey. The average energy consumption per unit of production (measured as masks x cm²) in 2011 has decreased dramatically to 1/3rd of its value in 1997. This energy consumed by fab tools per unit of production in 2011 has halved and non-tool energy consumption has become 1/4th of their values in 1997. This decrease in energy consumed per unit of production points to sharp improvements in energy efficiency of tools and other fab equipment. Because data was available for only three years, it was not possible to do any time series statistical analysis on factors that led to the decrease in this energy consumption per unit of production. A cross-sectional approach pooling together the data from three years was taken to understand the factors that influence the per unit energy consumption. Wafer size, geometry, fab area and fab age were all correlated with energy consumed per UoP, but in a multiple regression geometry retains the most explanatory power and other variables become statistically insignificant. This indicates that newer equipment, as proxied by geometry, is driving most of the decrease in energy consumption per EU. Since lower geometries are also associated with higher production levels, it is possible that geometry is picking up the effect of any economies of scale in energy consumption. However, there is not enough variation in the data to identify the effects of production levels and geometry separately.

Introduction

This report presents the results from a project undertaken to understand the factors that drive energy consumption in semiconductor fabs. The project was done for International Semiconductor Manufacturing Initiative (ISMI). The analysis is based on the data that was collected by ISMI during a survey of some of its member companies. The survey was first undertaken in 1997, and was conducted again in 2003, 2007 and 2011. In the survey, the participating fabs are asked to provide information regarding the characteristics of the fab (location, age, area etc), details of the technology used in the fab (wafer diameter, geometry, average number of mask layers), annual production in the fab, and finally the annual energy consumption in the fab. An analysis of the different factors that influence energy consumption in a fab was made, and this report contains the results of the analysis.

The fabs in the survey vary in the scale of production, and consequently the annual energy consumption varies simply owing to differences in annual output at the fabs. In order to isolate meaningful differences across fabs that drive their energy consumption, it is clearly necessary to normalize the energy consumption by some common measure of output and compare these normalized energy consumption measures. However, finding a common measure of fab production is not without difficulties. Although all these fabs produce silicon chips, some of them produce logic chips, and others produce memory chips, with more sub-categories within these two broad product lines. Even with these sub-categories, products differ in terms of complexity, which would in turn drive differences in product quality and prices. Clearly, it would be incorrect to simply add up microprocessors, which often sell for hundreds of dollars per chip, with memory chips which might sell for a few dollars. Keeping this in mind, the ISMI survey collects data on the average number of mask layers used in production at each fab, which is a measure of complexity of production. In the latest survey, ISMI also collected some information on the number of “wafer moves” per year, another measure of complexity in production. Hence the survey data contains three variables that can be used as a Unit of Production (UoP),

- (i) cm^2 of silicon produced
- (ii) Masks x cm^2 of silicon produced
- (iii) Wafer moves x cm^2 of silicon produced

Though the total silicon area of the more complex chips are same (or lower) than their earlier less sophisticated counterparts, they contain considerably more sophisticated circuitry and are more valuable for customers. This increase in quality of output has to be taken into account while considering any changes in energy consumption per unit of output produced. Hence the incorporation of masks and moves in measuring output.

This study will focus on the second measure of UoP (mask x cm^2). The third measure is left out because of data reasons, only 5 fabs in 2011 (and none in the earlier years) contain data on this

measure. The first measure has less data available than the second measure. Also, as shown in the analysis section, including masks in UoP significantly increases the quantity of output produced (and hence decreases energy consumed per UoP) in the later years. Hence, considering just cm^2 of silicon produced clearly understates the improvements in energy efficiency. The second measure will be referred to as the Electrical Utilization Index (EUI) in the following sections. Hence

$$\text{Electrical Utilization Index} = \frac{\text{Total kWh consumed in a year}}{\text{Average number of masks} \times \text{cm}^2 \text{ of silicon produced}}$$

In what follows, we analyse how the EUI has varied across years, and across fabs. We then perform a statistical analysis to understand the major drivers of differences in EUI across time, and across fabs.

Analysis of Fab Energy Consumption

In this section, we summarize the observations about fab energy consumption that is seen in the data and present an analysis of factors driving the energy consumption.

1. Variation in Electrical Utilization Index (EUI) Across Fabs and Years

Pooling together the data from all three years, it is clear that there is substantial variation in Electrical Utilization Index across fabs. The least efficient fab consumed more than 5 times the energy of the most efficient fab to produce 1 UoP. The coefficient of variation (standard deviation/mean) is 0.45. Figure 1 below plots the energy consumption across fabs. The first two letters in the fab name indicate the year the observation was made.





Figure 1 –: Substantial variation in energy consumption across Fabs

It is clear from the above graph that there has been a decrease in Electrical Utilization Index (EUI) over the years, with most of observations in later years having lower energy consumption. Before exploring this time effect, note that the dispersion in EUI is present even if production is measured in terms of cm² of silicon produced by the fab, rather than in (masks x cm²).



The variation in EUI with cm² of silicon as the UoP is less than that with (masks x cm²) as the UoP. This is partly because the output in cm² of silicon is not available for some of the newer, more efficient Fabs. However, for the ones that both data are available, the decrease in energy per unit of production in later years is more pronounced when UoP is in (masks x cm²). This is clearly because the mask layers have increased over the years as the products have become more complicated, and highlights the fact that including mask layers in UoP is important to understanding changes in energy efficiency at Fabs.

2. Variation in Energy Consumption per UoP over Years

Energy consumed per UoP has come down dramatically from 1997 to 2011. Energy consumed by both fab process tools and other equipment has decreased, but the non-fab tool equipment has contributed more to the decrease in EUI than fab tools.

It would have been very informative to trace the change in energy consumption per UoP for the same plant over time. However, this is not possible with the current data since the identities of the plants are hidden in the survey results, and there is a change in the mix of fabs participating. So, we restrict attention to just the change in average energy consumption per UoP across years. This is shown in Figure 2 below.

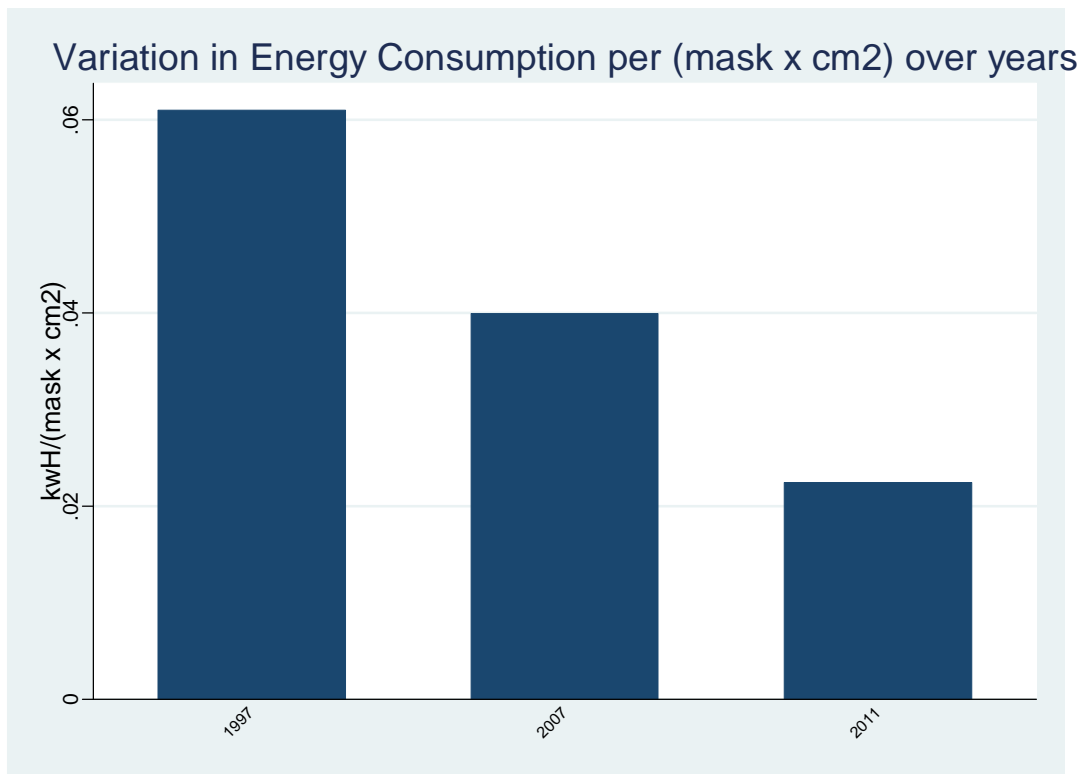


Figure 2 : Energy consumed per UoP has come down by a factor of 3.

Energy consumption per (mask x cm²) has become almost 1/3rd in 2011 of what it was in 1997. Clearly, this highlights that the industry has become much more energy efficient during the last 14 years. However, the energy consumption per cm² of silicon produced has remained roughly constant over the years, as is seen in Figure 3.

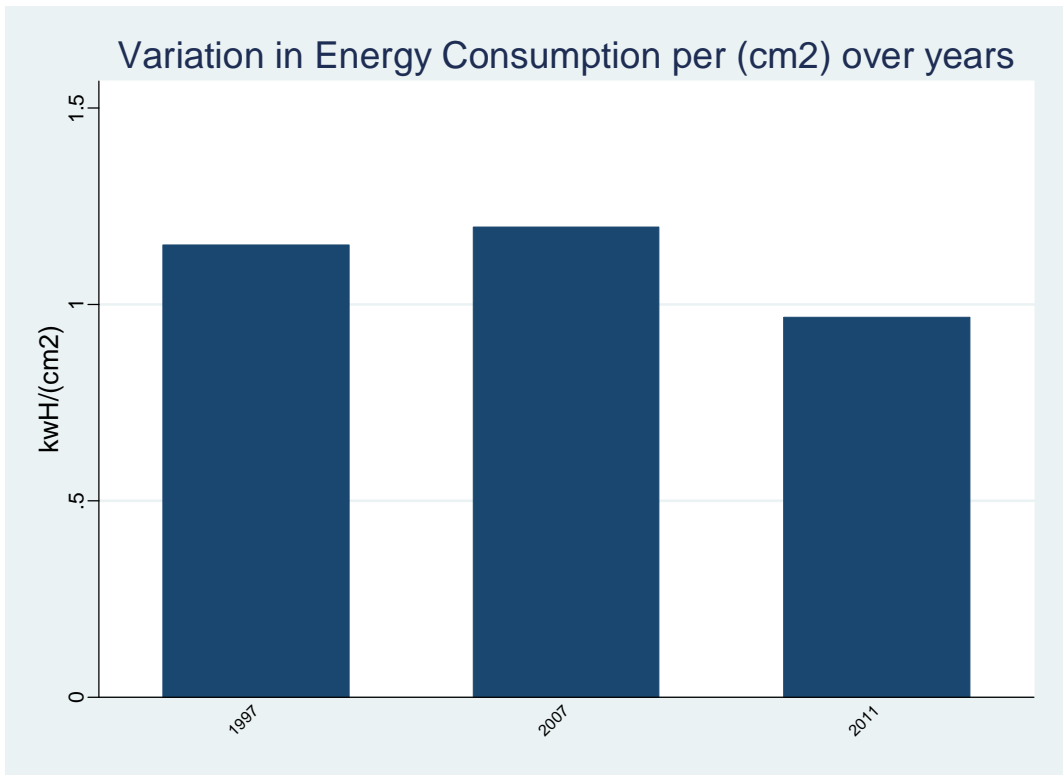


Figure 3: Energy consumption per cm² of silicon has remained roughly constant

This clearly indicates that the drop in EUI has been achieved by making more sophisticated chips, while keeping the energy cost of making a cm² roughly constant. Since the increased complexity of chips require more complex fab tools, one would like to know how the energy consumed by fab tools per UoP has changed over time. As shown in the graph below, the fab tool energy consumption per UoP has decreased by half from 1997 to 2011.

Variation in Tool Energy Consumption per (mask x cm²) over years





Figure 5 - The energy consumed by tools to produce a UoP has halved from 1997 to 2011.

However, when measured in terms of cm² of silicon output, the energy consumed by tools per UoP has slightly over years.



The non-tool energy consumption per (mask x cm²) has decreased more than tool energy consumption. The Non-tool related energy consumption has decreased by almost 1/4th from 1997 to 2011 (compared to 1/2 for tool energy consumption).

Variation in Non-Tool Energy Consumption per (mask x cm²) over years





Figure : Non-Tool Energy Consumption has decreased by 1/4th from 1997 to 2011

When the output is measure in cm² of silicon, the decrease in non-tool energy consumption per UoP is lesser, as shown in Figure -

Variation in Non-Tool Energy Consumption per cm² over years



From the above graphs it is clear that the tool energy consumption per cm^2 of silicon produced has increased slightly over the years, whereas the non-tool energy consumption per cm^2 of silicon produced has decreased over the years. Taking these two together, the total energy consumption per cm^2 of silicon has roughly remained constant over the years. However, when complexity of chips produced is taken into account using the mask layers as a proxy, then the total energy consumed per (mask $\times \text{cm}^2$) has fallen considerably over the years, with tool energy consumption falling to $1/3^{\text{rd}}$ and non-tool energy consumption falling to $1/4^{\text{th}}$ from 1997 to 2011. Hence the reduction in EUJ over the years, is clearly because of two factors

- (i) The ability of the improved tools to produce more complex chips while consuming roughly the same amount of energy as before.
- (ii) The significant reduction in non-tool energy per cm^2 of silicon over the years.

What could be the sources of this decrease in (i) and (ii)? There is only three data points available (corresponding to each year), which precludes any statistical analysis. There are two obvious candidate explanations for the decrease in EUJ over years.

First, there have been tremendous technological improvements in equipment used in fabs, both in fab tools and other cooling and accessorial equipment. The new fab tools are capable of etching smaller geometries and producing bigger wafers. But, they have also improved so that they are capable of producing more chips with less energy (and possibly less labor as well). These energy saving improvements might have come from improvements in throughput, or because of other specific energy saving technological innovations. It is beyond the scope of this report to analyze these other possible factors, since the survey contains no data on these other variables. But the new tools usually come out with smaller geometries, or bigger wafer sizes, or both. Hence one could use the geometry or wafer size to proxy for other energy innovations that have been developed over time. It should be clear that the argument is not that smaller geometries or bigger wafer sizes lead to lower energy consumption, but since energy-saving innovations are embedded in new generations of tools which usually come with smaller geometries or/and bigger wafer sizes, one can use geometry and wafer size to proxy for these other energy saving innovations.


Second, there could be the economies of scale in production. There is some fixed energy consumption involved in both running the fab tools as well as in running non-tool equipment like chillers and air recirculation equipment. For example, starting up these equipment may consume a lot of energy, which will be consumed no matter how many units are produced before the equipment is stopped. As the scale of production increases, this fixed energy consumption is divided over a larger number of output units and hence the energy consumed per unit would decrease as production increases. Also, many of the non-fabtool equipment like chillers, air circulators etc have volumetric nature of operations, implying that there are some physical economies of scale in production. In the next section, we explore these two, and other fab level variables, using individual fab level data.

3. Variation in Electrical Utilization Index within Year, across Fabs.

In this section, we explore in detail whether characteristics of fabs, characteristics of technology, or economies of scale have explanatory power or the differences in EUI seen across fabs. Each subsection looks at one variable.

(i) EUI and Wafer Diameter

The figure below shows the variation in EUI with wafer size, with each label standing in for the fab name.



The decrease to bigger wafers have been accompanied by lower EUI on average, although there are some exception. A simple regression (result below) shows that wafer size is a statistically significant explanatory variable, at a 1% level of significance.

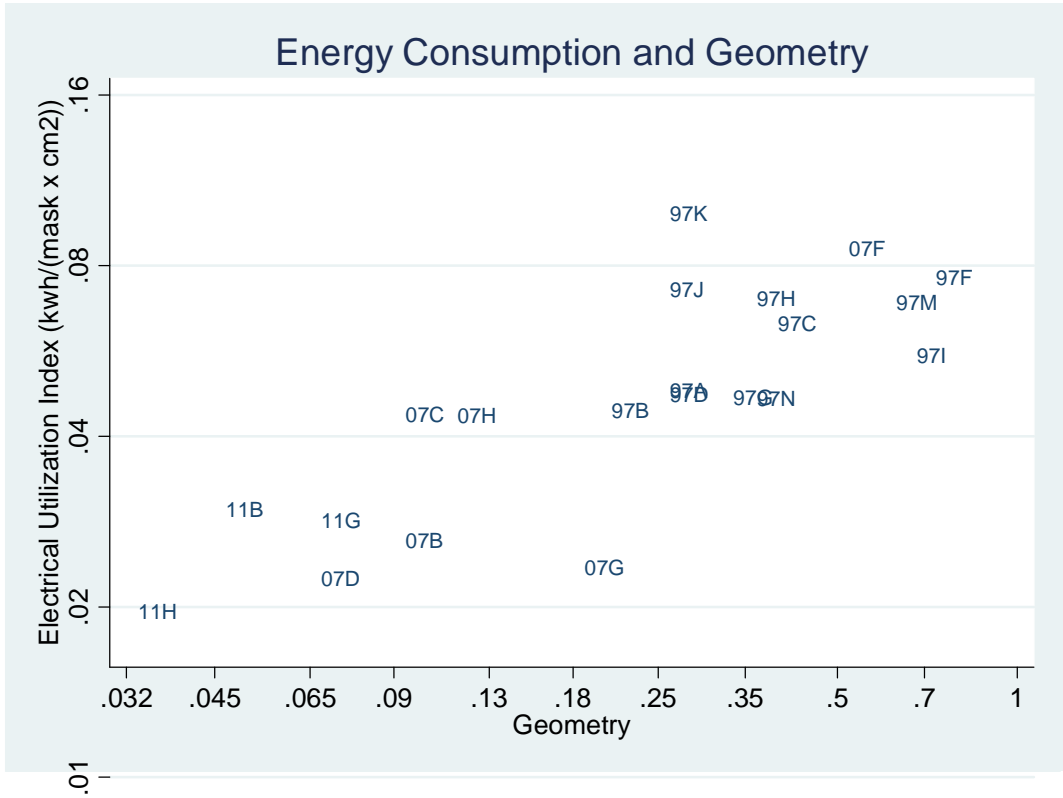
VARIABLES	ln(EUI)
Wafer Diameter	-0.172*** (0.0275)
Constant	-1.621*** (0.256)
Observations	27

R-squared 0.609

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

(ii) EUI and Geometry

The decrease in geometry has been accompanied by a clear decrease in EUI.



VARIABLES	ln(EUI)
ln(Geometry)	0.417*** (0.0705)
Constant	-2.405*** (0.130)
Observations	21
R-squared	0.649

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

As discussed before, it is not that a smaller geometry is by itself that is leading to smaller EUI. The geometries are standing in for other energy saving innovations that are accompanying new equipment. Since wafer sizes serve the same purpose, it is natural to ask whether geometry or wafer sizes are better proxies for energy savings that come with new equipment. The following regression shows that geometry has more explanatory power for changes in EUI than wafer sizes. In the regression below, $\ln(\text{EUI})$ was regressed on both $\ln(\text{geometry})$ and wafer diameter. As can be seen from the table, the coefficient on wafer diameter becomes insignificant, while the coefficient on geometry remains statistically significant (at the 5%) level, implying that geometry has more explanatory power than wafer diameter.

VARIABLES	$\ln(\text{EUI})$
$\ln(\text{Geometry})$	0.441** (0.179)
Wafer Diameter	0.0105 (0.0721)
Constant	-2.458*** (0.390)
Observations	21
R-squared	0.649


Another way of doing the above regression, is to do a fixed effects model with wafer diameter as the fixed effect. The fixed effect regression considers the observations within one wafer diameter as a group and finds out the effect of geometry on EUI for fabs with same wafer diameter.

VARIABLES	(1) $\ln(\text{EUI})$
$\ln(\text{Geometry})$	0.450** (0.181)
_Iwaf_dia_8	0.151 (0.225)
_Iwaf_dia_12	0.137 (0.448)
Constant	-2.470*** (0.180)
Observations	21
R-squared	0.661

As can be seen in the above table, geometry is significant even within the same wafer diameter. Since geometry has more explanatory power, we will omit wafer size and use geometry as stand for energy saving innovations introduced with new generations of equipment.

(iii) EUI and Output Level (Economies of Scale)

One problem with the analysis in (ii) above which suggests that newer generations of equipment have significant energy saving innovations is that the scale of production has also gone up alongside decreasing geometries. Therefore, the decrease in EUI might just be an effect of economies of scale, i.e the energy consumption decreases as scale of production goes up and this is erroneously attributed to decreasing geometries. As the graph below shows, there is a strong correlation between EUI and the scale of production.



.01

VARIABLES	(1) ln(EUI)
ln(uop_mask_cm2)	-0.445*** (0.0500)
Constant	0.465 (0.411)
Observations	27
R-squared	0.760

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

The above graph implies a relationship of the form

$$EUI = K UOP^{-0.44}.$$

Hence for every doubling of production (measured in mask x cm²), energy consumption per unit of production falls by around 26% . This can be simply because the underlying technology has economies of scale. It can also be because of learning effect, as companies produce more output they learn how to do so better and hence are able to reduce the energy consumed in production. Such relationships are also called learning curves and have been observed in a number of industries. Learning curves are often characterized by the how much the cost per unit (or in this case energy consumption per unit) falls when production decreases. This is often called the *progress ratio* and energy consumption graph above has a progress ration of 0.74).

But it is inaccurate to ascribe all of the reduction in EUI as simply coming from scale (or learning) effects. As we saw in (i) and (ii) the fabs in later years clearly had better technologies as well, in addition to higher scale. Hence one needs to control for technology changing factors like wafer size and geometry in order to understand the correct effect of scale. Since geometry was found to be a better proxy for changes in equipment than wafer size, we regress EUI on production level and geometry.


VARIABLES	(1) ln(EUI)
ln(uop_mask_cm2)	-0.290* (0.161)
ln(Geometry)	0.153 (0.161)
Constant	-0.529 (1.050)
Observations	21
R-squared	0.702

As can be seen geometry becomes insignificant, and the explanatory power of production level decreases as well. In fact, production level is significant only at a 10% level. With this level of significance it would be tenuous to state that production scale explains all the variation in EUI explained by geometry. One needs more data points, containing sufficient variation in in both production levels and geometry, to be able to say something more definitive.

(iv) Fab Age

It might be possible that the EUI in older fabs is higher because these fabs use older, less energy efficient equipment. There is a weak positive correlation, as can be seen in the graph below.

Variation of Energy Consumption per (mask x cm²) of Silicon with Fab Age



Regressing the EUI on fab age, after adding fixed effects for the years, we get a statistically significant (at 5% level) coefficient on fab age.

VARIABLES	(1) ln(EUI)
Fab Age	0.00154** (0.000643)
_Iyear_2007	-0.0260*** (0.00671)
_Iyear_2011	-0.0423*** (0.00791)
Constant	0.0507*** (0.00593)
Observations	27
R-squared	0.593

However, if our intuition is correct and the older fabs are less efficient because they use older technology, then adding a proxy for technology would decrease the explanatory power of fab age on EUI. This is indeed the case as can be seen from the regression results below.

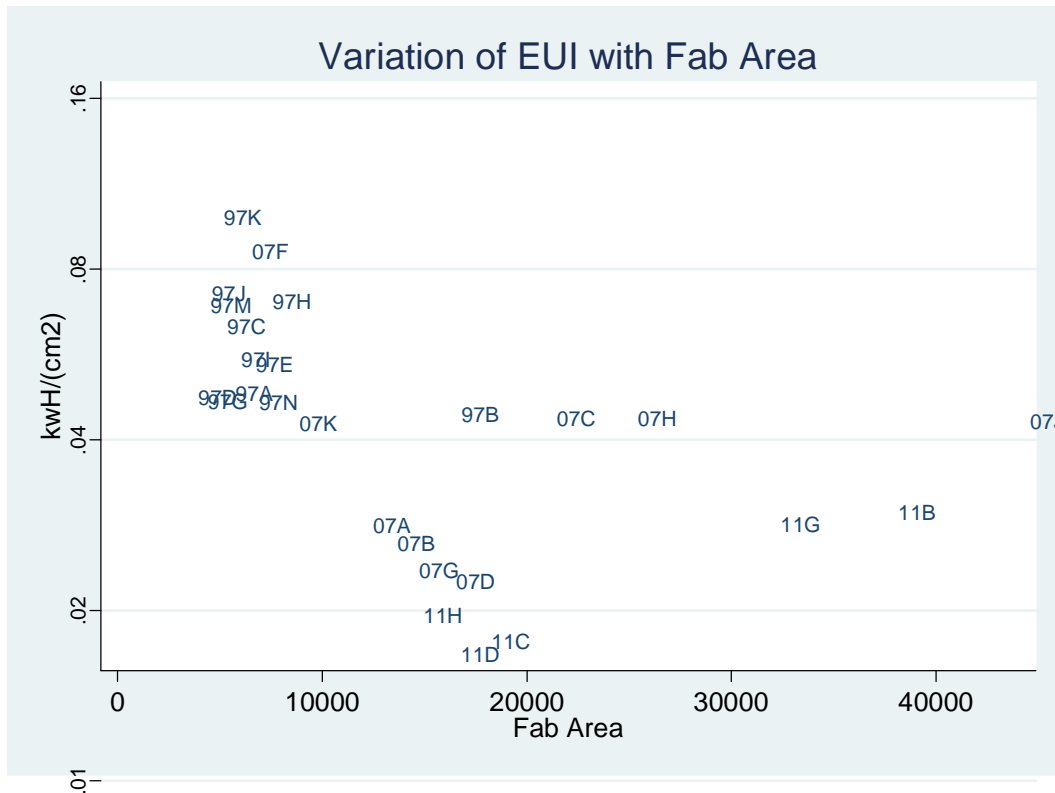
VARIABLES	(1) ln(EUI)
ln(Geometry)	0.404*** (0.0739)
Fab Age	0.0112 (0.0156)
Constant	-2.512*** (0.198)
Observations	21
R-squared	0.659

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

If geometry is put in the regression then Fab Age becomes statistically insignificant and geometry remains significant at the 1% level, indicating that the that fab age has on EUI is through new equipment.

(v) Fab Area

As the graph below shows, there is a weak negative correlation between EUI and Fab area. The regression result below confirms this, with a statistically significant (at 5%) level coefficient on fab area.



VARIABLES	(1) ln(EUI)
Fab Area	-2.02e-05** (8.16e-06)
Constant	-2.924*** (0.140)
Observations	26
R-squared	0.203

However, as with fab age it is necessary to check here whether the effect of fab area is really through new technology. The newer fabs are bigger, and they also have newer more energy efficient technology. To check this, we add geometry to the regression above. The result is below.

VARIABLES	(1) ln(EUI)
Fab Area	2.81e-06 (1.08e-05)
ln(Geometry)	0.444*** (0.124)
Constant	-2.393*** (0.151)
Observations	20
R-squared	0.627

As expected, when geometry is added the coefficient the effect of fab area becomes insignificant, while geometry remains highly statistically significant. Hence, the relationship between EUI and Fab Area is actually the effect of newer equipment used in the bigger (and newer) fabs.

Conclusion

The variation in the energy consumption per unit of production were examined in this study. The unit of production was chosen as (masks x cm²) because this was the UoP variable for which most data was available, and cm² of silicon was seen to be an insufficient measure of output.

There is significant variation in EUI over time, with average EUI decreasing in 2011 to 1/3rd of its value in 1997. This decrease in EUI over time is also accompanied by variability of EUI across fabs in a given year. An analysis of data pooled together indicated some of the variables that can explain part of this variation. Wafer size was found to be negatively correlated with EUI and geometry to be positively correlated. These indicate that newer equipment associated with bigger wafers and smaller geometries are also more energy efficient, with the wafer sizes and geometry possibly capturing the effect of other accompanying energy saving innovations. Once geometry is taken into account, the effect of wafer size on EUI becomes insignificant indicating that geometry better captures the energy saving innovations associated with new equipment than wafer size. Similarly, fab age and fab area are also correlated with EUI, with older fabs having higher EUI and bigger fabs having lower EUI (on average). However, once the effect of geometry is taken into account, fab area and fab age have insignificant explanatory power over EUI. Hence the older and bigger fabs have lower EUI because they are using newer equipment (as indicated by smaller geometries).

However smaller geometries are also associated with higher production levels and it is possible that geometry is incorrectly picking up the effect of lower EUI associated with higher production levels. A multiple regression however does not show convincingly whether one is more important than the other, and more data with sufficient variation in geometry and production level is needed to understand the relative importance of production scale and newer equipment on EUI.

There are many variables left out in this analysis which could possibly be affecting the EUI. The differences in the ambient temperature could be one, with fabs located in warmer environment probably requiring more energy for cooling than those in cooler environments. The type of product is another, with fabs producing some products consuming more energy than others, even with the same mask levels. The price of electricity is another missing variable, with fabs in areas with more expensive electricity possibly being more careful about controlling energy consumption. Adding these variables in the analysis, and including more data points would improve the analysis done in this project.